

The Exception Proves the Rule¹

RICHARD HOLTON, MIT

Abstract Legal rules admit of exceptions; indeed, it has been a legal maxim that one can infer the existence of a rule from exceptions that are made to it. Hart claims that the exceptions do not admit of exhaustive statement (a form of legal particularism) but that nonetheless rules can bind. This paper develops a logical framework which accommodates this position, shows that it is available to a positivist, elucidates the role of rules within it, and concludes by discussing the relevance to issues of judicial discretion.

When faced with a rule that they take to be true, and a recalcitrant example, people are apt to say: “The exception proves the rule”. When pressed on what they mean by this though, things are often less than clear. A common response is to dredge up some once-heard etymology: ‘proves’ here, it is often said, means ‘tests’.

But this response—its frequent appearance even in some reference works notwithstanding²—makes no sense of the way in which the expression is used. To insist that the exception proves the rule is to insist that whilst this is an exception, the rule still stands; and furthermore, that, rather than undermining the rule, the exception serves to confirm it. This second claim may seem paradoxical, but it should not, once it is realized that what does the confirming is not the exception itself, but rather the fact that we judge it to be an exception; and that what is confirmed is not the rule itself, but rather the fact that we judge it to be a rule. To treat something as an exception is not to treat it as a counterexample that refutes the existence of the rule. Rather it is to treat it as special, and so to concede the rule from which it is excepted. The point comes clearly in the original (probably 17th Century) Latin form:

Exceptio probat (figit³) regulam in casibus non exceptis.

Exception (i.e. the act of excepting) proves (establishes) the rule in the cases not excepted.

¹ Versions of this paper have been given at MIT, the University of Texas at Austin, the University of Oxford, the University of Birmingham, and at the Oslo Conference on Language and Law, where Scott Soames was the commentator. Many thanks to the audiences there, and to the referees for the *Journal of Political Philosophy* for some very helpful comments.

² For example K. G. Wilson, *Columbia Guide to Standard American English* (New York: Columbia University Press, 1993); *Hutchinson Encyclopedia*; (Oxford: Helicon, 1999). I don’t know how far back the mistake goes. It was made by Ambrose Bierce in *The Devil’s Dictionary* in an entry that would have been written in the late 19th Century. The confusion has yet to abate. After citing the very clear discussion in the *OED*, Robert Burchfield in his revised edition of *Fowlers* writes ‘In the context of the proverb, *proves* means “tests the genuineness or qualities of”, no more, no less’ *Fowler’s Modern English Usage*, Third Revised Edition (Oxford: Clarendon Press, 1998) p. 273.

³ Whilst ‘probat’ maintains the ambiguity of the English ‘prove’, which can mean ‘test’ or ‘establish’, the first citation given in the *OED*, from Samuel Collins in 1617, uses ‘figit’, which can only mean the latter.

Clearly this form of reasoning cannot apply when the rule that we are considering has the form of a simple universal generalization. Here there can be no exceptions, only counterexamples. So what we need, and what will be developed here, is some notion of a defeasible rule, one to which exceptions can be made without rejection.

However, even with such a notion in place, our classification of something as an exception to a rule is not actually evidence that the rule exists; it is merely evidence that we believe that it does. The further step can only be taken in cases where to concede that we accept a rule is to concede (or to come very close to conceding) that it exists. The obvious example is legal argument, and it is here, in an argument from Cicero, that the maxim appears to have its source.⁴

Lucius Cornelius Balbus, later to become Caesar's secretary and the first non-Italian consul, had been granted Roman citizenship (he was from Gades, i.e. present day Cadiz). Increasingly powerful and widely unpopular, his citizenship was challenged in the courts. Among many arguments, it appears that the prosecutor contended that, since treaties concluded with various peoples included clauses prohibiting the granting of Roman citizenship to their people, Balbus's citizenship was should also be judged illegal by a generalization of that approach, even though no such clause was present in the treaty with Gades.⁵ In response, Cicero successfully contended, that, since these prohibitions were clearly *exceptions* to the general policy, it could be inferred that the general policy allowed foreigners to be citizens: "if an excepting clause makes it impermissible, where there is no excepting clause, then it is necessary that it is permissible."⁶

ENGLISH LAW

Fowler observes, in a clear discussion of the maxim, that 'the value ... in interpreting statutes is plain'.⁷ It was certainly in use in English law by the 17th Century.⁸ By the 18th Century Lord Kenyon was using it with what was clearly its Ciceronian force, here arguing that the general prohibition on examinations is not undermined, indeed is supported, by the existence of acknowledged exceptions:

⁴ I owe the reference to Mark Israel's very useful discussion on the alt-usage-english.org newsgroup at <http://alt-usage-english.org/excerpts/fixtheexc.html>

⁵ The prosecutor's speech is not preserved, so we have to infer it from Cicero's response.

⁶ 'Quod si exceptio facit ne liceat, ubi non sit exceptum, ibi necesse est licere.' Cicero, *Pro Balbo* xiv 32; translation from Kimberly Anne Barber, *Rhetoric in Cicero's Pro Balbo*, (New York: Routledge, 2004) p. 56

⁷ H. W. Fowler *Modern English Usage* (Oxford: Clarendon Press, 1926). Fowler gives no examples; in the 2nd edition (1965) Gowers added the reference to Atkins mentioned below.

⁸ The earliest instance cited in Broom's *Legal Maxims* (8th U.S. edition, Philadelphia: Johnston, 1882) is from Edward Coke:

Know reader, that where it is said in this case, that a writ of error lies not upon an award, till the principal judgment is given; and where it is also said, that no writ of error lies till the whole matter in the original is determined; both these rules are regularly true; but yet each of them has exceptions; ... in this case, *Exceptio probat regulam, & sic de similibus*. (Coke *Institutes of the Laws of England* (1628) 11, 41. cited in Broom's *Legal Maxims*, xx.)

It has been said that there are cases where examinations are admitted, namely before the coroner, and before magistrates in cases of felony. That observation appears to me to go in support of the general law, rather than against it. Every exception that can be accounted for is so much a confirmation of the rule that it has become a maxim, *exceptio probat regulam*. Those exceptions alluded to are founded in the statutes of William and Mary; and that they go no further is abundantly proved.⁹

He also raises a note of caution, denying that an exception (again in a statute) should *always* be taken as evidence of a rule that would otherwise cover the case at hand:

As to the argument *exceptio probat regulam*; it seems to me that the anxiety of some members of the House induced them to insert the last clause, after the act was first drawn; but I think that the first section could never have been extended to the cases mentioned in the last, if they had not been excepted.¹⁰

Given that this last passage is generally (and reasonably) taken as precedent for the legal status of the maxim¹¹ we have a nice case of self-exemplification: it is exactly because Kenyon is denying the applicability of the maxim in this instance that we have reason to believe that it generally holds.¹²

Other uses of the maxim are more modest: it is invoked not to show that exceptions prove a rule, but that they are consistent with it. As we might more accurately put it: ‘The exception does not disprove the rule’. Indeed, the idea that one can find exceptions to a rule without giving up on the rule is central to the practice of common law. The origin of common law is not in statute but in precedent. The since the 19th Century the system has worked under the doctrine of *stare decisis*: the rulings that courts give are constrained by the rulings that earlier courts have given. Sometimes this precedent is, to all intents and purposes, binding: a lower court will normally treat as binding the decisions of higher courts within the

⁹ King v The Inhabitants of Eriswell, 3 T.R 722 [1790] per Lord Kenyon

¹⁰ Crespigny v Wittenoom 4 TR 793 [1792] per Lord Kenyon

¹¹ It is cited in Broom’s.

¹² It might seem that Lord Atkin is making the same point as Kenyon in saying that rules are not proved by exceptions ‘unless the exceptions themselves lead one to infer a rule’ (this is the line that Fowler quotes). However, when one sees the context of the line it is rather less clear what is going on. In discussing various judgments that raise the idea that a class of contracts may be voided if they have a general tendency to harm the public interest he writes:

I think that the substance of these judgements is that you must have a general rule, a general tendency to wrong to which there may be exceptions. But if the contract has to be applied to an act of social or other relations in which there will be generally no tendency to do wrong, though there may be exceptions, the contract will not be avoided. A rule is not proved by exceptions unless the exceptions themselves lead one to infer a rule. Lord Atkin in *Fender v St. John-Mildmay* [1938] AC, 14.

Other 20th Century Law Lords use the maxim in the more usual way. See Rodriguez v. Speyer Bros. [1919] AC 102 per Viscount Haldane: ‘... the exception proves the rule. It does not, as has, in effect been contended, destroy the rule.’

same jurisdiction. In other cases the precedent will be merely persuasive: for instance if it comes from a lower court, or from a court within another jurisdiction, or if the precedent in question was not actually used in reaching the earlier decision, but was merely an aside (that is, if it was not part of the *ratio decidendi*, but was merely *obiter dicta*). It is a contested question, to which we shall return, whether a court's earlier decisions can bind itself.

What is important at this point, however, is the fact that even apparently binding precedent can be avoided if the court can *distinguish* the case before it from the earlier cases. When a court distinguishes a case in this way, it does not overturn the precedent. Rather it points to some feature of the case at hand that makes it different, and on the basis of which a different decision can be given. In short the new case is treated as an exception to the rule determined by the precedent. So we need an account of rules that can accommodate exceptions in this way.

That is not the only option. An alternative is to give up on the idea that the common law works by means of rules at all, and a number of theorists have concluded just that.¹³ But such a view is not terribly attractive. Lawyers explaining those areas of law that are still basically governed by a common law approach—for instance, contract and torts, and, in England at least, large parts of the criminal law—will, in most cases, present their clients with a set of rules of what is and is not allowed: of how liability falls, of the situations under which a contract will be void, and so on. Equally, when courts develop a body of rulings that interpret statutes, these lend themselves to presentation by means of rules. And it is important, for familiar reasons of transparency and reliance, that this is so.¹⁴ Now it may well be that typically these rules are spelled out by the writers of legal texts rather than by judges.¹⁵ Nonetheless, the explicit formulations will come to be influential on judges as well as lawyers; and for large parts of the common law there is widespread agreement on what they are. So we should at least try to develop an account of rules that allows exceptions.

ACCOMMODATING EXCEPTIONS WITHIN RULES

Hart writes:

We promise to visit a friend the next day. When the day comes it turns out that keeping the promise would involve neglecting someone dangerously ill. The fact that this is accepted as an adequate reason for not keeping the promise surely does not mean that there is no rule

¹³ For instance, A.W.B. Simpson 'The Common Law and Legal Theory' in W. Twining ed *Legal Theory and Common Law* (Oxford: Blackwell, 1986)

¹⁴ On the importance of such considerations, see Larry Alexander and Emily Sherwin *The Rule of Rules* (Durham: Duke University Press, 2001). I think that the approach I advocate meets most of their desiderata, even though it does not construe rules in the strict way that they envisages. Clearly, as many have remarked, the law has to trade off between predictability and fairness.

¹⁵ See Neil Duxbury *The Nature and Authority of Precedent* (Cambridge: Cambridge University Press 2008) p. 23.

requiring promises to be kept, only a certain regularity in keeping them. It does not follow from the fact that rules have exceptions incapable of exhaustive statement, that in every situation we are left to our discretion and are never bound to keep a promise. A rule that ends with the word ‘unless ...’ is still a rule.¹⁶

So Hart at least implicitly accepts that some valid rules have exceptions that cannot be exhaustively stated. His comments come after his famous discussion of the open texture of law that results from the open texture of language—Is a child’s bicycle covered by a regulation that prohibits vehicles from the park?—so one might think that the ideas here can be similarly explained. But clearly they cannot. Whilst there may be vagueness in the idea of a promise—Is it a promise if made under duress, or if the promisor does not understand what they are committing themselves to, or if the promisee is not aware of it?—such vagueness is not relevant here. There may be no doubt that my promise was as clear and central an example as one is ever likely to find, and yet it still be true that in the circumstances I am not bound to keep it. So we need a different explanation.

One complicating factor is that Hart’s discussion here concerns moral rules and obligations, rather than legal. This is surprising, since Hart has been talking about legal rules, and it is far from obvious that what holds for moral rules also holds for legal. We’ll shortly return to the question of whether the two should be treated in the same way; for now let us follow Hart and keep our focus on the moral.

Hart’s example might put us in mind of recent discussions of moral particularism. Those who advocate such an approach contend that any putative moral rule is subject to exceptions; and as a result they tend to reject any role for rules. Take any rule that links the moral to the descriptive, they say, and we can find an exception to it; amend the rule to embrace the exception and we can find an exception to the amended rule, and so on.¹⁷

I am sympathetic to the drift of the position, but the as characterized it strikes me as too strong on several counts. First, how could we be sure that *every* moral rule is subject to exception? Our best grounds are inductive: that despite much effort, every rule that has been proposed has been shown vulnerable to exception. But it is a long and dogmatic step from that to the insistence that there are no exceptionless rules. A more plausible claim is a more modest one. Particularism should be couched in terms of its opposition to the theorist on the other side who insists that there are exceptionless rules. Given the failure of previous attempts to find them, the insistence on the existence of such rules is the piece of dogmatism. The particularist position should then be characterized as scepticism about that position, coupled with the positive claim that in the light of such scepticism we

¹⁶ H.L.A. Hart, *Concept of Law* (Oxford: Clarendon Press, 1961) p. 136

¹⁷ Jonathan Dancy, ‘Ethical Particularism and Morally Relevant Properties’ *Mind* 92 (1983) 530–47; *Ethics Without Principles* (Oxford: Clarendon Press 2004).

should not construct our moral theories on the assumption that exceptionless rules are there to be found.

Second, it does not follow from the supposition that there are no exceptionless moral rules that rules have no role to play. As we have seen, Hart suggests a model: perhaps rules end with an (often unstated) *unless*-clause; where this is not triggered, the rule applies. Hart, however, does not tell us how the *unless*-clause is to be completed. And here it might seem that he is faced with a dilemma. On one approach the *unless*-clause contains a full statement of all the factors that would defeat the rule; but that is clearly incompatible with the idea that the exceptions are incapable of exhaustive statement. On the other approach the *unless*-clause would be open ended: ‘One should keep one’s promises unless there is reason not to’. But that is clearly trivial. If we are to give substance to an open-ended *unless*-clause we need to find a middle way between these two approaches. I think that something is available for moral principles. In the next section I’ll present this account. Then I shall turn to the question of whether it is applicable to the law.

PRINCIPLED PARTICULARISM—MORALS

The intuitive idea that I shall work with is that a moral principle can be over-ruled if there is a moral justification for the exception.¹⁸ In Hart’s example there are moral grounds for tending to the seriously ill person rather than keeping my promise. But if so, then there is plausibly a moral principle that tells us that there are such grounds. So the *unless*-clause can be read as quantifying over other moral principles. It says that the moral principle will apply to the case unless there are other moral principles that apply to that case that render the verdict of the first principle wrong. That is: a moral principle like “Killing is wrong” applies to a case of killing if and only if there are no other moral principles—for instance, “Killing done in self-defence is not wrong”—that apply to the case and render the verdict of the first principle wrong. But the *unless*-clause doesn’t list all of the possible further principles that would defeat the application of the initial principle. That would be impossible if they are incapable of exhaustive statement as Hart supposes. It simply quantifies over them.

The crucial thought here is that what makes a moral principle apply to a case isn’t just what obtains; it is also what doesn’t obtain. So as well as adding an *unless*-clause to the moral principles, the full form of a moral argument will also require the addition of a premise to the effect that the *unless*-clause is not triggered. We can make these ideas more precise by defining a notion of what it is for one set of considerations to be *superseded* by a second; that is, for the second set to allude to some further consideration that would upset the conclusion that one would reach

¹⁸ I here follow the account that I gave in ‘Principles and Particularisms’ *Proceedings of the Aristotelian Society Supplementary Volume 67*, (2002) pp. 191-209. Readers wanting more details of the account, and further discussion of difficulties it faces, should consult that article.

on the basis of the first. Then we can construct the *unless*-clause, which I call ‘*That’s it*’, in terms of that notion of supersession.

I start with a first definition of supersession, one which will have to be modified, but which will serve to fix the main idea:

Supersession (first attempt)

Suppose we have a set of predicates $\{F_1, F_2 \dots F_m\}$; and suppose that these occur in a set of sentences $\{F_{1a}, F_{2a}, \dots F_{ma}\}$ and in a corresponding moral rule of the form $\forall x ((F_{1x} \& F_{2x} \& \dots \& F_{mx}) \rightarrow F_vx)$, where F_v is a predicate expressing a moral verdict. Then we say that those sentences and that moral rule are *superseded* by another set of sentences $\{G_{1a}, G_{2a}, \dots G_{na}\}$ and a corresponding moral rule $\forall x ((G_{1x} \& G_{2x} \& \dots \& G_{nx}) \rightarrow G_vx)$ if and only if:

- (i) $(G_{1x} \& G_{2x} \& \dots \& G_{nx})$ entails $(F_{1x} \& F_{2x} \& \dots \& F_{mx})$, but not *vice versa*;
- (ii) F_vx is incompatible with G_vx .

The first clause here requires that the second set of set of sentences says everything that is said by the first *and* something more; the second clause requires that the second moral rule brings one to a verdict that is incompatible with that of the first. For instance, the sentence ‘This was a killing’, and the moral rule ‘If x is a killing then x is wrong’ are superseded by the sentences ‘This was a killing’ and ‘This was done in self defence’ and the rule ‘If x is a killing and x is done in self defence then x is not wrong’. Now we can go on to define the *unless*-clause, making use of the idea of supersession that we have defined:

That’s it: There are no further relevant moral rules and facts; i.e. there is no true moral rule and set of *true* sentences that supersede those that appear in this argument.

What we want to say is that if *That’s it* holds, the moral rule holds. So we need to add *That’s it* to the rule, as a further conjunct of the antecedent: ‘If x is a killing, and *That’s it*, then x is wrong’. Then the arguments in which the rule occurs need a further premise saying that *That’s it* does indeed obtain.

So we get arguments like this:

- | | | |
|----|----|---|
| 1M | P1 | This is a killing |
| | P2 | $\forall x ((x \text{ is a killing} \& \textit{That's it}) \rightarrow \text{you may not do } x)$ |
| | P3 | <i>That’s it</i> |
| | C | You may not do this. |

Recall that Hart claimed that the *unless*-clause cannot be exhaustively stated: no matter how many exceptions are give to a rule, one can always imagine further

exceptions that have not been captured. Equivalently, no matter how much is built into the content of the rule itself, one can always imagine further factors that will render the rule invalid. In our current framework, this can now be understood as the claim that any moral argument like M1 is bound to be superseded by other valid arguments: take any moral argument, we can always find another that supersedes it.¹⁹ Or, at least, that is what we want to say; unfortunately we can't quite say it yet, since the result of inserting the *That's it* clause into the arguments is that the original definition of supersession no longer applies to them. So let us stop to fix that up by redefining supersession:

Supersession (second attempt)

A set of sentences $\{F_{1a}, F_{2a}, \dots, F_{ma}\}$ and a corresponding moral rule $\forall x ((F_{1x} \& F_{2x} \& \dots \& F_{mx} \& \textit{That's it}) \rightarrow F_x)$ are *superseded* by another set of sentences $\{G_{1a}, G_{2a}, \dots, G_{na}\}$ and a corresponding moral rule $\forall x ((G_{1x} \& G_{2x} \& \dots \& G_{nx} \& \textit{That's it}) \rightarrow G_x)$ if and only if:

- (i) $(G_{1x} \& G_{2x} \& \dots \& G_{nx})$ entails $(F_{1x} \& F_{2x} \& \dots \& F_{mx})$, but not *vice versa*;
- (ii) F_x is incompatible with G_x .²⁰

Given this revised definition we can say that 1M is superseded by the valid argument:

- | | | |
|----|----|--|
| 2M | P1 | This is a killing |
| | P2 | This is done in self defence |
| | P2 | $\forall x ((x \text{ is a killing} \& x \text{ is done in self defence} \& \textit{That's it}) \rightarrow \text{you may do } x)$ |
| | P3 | <i>That's it</i> |
| | C | You may do this. |

Similarly 2M would be superseded by an argument that added the claim that the killing was not necessary for the defence, and that in turn would be superseded by one that added that the defendant didn't realize this to be so, and that by one that he could have known it had he only paid due care, and so on. If Hart is right, no matter how complicated the rule gets, we will always be able to think of an argument that supersedes it. But the fact that every moral argument is superseded by some *valid* argument does not mean that it is superseded by a *sound* argument, that is, by a valid argument that has *true* premises. If the killing was not done in self-defence, and there is equally no other excusing condition, then the *That's it* premise in the original argument 1M will be true, 1M will not be superseded by any sound argument, and the conclusion, that the killing was impermissible, will be simply true. Put another way: the fact that every moral argument would be

¹⁹ I say that one *argument* supersedes a second iff the premises of the first argument, with the *That's it* content removed (both the premise, and the conjunct in the rule), supersede the premises of the second argument, with its *That's it* content removed. Since *That's it* includes a demonstrative, it changes its sense when it occurs in two different arguments, so we cannot define a notion of supersession that includes the *That's it* material.

²⁰ One consequence of this is that the definitions of supersession and *That's it* are now circular, in the sense that each makes reference to the other. I don't think that this is pernicious. For discussion see 'Principles and Particularisms' p. 200.

superceded were certain facts to obtain does not show that every moral argument is in fact superceded. A good moral argument is one that is not.

So the approach meets one of the desiderata with which we started: we have found a way to interpret the *unless*-clause that does not involve a simple list. The approach also meets the second, for clearly it does not lead to triviality. The *That's it* condition is a substantial one. Many real moral arguments go wrong exactly because it is not met. Very often when we reach a mistaken moral conclusion it is because there is some further relevant factor that we are overlooking.

PRINCIPLED PARTICULARISM—THE LAW

Can we take a similar approach in the legal case? Formally this is straightforward. We need a parallel account of supersession, defined as we defined it in the moral case, except that talk of moral rules is replaced with talk of legal rules. Thus, for instance, the legal rule ' $\forall x$ If x dishonestly appropriates property belonging to another with the intention of permanently depriving the other of it, and *That's it*, then x is guilty of theft' and the sentence 'A dishonestly appropriated property belonging to another with the intention of permanently depriving the other of it' are superseded by the legal rule ' $\forall x$ If x dishonestly appropriates property belonging to another with the intention of permanently depriving the other of it, and x is coerced, and *That's it*, then x is not guilty of theft' and the sentences 'A dishonestly appropriated property belonging to another with the intention of permanently depriving the other of it' and 'A was coerced'.

Then we need a version of *That's it* that applies to legal rules:

That's it: There are no further relevant facts and legal rules; i.e. there is no set of true sentences that, together with a legal rule, would supersede those that appear in this argument

Now we get arguments like this:

- 1L P1 Jones dishonestly appropriated property belonging to another with the intention of permanently depriving the other of it
- P2 $\forall x$ ((If x dishonestly appropriated property belonging to another with the intention of permanently depriving the other of it & *That's it*) $\rightarrow x$ is guilty of theft)
- P3 *That's it*
- C Jones is guilty of theft

The claim is that this argument could be superseded by other *valid* arguments, for instance

- 2L P1 Jones dishonestly appropriated property belonging to another with the intention of permanently depriving the other of it
- P2 Jones acted under duress

- P3 $\forall x$ ((If x dishonestly appropriated property belonging to another with the intention of permanently depriving the other of it & x acted under duress & *That's it*) \rightarrow x is not guilty of theft)
- P4 *That's it*
- C Jones is not guilty of theft

But again the fact that every legal argument can be superseded by valid arguments does not mean that every legal argument is superseded by sound arguments. A good legal argument in a particular judgment is one in which the *That's it* clause is true.²¹

DISANALOGIES?

It is easy enough then to construct an account of legal argument that is formally analogous to that I gave for moral arguments. What is harder is to see whether the substance is the same. There is one glaring difference.²² Moral particularism is typically associated with a form of objectivism about morals. On this approach it is exactly because morality is independent of us that it is uncodifiable. Though the picture is probably not compulsory, it is natural to imagine the moral truths there, ranged in infinite ranks ready to supersede any codification that we attempt.

There are views on which the law is similar: natural law theories that treat the law as not exhausted by the practices of the courts. But such views are contentious, and they are clearly not available to positivists like Hart, who see law as, in some sense, an artifact. Can the account be made compatible with a positivist approach?

The central issue is how we should understand the reference to a 'legal rule' in the definition of supersession and the *That's it* clause. It might seem that to a positivist like Hart the list of legitimate legal rules must be finite. After all, they are created as a result of human actions, and there has been a finite number of them. So if each action can itself create only a finite number of rules, then the number of rules must be finite. And one might think that this means our *unless*-clause will not be uncodifiable after all, and it will not be true that each legal argument will be potentially superseded.

One response here would be to soften the positivism. We might, for instance, broaden the account of what the law is, so as to include not just the extant sources, but also what follows from those sources, or what provides the best justification for those sources. Alternatively, leaving the law as it is, we might want to soften the

²¹ Point of clarification: I am primarily talking here about exceptions that involve a change to the legal rule defining what is necessary for an offence. There are other cases in which an exception will involve conceding that the law has been broken, but will provide a full or partial defence given the nature of the action. I would suggest treating these in like way, although here the legal rule defining the offence would not be qualified; the qualification would rather apply to the further rule defining the appropriate treatment of the defendant. My thanks to a referee for this point.

²² Perhaps there are other differences that tell in favour of the account for law rather than for morals. Jonathan Dancy objects that the account can make no sense of contributory reasons: *Ethics Without Principles*, 27–9. I'm not sure what to make of that worry in the moral case; in the legal case, given the role I assign to legal rules, I doubt that we need any such notion.

definition of supersession and the *That's it* clause, so that the relevant legal rules wouldn't be the extant rules, but rather those that would be included in a legitimate extension of the extant rules, where again this might make use of some notion of inference from, or justification of, the extant rules.

There are things to be said for and against these approaches.²³ I shall not, however, pursue the matter here, for neither, I think, is necessary. That is, there is an important sense in which we can think of the law as being uncodifiable even if we stick with a restrictive, positivist, source-based account of what the law is, and think that the definition of supersession and the *That's it* clause quantifies over only such law. The trick involves thinking of the law *from the judge's perspective*.

When a judge decides a case, the law, at the moment of deciding that case, will include all of the previous law, *and* the very decision that the judge is making. If the judge has distinguished the case from those that went before, the decision will itself constitute a new legal rule. Suppose, to take an absurdly simplified example, that the law on theft had never confronted a case of coerced theft, so that every legal decision looked like 1L. Then suppose a judge did confront such a case, and distinguished it from those covered by the existing legal rule to produce an argument like 2L. It would then be true at the time that the decision was made that there was a legal rule that superseded the earlier simple rule, since it would be the rule that appeared in that very decision. Moreover, of course, the antecedent of that rule would be true, so the argument would be sound.

If the possible decisions, even the possible legitimate decisions, that a judge could make are uncodifiable, then, when she comes to list the rules that she might use in making a decision, that list would be uncodifiable. Why might the possible legitimate decisions be uncodifiable? The obvious explanation is provided by the uncodifiability of morals. If, as a matter of fact, the legal rules contain moral requirements as soft positivists like Hart allow²⁴—if, for instance they require that a given expectation be *reasonable*, or a burden not be *unfair*, or that there be *due care*—and if the application of those moral requirements is uncodifiable as the moral particularists allege, then the uncodifiability will be inherited by the legal rules.

I described this move as a 'trick' but I didn't mean that in a derogatory sense. There is no real trickery here. We are not saying that the legal rules at any one moment are uncodifiable; an observer could take a legal snapshot of how things

²³ For arguments against the first, see J Raz, 'Authority, Law and Morality' *The Monist*, vol. 68, 295-324; similar arguments could be applied to the second. My own primary worry with either approach stems from a scepticism that they would result in a unique consistent extension.

²⁴ There is a weaker and a stronger thesis that soft, or inclusive, positivists might accept. The weaker is the idea that certain moral requirements may have entered into the law by precedent; the stronger is the thesis that the rule of recognition might allow moral requirements to be incorporated into the law directly, with no basis in precedent. I need only the weaker thesis here. (For a nice presentation of the distinction see Brian Leiter, 'Beyond the Hart/Dworkin Debate: The Methodology Problem in Jurisprudence', *American Journal of Jurisprudence* 48 (2003). He takes the Hart/Raz dispute to concern only the stronger thesis. It would certainly suit me if the weaker thesis were uncontroversial, but it frequently seems as though some of those on the Raz side of the debate were denying it.)

stand, and could exhaustively characterize the content of the law, though of course this content might include moral terms that are themselves uncodifiable. We are rather saying that the possible legal rules confronting a judge at the moment of making a decision are uncodifiable; and since these rules would be made actual if the judge endorsed them, something that it is in her power to do, it is the uncodifiability of the possible rules that matters. The point is simply that, once we concede the degree of discretion to judges that is needed if they are to distinguish cases, then judges cannot regard the law as fixed at the point at which they decide a case.

We must not confuse this sense of discretion with another notion of discretion that is employed by Hart, a notion that is sometimes called ‘strong discretion’.²⁵ This is the idea that in certain cases the law gives no right answer: judges can legitimately decide either way. Hart accepts that are instances of strong discretion—cases when the law simply leaves a gap—but he wants to greatly limit their extent. As we saw, he wants to say that rules can contain *unless*-clauses, and yet, when these clauses are not triggered, they will be binding.

Now though we might worry whether this position is really available to Hart. If every legal rule contains an *unless*-clause, and if judges have the discretion needed to distinguish exceptions, then what is to stop them at any point deciding to create a new rule that supersedes the rule that otherwise would have applied? Of course there might be a rule governing when an exception can be made; but if that rule in turn contains an *unless*-clause, that can be superseded; and so on all the way up. Such an approach might be attractive to legal realists, but it clearly would not have been attractive to Hart.

These issues take us into the much contested debate about judicial discretion. I shall leave this debate for now though. First I want to look at the question of what role legal rules can play if they are understood in this particularist way. With that in place we can return to issue of discretion.

THE FUNCTION OF LEGAL RULES

It is often assumed that rules, especially on a positivist conception, must act as a decision procedure. Thus for example, Scott Shapiro argues that positivism is committed to the idea that rules guide judges’ decisions, and Fred Schauer’s book on rule-following is subtitled *A Philosophical Examination of Rule-Based Decision Making in Law and in Life*.²⁶ But if the kind of legal particularism that I have outlined is correct, this is a role that they can fulfil at best only partially.

In its exact sense, the idea of a decision procedure is the idea of a purely mechanical procedure for arriving at the conclusion that a given thing either does or

²⁵ The term comes from Dworkin; see *Taking Right Seriously* (London: Duckworth, 1977) pp, 32ff.

²⁶ Scott Shapiro, ‘On Hart’s Way Out’ in J. Coleman (ed.) *Hart’s Postscript* (Oxford: Clarendon Press 2001); Frederick Schauer, *Playing by the Rules* (Oxford: Clarendon Press 1993).

does not have some property. In that sense decision procedures are rather hard to come by. Famously, there is not even a decision procedure for assessing the validity of an argument of the predicate calculus. But the problem with using particularist rules is much more basic than that. Even in a very loose sense of decision procedure (whatever that might be), particularist rules are clearly hopeless: one cannot start by assembling one's particularist rules and seeing how they apply to a given set of premises, since the set of potentially relevant rules is uncodifiable. We are thus far worse off than we are when trying to assess the validity of an argument in the predicate calculus, for at least there we have the argument, and the rules of inference are known.

For this reason, those who try to construct formal models in AI that mimic the reasoning employed in moral or legal domains do not employ anything like the apparatus that I have suggested.²⁷ My response is to concede that I am not trying to construct a decision procedure. Many people have objected that actual legal reasoning cannot be captured by an overly mechanical model; I am happy to agree with the point.

A related consideration comes from social psychology. There is a growing movement, led by psychologists such as Jonathan Haidt, to see moral arguments as typically *post hoc*. People do not arrive at moral conclusions as a result of going through an argument: rather, they arrive at a conclusion, and then construct an argument to support it. There are various quick, unconscious, inaccessible, systems that do most of the work of arriving at the initial judgment: emotional responses, employment of stereotypes etc.²⁸

One should be careful in applying this directly to the legal case. Preliminary neurological evidence suggests that Haidt's thesis is only true of certain moral judgments—those that directly involve the agent. In contrast, impartial moral judgements, which are similar to the kind of impartial judgments that judges typically have to make, look, at the neurological level, more like non-moral judgments.²⁹ But again I would be happy to concede that the factors that bring judges to their initial judgment are very diverse, and rarely if ever involve the kinds of rules that I have been describing.

So what are the rules good for if, in many cases, they are of no use in arriving at a legal decision? I suggest that they have two roles: (i) as confirmation to the agent that their initial reaction was good; (ii) as a justification to others that the decision was correct, and a guide to their use as precedent.

²⁷ See for instance John Horty's discussion in 'Nonmonotonic Logic', *The Blackwell Guide to Philosophical Logic*, L. Goble (ed.), (Oxford: Blackwell, 2001), pp. 336–361, in which he complains about a proposal along the lines of that suggested here, on the grounds that the possible *unless*-clauses would be open-ended, and that we would be in no position to verify that they didn't obtain. Instead he suggests using a default logic like that of Reiter's.

²⁸ The best recent presentation of such a view is in Jonathan Haidt, 'The Emotional Dog and Its Rational Tail' *Psychological Review* 108 (2001) 814–34

²⁹ Joshua Greene *et al.* 'The Neural Bases of Cognitive Conflict and Control in Moral Judgment' *Neuron* 44 (2004) 389–400.

(i) *Rules as aiding confirmation*

There are two ways of interpreting Haidt's position, or, more plausibly, a continuum of ways between two poles. On the cynical view, the moral intuitions are little more than blind prejudice, which the *post hoc* moral arguments then cover with a veneer of respectability. On the uncynical view the moral intuitions provide a quick, defeasible but generally reliable first pass, which may then be corrected in the light of later deliberation and argument construction, whether this is done on one's own or with others. I suspect that for moral decisions the truth lies somewhere between the two extremes. Certainly this is Haidt's position, though he thinks that things are closer to the cynical pole than many people have realized.³⁰

I suggest that for legal reasoning the position is closer to the uncynical pole. For a judge the ability to construct a plausible legal argument is crucial. It may be that a competent judge could construct a plausible argument for just about any conclusion, though I rather doubt that (I'll discuss it later). But even if that were true it doesn't show that any argument they construct is just window dressing. The argument remains a crucial part of the judgement. And if they were unable to construct one (however unlikely that might be), their initial judgment would have to be revised.

(ii) *Rules as aiding justification*

It is scarcely controversial that legal arguments are used to justify decisions. So long as one accepts that decisions have a *ratio decidendi*, the obvious thing to identify it with is an argument, structured round a legal rule. I suggest then that the second, and perhaps most important, role of a legal rule is to structure an argument so that it provides a justification of the decision. To play this role the rules do not need to give us a decision procedure; it is enough, once the argument is constructed, that we can recognize it as good. John Dewey argued for the centrality of such a role:

Courts not only reach decisions; they expound them, and expositions must state justifying reasons. The mental operations herein involved are somewhat different from that involved in arriving at a conclusion. The logic of exposition is different from that of search and inquiry. ... [Exposition's] purpose is to set forth the grounds for the decision reached so that it will not appear as an arbitrary dictum, and so as to indicate the rule for dealing with similar cases in the future.³¹

The presence of an open-ended *That's it* clause means that it is hard to be certain that an argument is good: could there not be some other rule that one should have considered that, together with certain facts would supersede the current argument? But to provide adequate justification we do not need certainty.

³⁰ One shouldn't, however, overlook how much credence he gives to the uncynical pole. He describes his position as 'social intuitionism', and then characterizes intuitionism as a kind of cognition, parallel to the philosophical position that sees it as a way of grasping moral truths. And in response to various critics he has stressed how much there is for reason to do; see 'The emotional dog does learn new tricks' *Psychological Review*, 110, (2003) 197-198, and 'The emotional dog gets mistaken for a possum' *Review of General Psychology*, 8 (2004), 283-290.

³¹ J. Dewey 'Logical Method and Law' *The Philosophical Review* 33 (1924) 560-72, at p. 570. The piece was simultaneously published in the *Cornell Law Quarterly*.

To whom is the justification directed? It is commonly held that courts are required to justify their actions to those they might coerce. Perhaps they do owe a justification, though as a thesis about the behaviour of courts it doesn't seem to hold generally: the decisions of French courts, for instance, are often little more than statements of the verdict. I suspect that much of the reason that courts in the common law system give such lengthy decisions is specific to that system: they are concerned to justify them to other courts, partly because they are vulnerable to appeal, and partly because they know that they are setting precedent, and so know that they need to articulate them sufficiently fully for them to serve in that role.³²

CONSTRAINTS ON DISCRETION

Let us return now to the difficult issue that was shelved above. We wanted legal rules to provide a constraint on what a court could decide: as Hart says 'It does not follow from the fact that rules have exceptions incapable of exhaustive statement, that in every situation we are left to our discretion'. So if rules are to contain a *That's it* condition, that condition had better provide a constraint. And it won't if judges can distinguish whenever they want.

An obvious first move is to say that there are further meta-rules, something like Hart's 'rules of change', governing how judges can amend rules. For Hart these are part of the secondary rules.³³ Such rules, we might suppose, will constrain when judges can legitimately distinguish, and when they cannot. The problem comes back if we think that these rules themselves contain *That's it* conditions, and any further rules binding those rules contain *That's it* conditions, and so on all the way up. Then judges would always have strong discretion, since they could distinguish all of the rules all the way up.

We then need to think about the nature and function of these meta-rules. It is surely unproblematic to hold that there is a broad distinction between the particular rules of the common law—those that define theft, for instance, or the defence of provocation—and the kind of rules that judges use in deciding how to apply and distinguish those rules. Since we are thinking of the *That's it* clauses as implicit, it is unlikely that there will be explicit rules governing when they may be triggered.³⁴ Instead the meta-rules are better thought of as very general principles: that no one should benefit from their own wrongdoing for instance, or that the act doesn't make for guilt unless the mind is also guilty. And clearly these rules are sufficiently open-ended that we can think of them as in turn containing implicit *That's it*

³² For a lengthy and enlightening discussion of precedent which sees such reason giving as central within the common law tradition see Duxbury *The Nature and Authority of Precedent*, esp. Chs. 2 and 3. Duxbury goes on to say that the law should not be understood as needing to be in conformity with logic (see esp. pp. 140–9). I think that this involves an over monolithic view of logic; the important question is rather the kind of logic that the law has.

³³ See *The Concept of Law* pp. 93–4.

³⁴ Although there are a few places where something along these lines is explicit; for instance the first clause of the Canadian *Charter of Rights and Freedoms* reads: 'The Canadian *Charter of Rights and Freedoms* guarantees the rights and freedoms set out in it subject only to such reasonable limits prescribed by law as can be demonstrably justified in a free and democratic society.' Thanks to Dénise Réaume for discussion here.

conditions. What must be mistaken, however, is to think that judges are only constrained in the distinctions they can make insofar as there is a meta-rule governing the distinctions they can make.

The point comes sharply into focus when we consider the courts' attitude to the notion of precedent itself. In 1898 the House of Lords made a famous ruling for itself in *London Tramways*, deciding that it would be bound by its own precedent. In 1966 the Lords reversed the decision, declaring in the Practice Statement that they would depart from previous decisions when it was right to do so. Many authors have argued that there is something paradoxical, or at least question begging, about these decisions. After all, why should the Lords have ever felt themselves bound by the *London Tramways* decision if they did not already acknowledge the power of precedent? And if they did acknowledge that power, then what did the decision add?

But we mustn't let the attractions of paradox-mongering obscure the fact that clearly *London Tramways* and the Practice Statement did play some important role. The obvious lesson to draw is that respect for precedent must indeed precede any attempt to codify it in a rule. Once a practice of respect is there, then legal rules, including legal rules about the very nature of precedent, make sense as ways of making that practice both more precise and more firm. I think that they are best understood as akin to resolutions: the Lords resolved not to overturn their own judgments, and the resolution got its force not just from its status as a rule, but also from the attitudes of the Lords themselves. Similarly, the practice statement can be understood as a new resolution, itself prompted by, and ineffectual without, the changing attitudes of the Lords. The role of resolutions is to block reconsideration: not to block it absolutely, since circumstances can change, but to raise the threshold for reconsideration.³⁵ This is what the Lords did. In effect they resolved in *London Tramways* to be bound by precedent even in cases where disregarding it would otherwise have led to a more just outcome. And they did this because they thought that there was more to be gained by a system that was transparent and predictable in this way, than would be lost from the occasional less than perfectly just decision. Their resolutions could have had no effect if they did not already have appropriate attitudes towards precedent, but given those attitudes they do have effect: *London Tramways* increased the influence of precedent even in those cases in which the Lords would have otherwise ignored it.³⁶

Hart held that we only get a legal system when the officials take the internal point of view towards the secondary rules, where this involves thinking that they have *reasons* for obeying and enforcing the law, reasons that give rise to *obligations*. Reflection on the Lords' handling of precedent shows that such attitudes should not be confined to rules. It is not simply that judges must think they have an obligation

³⁵ For the idea that resolutions must block reconsideration because temptation tends to corrupt judgment see my *Willing, Wanting, Waiting* (Oxford: Clarendon Press, 2009) esp. Chs. 5 and 6.

³⁶ There is an interesting question whether making the *That's it* clause explicit would tend to make it feel less binding, just as focusing on the fact that resolutions can be amended makes them feel less binding. An interesting test case would be provided by the Canadian *Charter* mentioned above.

to obey certain rules; it is also that they must think they have an obligation to behave in certain ways, independently of whether those ways are described in rules.³⁷

My suggestion is that considerations of this kind are what block the threat of complete indeterminacy. Judges just have attitudes to what is and isn't a legitimate distinction, attitudes that are independent of legal rules. So even if the legal rules all have *That's it* clauses, with the result that, consistently with the legal rules, any case could be decided any way, this does not entail that judges would think that they could decide any case any way, since they think themselves bound by more than just the legal rules.

However, this still leaves open the issue of how tightly judges are bound by this combination of these attitudes and the legal rules. It is one thing to say that we are subject to someone else's authority, quite another to say that they leave us no choice in what to do, since it could be that the authority hardly ever issues any commands. Things are similar in the case of the law: to say that judges are constrained by the legal rules and their attitudes is not to say that they are constrained their every decision, for they might leave things quite indeterminate in many cases. If we reject the position that says that they are bound in every case, we are left with a continuum from Hart's position that they are bound quite tightly in most cases, to the Legal Realists' position that, at the appellate level and above, such binding is largely illusory. I cannot do anything like justice to this disagreement here. Let me just say two things, one about the psychology of the matter, and one about the facts that stand behind it.

First, as a psychological thesis, it strikes me as very plausible that even judges in appellate courts *think* of themselves as bound in most cases to find as they do. The strong version of legal realism which depicts legal reasoning as mere pretense asks us to believe that judges behave in ways in which human beings rarely behave. Judges write decisions in which they present themselves as having been required by the law to come to the verdict to which they have come.³⁸ There is a large psychological literature on self-presentation, which makes it very clear that people are highly motivated to present an impression that they are behaving well to other people, and that, having presented such an impression, they subsequently tend to conform to it. Publicity makes all the difference: if we are seen by others to behave in a way that exhibits certain attitudes, we are much more likely to internalize those attitudes.³⁹ It is hard to think of a forum that is more public than a judge delivering an opinion, an opinion that will be recorded and made freely available.

It is true that there is typically no sanction (or no sanction other than the disapproval of the public and their peers) if judges break the rules that should govern their decisions. In that sense they are free to do as they please. But from that it should not be concluded that they have no motivation to adhere to the rules.

³⁷ Matthew Kramer makes this observation about the internal point of view, *In Defence of Legal Positivism: Law Without Trimmings* (Oxford: Clarendon Press 1999). For an interpretation of the Lords' behaviour along these lines see Neil Duxbury, *The Nature and Authority of Precedent* pp. 122ff.

³⁸ Though admittedly they do sometimes make comments outside the courts that undermine this impression. See for instance Richard Posner's recent comments on U.S. Supreme Court decisions in the 'Foreword' to the *Harvard Law Review* 119 (2005) pp. 31–102. Here I am tempted to say that the exception proves the rule.

³⁹ See D. Tice 'Self-Concept Change and Self-Presentation: The Looking Glass Self Is Also a Magnifying Glass', *Journal of Personality and Social Psychology* 63 (1992) 435–51.

Empirical work on why people obey the law suggests that threat of sanctions is never a very important factor. Far more important is if they feel that the system is legitimate and fair, that it treats them with respect, has mechanisms for listening to their side of the case and so on.⁴⁰ On all of these counts the legal system surely serves judges very well. One would expect them to be well motivated to obey any demands that they think it puts upon them.

This doesn't decide the question of whether judges think that they are typically bound by the law to come to the decisions to which they come—though it is hard to know what could, given the question's nature—but I think it does lend some credibility to Hart's modest contention that, for the most part, legal decisions

are reached either by genuine effort to conform to rules consciously taken as guiding standards of decision or, if intuitively reached, are justified by rules which the judge was antecedently disposed to observe and whose relevance to the case in hand would generally be acknowledged⁴¹

However, that is only the psychological question. Even if we concede that judges *think* that in most cases there are legal rules that bind them in most cases, there remains the question of whether they really are. Perhaps, though judges are not dissimulating, they are systematically self-deceived. The cognitive dissonance literature is full of examples of people fabricating stories in order to reassure themselves that they are more consistent or honest than they really are. Couldn't this be another example?

Of course it could, but even here I think that we currently have very little evidence to tell one way or the other, and that what it commonly held to tell in favour of a legal realist approach does not. For a good representative position of the kind of evidence I mean, consider the following post to Brian Leiter's blog from a practicing lawyer:

I wonder at the sanity of those who think that it is realistically possible in close and controversial cases on issues of public import (and most Supreme Court cases are precisely that or they wouldn't have had cert granted) to ignore your own predilections and beliefs. In fact, it is just such cases that seem to call for a leap of faith or logic at some critical point; the means to accomplish that leap often stems from personal beliefs. It is so obvious to those of us who live in the courtroom that it is hardly worth discussion. No single fact or circumstance is more important than to know the judge who will decide the case. To take another example, every firm that I know of with a significant Supreme Court practice brings to bear on a Supreme Court brief and argument the views of former clerks to as many current Justices as possible. There is only one reason to do that, and it is not to consider the precedents or even "super precedents."

⁴⁰ Tom Tyler, *Why People Obey the Law* (New Haven: Yale University Press 1990)

⁴¹ *Concept of Law* 137. The second disjunct is remarkably weak; it is hard to think that anyone would deny that the rules were *relevant*; the contentious question is whether they *determine* the decision.

We can concede that different judges will decide differently, and in characteristic ways. And we can also concede that one of the best ways to predict how a given case will be decided is to find out who will be deciding it and then to come to know something about them. The same is doubtless true of any field in which there are controversial issues. It is surely true in philosophy, and I expect that it is true of much of natural science: the best predictor of whether someone will agree with a controversial claim will come from looking at what they say about other related claims. This clearly does not show that there are no right answers in science. If the same does not appear to be true of mathematics, that is surely because, in large areas of mathematics, there is so little controversy.

Clearly though the writer is after something else, as becomes clear in their references to the importance of the judges' 'own predilections and beliefs'. What exactly is at issue here? Of course if people interpret the law differently, they will have different beliefs about how the law should be interpreted. No one could be claiming that judges could ignore those beliefs. So the charge must be that there is a particular kind of belief, a 'personal belief' as the writer characterizes it, that influences how judges decide: most plausibly, that judges are highly influenced by moral, political or religious beliefs. Here arguably we do get a contrast with the natural sciences, and perhaps with some areas of philosophy. It would be surprising if, for instance, the holding of different interpretations of quantum mechanics correlated to the holding of different moral or religious beliefs; and if they did we should be sceptical that people were being moved by the evidence that should move them.

But it is quite unclear to me that a similar conclusion can be drawn in the legal case. Suppose that the law were completely determinate, but that it were hard to work out what it required. Then we would expect there to be disagreement, and it would not be surprising that that disagreement would correlate with disagreement on issues of morals or politics or religion. For it is quite plausible that, unlike with quantum mechanics, skill in one of these areas would auger well for skill in another.

I suspect that if we are to construct an argument for the degree of determinacy in law it will come from the other direction: not from reflections on the behaviour of judges, but from reflections on the nature of law. Here I think that an approach that treats law as an artifact will find good grounds for rejecting the idea that there will always be a right answer in hard cases; and perhaps it will say enough to show that the different behaviour of different judges cannot after all be explained as a result of differences in their skill. But clearly that is a project for another paper.