

Defeasibility and Pragmatic Indeterminacy in Law

Andrei Marmor

Defeasibility, in one standard sense, is a feature of inferences and one that seems to defy classical first-order logic. One of the fundamental postulates of classical first-order logic is *monotony*: If there is a sentence φ that can be inferred from a set of sentences X , then φ should also be inferable from any set of sentences that is an extension of X . Monotony captures the idea that the addition of premises to a valid inference cannot detract from its validity: If premises $\{a, b \& c\}$ entail φ , any set of premises including $\{a, b \& c\}$ as a subset would also have to entail φ . But that seems to be exactly the principle that is violated by the defeasibility of an inference. The standard case of defeasibility runs like this: We have an inference whereby premises $\{a, b \& c\}$ warrant the conclusion that φ . However, with the addition of another premise, d , the conclusion that φ is no longer warranted. In short, defeasibility occurs when the addition of premises to a valid inference modifies the conclusion, which is exactly what monotony would not allow.

Here's a favorite textbook example: Suppose you are told that Tweety is a bird; you would be warranted in concluding that Tweety can fly. However, upon learning that Tweety is an ostrich, the conclusion is reversed. Now you would have to conclude that Tweety cannot fly. Examples seem abundant in law: Suppose that X intentionally killed a person. Under normal circumstances, this would be murder and punishable as such. However, further premises might defeat the legal conclusion: If X acted in self-defense, his intentional killing is not murder; it may not be an offense at all. Pragmatic inferences, particularly conversational implicatures, are also defeasible; the implicated content can be defeated or canceled by the addition of further pieces of information. For example, if I tell you that "John and Mary went to Paris last summer," you'd be warranted in inferring that they must have gone to Paris together. But suppose I add some further information,

saying, “You know, they never met while in Paris.” So now you know that they did not go there together.¹

The question of whether all defeasible inferences are non-monotonic is a highly contentious issue and I will not attempt to resolve it here.² I will certainly assume that many of them are. But, in any case, it is quite obvious that part of what allows for the possibility of defeasibility of an inference consists in the fact that the inference contains some *typicality* premise. The inference from “Tweety is a bird” to the conclusion that “Tweety can fly” is premised on the assumption that birds typically can fly; most of them fly, or, by and large, birds fly, or something along those lines. Similarly, the inference from “X intentionally killed person Y” to the conclusion that “X committed murder” is based on the premise that typically, under normal conditions, etc., intentional killing of a person is murder. The inference is defeated when some additional piece of information (e.g., that Tweety is an ostrich, or that X acted in self-defense) demonstrates that the typicality premise *fails on this occasion*.

Epistemologists have long identified two types of defeat: *rebutting* and *undercutting*. A rebutting defeat is one in which the additional (henceforth: superseding) premise to a prima facie warranted inference is such that it negates the conclusion of the inference. The superseding premise to “Tweety is a bird,” that “Tweety is an ostrich,” is such that it simply negates the conclusion that “Tweety can fly”; ostriches cannot fly. Hence it is a rebutting defeat. In other cases, however, the superseding premise might be such that it undermines the initial evidence we had for the conclusion. For example, suppose an object in front of me looks red. Under normal conditions what looks red is red, and thus I may conclude that the object is red. However, suppose I learn the superseding fact that I was given a drug that makes me see everything in red. This superseding premise undercuts my evidence for concluding that the object in front of me is red. Notice that it does not necessarily negate the conclusion that the object is red; for all I

¹ Notice that these conclusion reversals can go on for a while; suppose I add “they separated as soon as they arrived at the airport,” so now you would realize that they did go together after all.

² Probabilistic reasoning, such as Bayesian epistemology, is in many cases the main alternative model, and it is sometimes very difficult to nail down the boundaries between inferences susceptible to non-monotonic logic and those that are instances of credence adaptations.

know, it might be red. What is defeated here is the reliability of the initial evidence I had for concluding that the object is red. These are called undercutting defeats.³

My main argument in this paper is that, at least in the practical domain, in addition to rebutting and undercutting defeats, a third type of defeat exists, which I will label *conflicting* defeats: cases in which the superseding premise renders the initial inference genuinely *indeterminate*. A conflicting defeat neither negates the conclusion nor undercuts the initial evidence for it. The defeasibility in such cases consists in the fact that it becomes indeterminate whether the putative conclusion follows or not – namely, it is a conclusion that one would not be unreasonable to deny, nor unreasonable to affirm. My discussion, and most of the examples, will focus on pragmatic inferences, legal inferences and on some overlapping cases, that is, cases in which legal defeasibility is actually a matter of pragmatics. The upshot of this discussion is to show that defeasibility in law sometimes generates a genuine kind of legal indeterminacy. From a legal point of view, the conclusion would be neither here nor there. In such cases, decision-makers must make their judgments on the basis of considerations not dictated by the relevant law.

1. Pragmatic Defeasibility

The general defeasibility of pragmatic inferences is widely recognized by linguists and philosophers of language. Perhaps less widely recognized is that the defeasibility of pragmatic inferences probably shows that their underlying logic is non-monotonic.⁴ But the non-monotonic logic that may apply here is not my concern. What we need to see is what makes pragmatic inferences defeasible, in what ways, and what kinds of defeasibility are in play. For the most part, I will focus on standard instances of conversational implicatures, and I will largely work within the original Gricean

³ The idea of undercutting defeats was probably presented first by J. Pollock, "The Structure of Epistemic Justification", *American Philosophical Quarterly*, 78: 62-78 (1970). It has been widely used since.

⁴ Part of the problem here stems from the fact that the development of a formal non-monotonic calculus proved to be a very tricky matter (though there is a huge interest in it in computer sciences). There are several suggestions in the literature but they all suffer from shortcomings and limitations. The most famous, perhaps, is R. Reiter "A logic for default reasoning" *Artificial Intelligence*, 13:81-132 (1980) Some linguists and computer scientists have applied default logic to conversational implicatures; see, for example, M. Morreau, "How to derive conveyed meanings" (1995, posted online), M.A. Walker, "Inferring Rejection by Default Rules of Inference" (1996, posted online).

framework,⁵ though two modifications may be needed. First, Grice probably thought that conversational implicatures, and perhaps pragmatic enrichment more generally, concern content that goes beyond what is said (viz., asserted) by the speaker in a given context of utterance. It is fair to say that it is now widely recognized that pragmatic enrichment also, and quite often, plays a crucial role in determining what is said, and not only the kind of content that goes beyond it. In other words, the assertive content of an utterance on a given occasion of speech is often pragmatically enriched content, not fully determined by semantic and syntactical determinants. Therefore, for our purposes, the distinction between pragmatic inferences that contribute to the asserted content of an utterance on a given occasion of speech, and those which concern implications beyond what is said, will be of little consequence. I will not try to be very precise about this.⁶

The second modification may be a bit more controversial and more relevant to our concerns. Suppose a speaker S utters a sentence P in context C, thereby conversationally implicating that Q. There are probably two ways to answer the question of what “implicating that Q” stands for. I think that for Grice it meant something like the subjective communication intention of S: By saying P in context C, S *intended* to convey the additional content that Q. In other words, grasping the implicated content consists in grasping some content that the speaker intended to convey. I am not entirely sure that this is what Grice had in mind, but given his overall linguistic framework, it is a very plausible assumption. However, it is doubtful that it is also the best option. On a more objective understanding of communicated content, we can think of the idea of implication in terms of commitments, rather than intentions. On this understanding, content that is conversationally implicated by saying P in C is the kind of content that the hearer can *reasonably* infer that the speaker *is committed to* by expressing P, given the relevant conversational maxims that are taken to apply. And of course it matters which one of these options we assume; the criteria for warranted inferences would vary accordingly.

⁵ P. Grice, *Studies in the Ways of Words* (Harvard University Press, 1989).

⁶ The same goes for the distinction between implicatures and utterance presuppositions. On the differences between these various forms of implication, and their potential relevance to law, I have elaborated in my *The Language of Law*, (Oxford, 2014), chs 1&2.

Let me use one of Grice's famous examples to illustrate the difference. Think about the scenario in which S observes H standing near his stalled car, which ran out of gas, and says to H: "There is a gas station around the corner." Now of course S has not actually said that, for all he knows, the gas station is open and would have gas to sell, but given the ordinary maxims of conversation, particularly the maxim of relevance, and the context of the expression, the utterance clearly implicates that, for all S knows, the gas station is likely to be open. It is the kind of inference that the hearer would be quite warranted in drawing. Suppose, however, that as a matter of fact S actually knows that the gas station is closed, and has no intention to implicate otherwise. (Notice that S did not say anything that is false; after all, he only said that there is a gas station around the corner, and that is true, there is one.) Now the problem for the subjective intention-based interpretation here is fairly obvious: If we think that implicated content is constituted by the subjective intention of the speaker, then the truth of the matter here is that there is no implication that the gas station is open; the speaker had no intention to convey such content. However, that does not seem right. Naturally, we would want to say that, regardless of what the speaker had in mind, the utterance in this context commits the speaker to the content that, for all he knows, the gas station is likely to be open. It is a commitment that follows from the content of what S said, the factual circumstances and the normative framework of a cooperative discourse that is presumed to apply here. Generally speaking, a speaker can be expected to be committed to a certain implicated content if and only if a retrospective denial of that content by the speaker would strike any reasonable hearer, sharing all the relevant background knowledge in the circumstances, as contradictory, perplexing or disingenuous. Thus, on this understanding of what implication consists in, the speaker in our example would have implicated that for all he knows the gas station is likely to be open, even if, for some reason or other, he expressed the utterance not actually intending to implicate this content. And this seems like the right result.⁷

None of this means, of course, that a speaker cannot cancel the implication of his utterance by some additional clarification. Even if we think about implicated content as

⁷ I will have more to say about this objective standpoint in the last section.

the kind of content that the speaker is committed to, given what she said in the circumstances, etc., the commitment can be canceled explicitly. In our example, the speaker can simply add a comment suggesting that the gas station may not be open – in which case, of course, the implication would have been canceled. As Grice and others have rightly noted, cancelability is an essential aspect of conversational implicatures.⁸ Cancelability, however, is only one way in which an implicature can be defeated; it is defeated by an explicit denial by the speaker himself. And note that canceling an implicature is typically (though not necessarily) a rebutting defeat; by canceling the implication the speaker typically negates a certain content that would have been otherwise implicated. Cancelability, however, is a special case of defeasibility. Implicatures can be defeated in other ways. The cases I want to highlight concern contexts in which an inference to implicated content comes into conflict with a salient contextual fact presumed to be common knowledge between speaker and hearer. Let me give some examples:

(1). Consider the utterance, “Senator McCain and the Republicans voted against the bill.” Under normal circumstances, the implicature (or utterance presupposition) of this utterance would be that McCain is not a Republican. But if it is known to the hearer, presumed by the hearer to be known by the speaker, and known by the speaker to be known by the hearer, that McCain *is* a Republican senator, the implicature becomes perplexing and doubtful. All we can tell here is that the speaker must have had a reason to single out McCain, but we do not know what it is. Some content would seem to be implicated by the speaker, but it is hanging in the air, in need of specification.

(2). Consider the utterance, “John and Mary went to Paris last summer.” Under normal circumstances, one would infer that they went to Paris together; it would

⁸ The notable exception concerns what Grice called “conventional implicatures.” As I explained elsewhere in some detail (*The Language of Law*, ch 2), I think that conventional implicatures are the kind of implications that are semantically encoded in the expression used and, thus, not cancelable.

be the natural implication of such an utterance. Suppose, however, that it is common knowledge between speaker and hearer that John and Mary went through an ugly and bitter divorce last year. Under such circumstances, the implicature becomes doubtful. Notice, however, that the implication is not necessarily rebutted; it just becomes very doubtful.

(3). Suppose your colleague says, “All the graduate students who passed the exam went to the party last night.” Naturally one would assume that some students did not pass the exam. But suppose that you know that all the graduate students passed the exam, and you know that your colleague is perfectly aware of that. The natural implication would not follow here; something may have been implicated, but we cannot quite say what it is.

What we see in such cases is that a pragmatic inference can be defeated by the presence of some contextual facts, presumed to be common knowledge between speaker and hearer. Such contextual facts can either undercut the inference or, at other times, render the inference indeterminate or underspecified. And this should not be surprising. After all, the determinants of a pragmatic inference include contextual background that is presumed to be common knowledge between speaker and hearer. When some of the background facts somehow conflict with the inference to implicated content, the inference might be defeated.

I mentioned earlier that the defeasibility of an inference is enabled by the fact that such inferences contain typicality premises. The inference is normally defeated when the superseding premise demonstrates that the usual or the typical fails to obtain in the particular case. I hope we can now see that this is precisely what happens in the cases of conversational implicatures that are defeated by some salient fact that obtains in the particular context of the conversation. Let me use the implicature from example (2), above. When somebody utters (2), the normal implication would be that John and Mary went to Paris *together*. Why? Because otherwise the speaker would have said too little. If

the speaker knows that Mary and John went to Paris separately, the speaker would have omitted some relevant information. But notice that the maxim of quantity comes into play here already entangled with some typicality premise in the background. It is a premise about reasons: The idea is that one would normally have no reason to mention in the same sentence two people doing something unless they were doing it together. Generally speaking, then, the idea is that we say things for reasons, and assumptions we make about typical (ordinary, common, etc.) reasons often play a role in the pragmatic inferences we draw from what we hear. Thus, generally, we would assume that if you say something like (2), you implicate that John and Mary went to Paris together. But now, if we happen to know (given their bitter divorce and all) that it is very unlikely that those two did anything together last summer, the typicality premise fails on this occasion. The hearer would still surmise that there must be a reason for the speaker to mention the pair in one sentence, but it is not the usual, typical reason; it must be something else. Similarly, in example (1), when a speaker singles out an individual from a set, we would assume that the individual is not a member of the set; otherwise, the speaker would have said too much. But again, the maxim of quantity comes into play here on the basis of the assumption that there are reasons to single out the individual from the set mentioned in the same sentence, and the most typical reason would be the intention to indicate (or to acknowledge) that the individual is not a member of the set. When it is common knowledge that the singled-out individual is a member of the set, the typicality premise fails; we can no longer rely on it. And I take it that very similar considerations apply to (3) as well. When a speaker says, “All X’s who are F ...,” we assume that for all the speaker knows there might be an X who is not F. When we know that the speaker believes that all X’s are F, the typical reason to predicate the X’s in this way fails on this occasion, and the assumption has to be that there is something else going on.

The fact that assumptions about typical reasons play a role in pragmatic inferences is of crucial importance. It explains why defeasibility can generate genuine indeterminacy. Generally speaking, reasons can come into conflict. They also have something we call a dimension of weight. When reasons conflict, we can often determine that one of the reasons outweighs the others in the circumstances. Indeterminacy comes about when conflicting reasons are on par with each other – either because they have

roughly equal weight in the conflict, or else, perhaps, because they are incommensurably on par. In both cases, the situation is that no reason in the conflict outweighs the other. So we get indeterminate results. This simple model, I submit, explains why pragmatic inferences can be defeated in ways that neither rebut the initial conclusion nor undercut the initial evidence for it, but render the conclusion indeterminate.

Admittedly, a distinction exists between indeterminacy and under-determinacy. Results are under-determined by some facts or evidence if the relevant facts do not provide sufficiently robust or unambiguous support for the particular result in question. In other words, under-determinacy is a relation between premises/evidence and a conclusion. Indeterminacy is a relation between results and some particular question the result is purported to answer. Notice, however, that under-determinacy can be the reason for the indeterminacy of the result. In some of the examples above, that is exactly what happens: The pragmatic factors that should enable the hearer to infer the implicated content under-determine the content of the implicature, which is what makes the implicature indeterminate relative to some pertinent question or interest. In principle, under-determinacy and indeterminacy are logically independent aspects of an inference. In practice, however, the former is often the reason for the latter. A particular result would be indeterminate relative to some relevant question, such as what exactly the utterance implicates in the circumstances about X, because the implicature about X is under-determined by the combination of the semantic and pragmatic inputs. Given the relevant semantic and pragmatic determinants in this case, the implicated content about X is neither here nor there, which is to say, the utterance is indeterminate about the implication of X.⁹

2. Forms of Defeasibility in Law

Logicians who write about non-monotonic logic like to give legal inferences as examples of what they have in mind. They think that law provides many examples where we

⁹ I am not suggesting that this is the only kind of relation between indeterminacy and under-determinacy generally; it is the one, however, that is relevant to our cases.

obviously employ non-monotonic logic.¹⁰ And they are probably right about that, but the examples are potentially confusing because they come from very different kinds of legal inferences. Only some of those cases concern the kind of defeasibility we are discussing here. So let me start by drawing some distinctions. To begin with, I would like to bracket inferences about matters of fact. Legal cases are based on evidence about matters of fact; facts need to be established in legally recognized ways. In drawing factual conclusions from evidence, we rely on many generalizations and typicality premises, and obviously in ways that are potentially defeasible. The theoretical framework of such factual inferences, however, is controversial. Epistemologists have different theories to model such inferences, and I will have nothing to say about the epistemology of factual inferences in law or elsewhere.¹¹ My concern is focused on legal inferences, that is, cases in which we infer a legal verdict (viz., a particular legal conclusion) from legal norms and established facts. The questions concerning ways in which those facts are legally established is something that I will not discuss here.

Even within the category of legal (as opposed to factual) inferences, however, there are several, potentially overlapping, categories of cases. The three main types that come to mind are legal inferences based on presumptions, legal inferences involving exceptions to legal rules and legal inferences about the content of legal rules.¹² However, I think that only the third category, concerning inferences about what the law says and implicates, are genuinely susceptible to non-monotonic defeasibility; presumptions and exceptions to rules concern different types of defeat.

The most misleading cases that may seem like standard examples of defeasibility in law concern rebuttable presumptions. The law often contains, as a matter of legal rule, or sometimes as a matter of entrenched legal doctrine, a general presumption, postulating an inference from X to Y, typically rebuttable in some recognized ways. But as lawyers know very well, presumptions come in different forms and function in different ways.

¹⁰ See, for example, Burges, *Philosophical Logic* (Princeton, 2009) p. 123. See also the entry on non-monotonic logic in the *Stanford Encyclopedia of Philosophy*.

¹¹ Bayesian epistemology is of course the main candidate here.

¹² There might be a fourth type of case that is unique to common law systems concerning the inferences courts draw from precedent and ways in which precedents are "distinguished". I will not discuss reasoning from precedent in this paper, it is complicated enough to deserve separate discussion.

Some presumptions, though not many, are not rebuttable. In effect, non-rebuttable presumptions amount to a legal stipulation that, for the relevant legal purposes, X counts as/or entails that Y.¹³ Such rules are not really presumptions. When the law says that, for purposes P, X counts as Y, the law just stipulates a certain result, and like any other legal rule, if the antecedent conditions are met, the legal conclusion is to follow. There is nothing special about such cases; they do not rely on typicality premises.

Most legal presumptions, however, are recognized to be rebuttable. But again, not all rebuttable presumptions in law are of the same kind. Consider, for example, the famous, and obviously rebuttable, *presumption of innocence*. The legal presumption that a man is to be considered innocent until authoritatively proven to be guilty, functions as a rule about allocation of burden of proof: The rule imposes the burden on the prosecution to provide legally permissible evidence and positively prove the guilt of an indicted defendant. In other words, the so-called presumption of innocence is essentially a rule about the allocation of burdens of proof.¹⁴ And it is just one among many such rules, often called presumptions. For example, there is a well-known (rebuttable) presumption in criminal law that a person intends the natural consequences of his action. What this rule essentially prescribes is, again, a matter of burden of proof. In order to show that the defendant on trial had the relevant criminal intention to X, the prosecution only needs to prove that the agent's action was such that its natural, typical consequence is X; the conclusion that in such cases the agent intended X is something that the prosecution does not need to adduce additional evidence about. The burden to rebut the conclusion is shifted to the defendant. The defendant needs to provide evidence and convince the court that, in spite of the natural consequence of his action, the consequence is not one that he intended to bring about. Many rules in law of evidence about shifting burdens of proof are defined in terms of rebuttable presumptions. It is a matter of convenience, entailing nothing special about the defeasibility of such rules.

¹³ Some types of sovereign immunity are legally defined as non-rebuttable presumptions, probably for rather archaic reasons.

¹⁴ Admittedly, there is also a colloquial understanding of this presumption that is very different from the legal one. In everyday, non-legal sense, the presumption of innocence is not a presumption at all, it is a kind of moral requirement to abstain from judgment about a person's guilt until proven otherwise.

Many other legal presumptions are rules of statutory interpretation. For example, a rule determines that statutory provisions formulated in the masculine are presumed to include the feminine, or a rule determining that a criminal offense is presumed to require mens rea.¹⁵ These, and many other such presumptions, operate as default devices. They provide a general default that legislators and administrative agencies are free to deviate from, but they would need to do so explicitly, indicating the deviation from the default by some recognizable means. The formulation of such rules as rebuttable presumptions is essentially a technical device, a matter of convenience. It is like telling the legislature: If you want to prescribe X in circumstances C, you do not need to say so; we will assume that if C then X. Therefore, if you want the legal rule in C to prescribe that not-X, you need to say so. To the extent that such rules are defeasible, and of course they may be, their defeasibility has nothing to do with their formulation as presumptions. To sum up, we should not take the formulation of rules or legal doctrines in terms of rebuttable presumptions to indicate that such rules are somehow more defeasible than any other legal rule.

The general defeasibility of legal inferences would seem to be very closely tied to the second category I mentioned, namely, cases in which courts need to consider an exception to a legal rule. First, however, we need to narrow the issue here. Countless legal rules, at various levels of generality, have legally recognized exceptions, either originally contained within the relevant enactment or added later by the legislature or by the courts in their interpretative decisions. Either way, when a case concerns the putative applicability of a recognized exception to a legal rule, the role of the court is only to determine whether the conditions for the exception are met in the particular case at hand. In this the role of the court is not different from any other type of case in which it needs to determine whether the conditions for the application of a rule obtain or not.

Defeasibility is thought to become the issue when we focus on the open-ended nature of possible exceptions to a legal rule. The thought is that perhaps any legal rule, whether it already has legally recognized exceptions or not, may become applicable to a

¹⁵ U.S. lawyers and judges often refer to these rules of interpretation as *canons of construction*. U.S. jurisprudence recognizes dozens of such canons of construction, but their nature is far from uniform and many of them are just rule-like formulations of well-established legal doctrines or precedents.

situation that might justify an exception to the rule, and there is no way of determining all those possible situations in advance. Legal rules have exceptions that cannot be exhaustively stated *ex ante*. One way of expressing this idea is to say that legal rules are defeasible. Some justified exception to the rule might always come up that would defeat the rule's original verdict.

Let me say from the outset that I would not wish to object to the label of defeasibility here, but I want to explain why the logic of exceptions is different from the logic of defeasibility with respect to inferences that contain typicality premises. These are two distinct types of cases. Richard Holton has suggested a way of modeling exceptions to legal rules and, though I have some reservations about the scope of his account, I think that he is quite right to argue that the logic of open-ended exceptions to legal rules need not defy monotonicity.¹⁶ I will try to show that monotonicity can be maintained here precisely because the inferences in question do not rely on typicality premises. Something else is going on.

So what is the phenomenon in question? The idea is simple enough, and it is not unique to law. The idea is that, when we formulate a general normative prescription by way of a rule, we would normally recognize certain exceptions to the rule. So we can normally state in advance that "R: In circumstances C, all X ought to φ , *unless* conditions a, b or c obtain." It seems reasonable to assume that every plausible, justified or sensible rule would have some *unless* clauses attached. So far there is not much of a problem; we can think of the *unless* clauses as narrowing the scope of the rule, that is, providing a more fine-grained definition of its conditions of application. The problem comes about from recognizing that the *unless* clause of any plausible rule would have to be open-ended. We cannot formulate an *unless* clause that would state all the possible factors that would defeat the rule in question. But then the danger is that if the *unless* clause does not state all the defeating factors, it is bound to be vacuous or empty, like saying that the rule applies unless it turns out that it should not apply. That is not very helpful.

¹⁶ R Holton, "Modeling Legal Rules" in Marmor & Soames (eds), *The Philosophical Foundations of Language in the Law*, (Oxford 2011), 165-183.

Now one can see why a non-monotonic logic would seem like a tempting model to apply here.¹⁷ We could think of rules as implicitly but necessarily containing a typicality premise. So the idea would be that “R: In circumstances C, all X’s ought to φ ,” actually says something like “R: In circumstances C, most/typical X’s ought to – and/or typically ought to – φ .” The unless clause, on this model, is understood as an indication that the typicality premise may fail on any given occasion, in which case, the inference from the rule to its application may not be warranted on that occasion. But I think that Holton is quite right to argue that this model is not all that plausible. Even if we are willing to give up monotonicity, which is always a high price to pay, the model misconstrues the ways we normally think about legal rules. Consider, for example, the legal rule *pacta sunt servanda*, in common law meaning that a valid contractual obligation must either be performed or paid compensation if not. We recognize, of course, that there are bound to be some exceptions whereby a breach of a contractual obligation need not be compensated. And we may also recognize that there is no way of stating all the exceptions *ex ante*. But still, we would not think about the rule as premised on some typicality condition; we would not think that the content of the rule is that “typically or mostly contractual obligations need to honored.” We would think that contractual obligations need to be performed (or compensated for non-performance) unless there is a good reason not to. So this is precisely the intuition that Holton aims to model. His model assumes that we can think of *unless* clauses as quantifying over other rules. A rule applies *unless* there is another (sound, justified) conflicting rule that renders the application of the first one wrong.¹⁸ And then we need a suppression tool, which Holton calls “*That’s it*,” to indicate that the application of a rule in any given case presupposes that no justified superseding rule is relevant. A rule is defeated when the “*That’s it*” clause is not warranted under the circumstances, namely, when there is some superseding rule that, under the circumstances, justifies the defeat of the rule in question.

¹⁷John Horty has developed a detailed and interesting account based on non-monotonic logic; see, for example, his “Reasons ad Defaults”, *Philosophers’ Imprint*, Vol 8, no 3 (2007). Horty’s work is partly the target of Holton’s argument about exceptions to legal rules. As will be seen below, in this particular context I side with Holton in this debate.

¹⁸Holton’s account is reminiscent of ways in which the concept of *prima facie* obligation has been thought of. The idea is that a *prima facie* obligation is a reason for action, such that failure to act on it is wrong unless there is a conflicting reason not to act on it that prevails in the circumstances. The account traces to W. D. Ross, *The Right and the Good*, (Oxford, 1930) ch 2.

I wish to remain agnostic about the plausibility of Holton's model with respect to moral rules. I am not sure that there are rules, strictly speaking, in the moral domain, and I am not sure that we can model the way we think about exceptions to moral rules by assuming that they all contain *unless* clauses that quantify over all other moral rules.¹⁹ Holton's model does seem more plausible, however, in the legal context, and precisely because it is different from the moral one, in two crucial respects: First, legal rules, unlike moral ones, typically have canonical formulations. Rules become part of the law by way of some enactment or other; they are enacted by a legislature or some administrative agency or, in the common-law system, by way of legal precedents decided by higher courts. Generally speaking, rules in the law, even rules at high levels of generality, tend to have some canonical, legally recognized, formulation. Second, and more important, exceptions to legal rules require authoritative determination, usually by the courts. Exceptions to legal rules are not added by reason; they are added by authoritative judicial (and sometimes legislative) decisions. Thus, in the legal case, we can interpret the rules as implicitly containing a *That's it* clause, with the additional premise that the courts would normally have the legal power to revisit the *That's it* clause and add an exception to the rule that had not been previously recognized. The idea here is that courts typically have a kind of quasi-legislative power to add items to any list of an *unless* clause in a legal rule, thus *modifying* the conditions for the rule's applicability. In other words, precisely because, in the legal case, any addition to a list of *unless* items is a form of institutional, authoritative modification of the original rule, the *That's it* suppression makes some sense; there is some legal reality to it. Until a court authoritatively determines otherwise, legal actors must presume that existing exceptions to a given rule are not superseded by a legal argument; *That's it*, in the legal case, attaches to a legal rule's list of recognized exceptions, as an institutional matter. Only courts or other legal authorities have the power to determine that, in a given case, *That's not it*.

Be this as it may, I share a crucial intuition with Holton, at least with respect to legal rules. The intuition is that it would be wrong to infer from the open-endedness of

¹⁹ The idea that there are rules in the moral domain is famously challenged by J. Dancy's particularism. And of course, though for very different reasons, by act-utilitarianism as well. As I say in the text, nothing that I say here about legal rules should be assumed to apply to moral rules as well.

possible justified exceptions to legal rules that the rules must implicitly contain some typicality premises. Some exceptions to this may exist. There might be some legal norms that are taken to capture typical or paradigmatic types of cases, defeasible in ways that typicality premises usually are.²⁰ Mostly, however, we understand legal rules to determine some results unless there is some conflict that justifies an exception. And even then, the exception would have no legal significance without being authoritatively recognized as such. A clear example of this way of thinking comes from cases in which the application of a legal rule in some peculiar set of circumstances yields patently absurd or unacceptable results. In some cases, courts would decide that the rule does not apply and thus create an exception, but, in others, they decline to modify the rule and apply it anyway (sometimes expressing the hope that the legislature will intervene and enact the requisite modification).

So what does the open-endedness of possible exceptions to legal rules amount to? In the legal case, at least, the answer is that it is essentially a matter of legal authority. Legal systems need to assign the authority to modify rules and adapt them to varying circumstances. Why? Because legal rules are enacted for reasons, aiming to achieve some particular purposes, and it may happen that the reasons for a given legal norm are either not well served by applying the law in a particular case of its putative application, or else they conflict with other reasons that apply. If law is to be responsive to reasons, it must have some mechanism that allows for adaptations and modifications of rules; somebody should have the authority to modify the laws when the relevant reasons call for a modification. In common-law systems, this authority is usually vested in the courts. Courts get the power to determine which exceptions to legal rules to recognize, and their power in this respect is, indeed, pretty much open-ended. If we want to call this aspect of legal reasoning defeasibility, so be it. It is a different kind of defeasibility, however, from what I will describe in the next section.

²⁰ Lest somebody assumes that I have in mind legal principles *ala* Dworkin, I should clarify that I do not. Elsewhere I argued that there are no such things as legal principles in the sense Dworkin has argued for; see my *Philosophy of Law* (Princeton, 2011), ch 4. I am thinking about some well-entrenched legal “tests,” such as foreseeability for proximate causation in tort law. The way the foreseeability test is used by courts does lend itself to be the kind of rule that has an implicit typicality premise embedded, one that is recognized to fail on occasion.

3. Conflicting Defeats in Law

Elsewhere I argued in some detail that an act of legislation is a collective speech act, and that the legal content of enacted law is the content communicated by the speech act in question.²¹ As with other cases of linguistic communication, the content conveyed by a speaker is often pragmatically enriched content. The hearer's ability to grasp the pragmatically enriched content involves a *defeasible inference* from the semantic content of the expression used, its syntax, the factual context that is common knowledge between speaker and hearer, and the relevant normative framework governing the conversation in question. We have already seen why pragmatic inferences tend to be defeasible, and I argued that some pragmatic defeats are of the conflicting kind, rendering the communicated content genuinely indeterminate. It is time to give some legal examples now and see how these pragmatic forms of defeasibility may render legal content indeterminate. I will start with two legal cases.

Consider, first, the famous case of *FDA v. Brown & Williamson Tobacco Corp.*, decided by the U.S. Supreme Court in 2000.²² The litigation concerned the question of whether the Food and Drug Administration (FDA) has the legal authority to regulate tobacco products. The background of this case is rather complex. The FDA was reshaped by Congress in 1965 in an amendment to the Food, Drug and Cosmetic Act of 1938 (FDCA), giving the federal agency the authority to regulate, in the relevant section, any “articles (other than food) intended to affect the structure or any function of the body.” Needless to say, in the 1930s (or even the '60s, for that matter), nobody thought that tobacco products would fall within the ambit of the FDA's regulatory authority. And indeed, for decades, the FDA explicitly declined to assert any authority to regulate tobacco products. This changed in 1996, when the newly appointed director of the FDA changed course and declared that the FDA did have such authority, granted to it by the original FDCA.²³ However, the FDA faced a serious problem. Other sections of the FDCA made it clear that if tobacco products fall within the ambit of its regulatory

²¹ See my *The Language of Law*, ch 1.

²² 529 U.S. 120 (2000).

²³ For an historical account of the facts and circumstances that brought about this litigation, see T. Ruger, “The Story of *FDA v Brown and Williamson* (2000): The Norm of Agency Continuity” in *Statutory Interpretation Stories*, Eskridge et al (eds), (Foundation Press, 2011), 334.

authority, the FDA may have no choice but to prohibit the sale of tobacco products entirely.²⁴ Obviously, the FDA wanted to avoid this legal result, and argued in its briefs that it had legal ways to avoid banning the sale of tobacco entirely, even if regulation of tobacco products were to fall within its jurisdiction and were found to be harmful. The argument about this particular point was rather ingenious, but a bit too much so. The courts did not buy it, and proceeded to examine the legal question with the assumption that if they found the FDA to have the authority to regulate tobacco products, the FDA would have little choice but to prohibit their sale altogether.

So far, so good; it seems to be an easy case, at least from a linguistic perspective. The relevant section of the law gives the FDA a very wide authority to regulate any product that is “intended to affect the structure or any function of the body.” Surely cigarettes and other tobacco products fall within the extension of this expression. Tobacco products certainly intend to “affect the function of the body”; it is what they are made to do. But this is not the whole story. Over the years, between 1965 and 1996, Congress enacted six pieces of legislation explicitly regulating the sale and advertisement of tobacco products. These laws imposed various restrictions on the ways in which cigarettes and other tobacco products can be sold, prohibiting their sale to minors, restricting advertisement in mass media and imposing various labeling requirements. Now, evidently, all these laws implicate that the sale of tobacco products, albeit restricted, is not illegal. If Congress says that you can only sell a product X if it is labeled as Y, it clearly implicates that if the product is labeled as Y you may go ahead and sell it. Or if Congress says that you may not sell X to minors, it clearly implicates that you are allowed to sell X to adults.

The time sequence is important here. Suppose, for the sake of the argument, that there is no doubt that the original FDCA, from 1965, actually asserts – regardless of what Congress may or may not have intended – that tobacco products fall within the

²⁴ That is so, because the FDCA created two separate regulatory schemes for products (other than foods and cosmetics) that purport to have some medicinal benefits and those that do not. If a product is claimed by its manufacturer to have some medicinal benefit, the FDA needs to conduct a series of hearings, based on scientific research, to determine whether to approve the product or not. However, if the manufacturer does not claim any medicinal benefit with respect to a given product, which is clearly the case with tobacco products, and the product proves to be harmful, the FDA must prohibit its sale.

jurisdiction of the FDA, and, consequently, due to other parts of the statute, and the fact that tobacco products are undeniably harmful, it follows that the FDA must ban their sale. Now, it is a widely accepted principle of democratic legislation that Congress has the authority to amend or modify its previous laws. Therefore, the real question here is whether the later pieces of legislation, regulating the sale and advertisement of tobacco products, actually withdrew the putative authority of the FDA to ban tobacco products or not. Remember that these later pieces of legislation would seem to clearly implicate, even if they do not quite say so, that the sale of tobacco products is legal. But here is exactly where we run into trouble. The contextual background of these later pieces of legislation is muddled; when they were enacted, the FDA *did not claim* authority to regulate tobacco products. The question of whether the FDA might claim such authority, and whether it would have it, was hanging in the air, suspected by legislators over the relevant years but not established. In fact, Congress attempted to pass legislation in both directions, explicitly granting and explicitly withdrawing such authority, that failed to pass. In other words, the contextual background was one of uncertainty. Given this background, the inference from the enactments regulating the sale and advertisement of tobacco to the implication that the FDA has no authority to ban the sale is defeated. But notice that the defeat is not of a rebutting kind. The contextual uncertainty does not rebut the putative implication of those laws; it only renders them inherently uncertain.

A second example is *West Virginia University Hospitals v. Casey*,²⁵ a textbook case for textualism in action, in which Justice Scalia, writing for the court's majority opinion, reasoned to the decision on linguistic grounds. *Casey* stems from a civil suit filed by the hospital of West Virginia University against the Medicaid system adopted by the state of Pennsylvania, concerning remuneration of Medicaid costs for out-of-state services, which, the hospital claimed, was in violation of federal and constitutional law. The hospital prevailed at the trial court and pursuant to the provision of a federal statute was awarded the cost of its attorney's fees, which included the cost of expert fees paid by the attorneys to their nonlegal experts. The case went to the Supreme Court only on this last point: The defendant argued that expert fees are not included within the expression of

²⁵ 499 U.S. 83 (1991).

the federal statute allowing the court to award “a reasonable attorney’s fee.” The experts are not attorneys. Justice Scalia, speaking for the majority, agreed, but not because the ordinary meaning of the expression under consideration would naturally exclude the cost of experts to the attorneys in question. It is pretty clear from the decision that none of the justices on the Supreme Court thought that the matter could be settled by simply looking at the semantic content of the relevant expression. Scalia’s argument was based on a kind of pragmatic inference: the fact that, in many other acts of Congress (though not all of them) awarding attorney’s fees to a prevailing party in civil litigation, the act explicitly mentions attorney’s fees *and* expert witness fees. *Ergo*, if Congress chose to use only the expression “attorney’s fees” without the addition of expert fees, the latter were meant to be excluded.

So far, so good; the pragmatic inference sounds very reasonable. But, as the dissenting opinion made clear, this pragmatic inference ignores the fact that Congress clearly “intended to make prevailing plaintiffs whole.” There is, indeed, a great deal in the context of such legislation to suggest that the purpose of the legislative provision was to make sure that a plaintiff in a civil lawsuit who prevails in trial can recover its reasonable litigation costs from the losing party. And, of course, it is widely known that litigation costs often include much more than the legal fees paid to the attorneys; attorneys often employ various experts in the service of the litigation. In other words, what the dissenting justices argued here is that the pragmatic inference Scalia had in mind is defeated by the context and the purpose of the legislative speech act in this case. Given the reasons the legislature had in enacting this provision, the inference is rebutted.²⁶

I think that the dissenting opinion is right, but only up to a point. The dissent is right to conclude that the implication is defeated in this case, but I think that it is wrong to assume that the defeat is of the rebutting kind, as they must have thought. The superseding premise here, concerning the contextual background and putative purpose of the legislative act, is only part of the story. The other part is the one Scalia emphasized, namely, that when the legislature wanted to shift expert fees as well, it explicitly said so.

²⁶ Truth to be told, the dissent’s view has been later vindicated by Congress itself, since it swiftly enacted an amendment to the act, overruling the decision in *Casey*.

So the background here is a bit more muddled than the dissent's view would have it. I would suggest that the defeat here is, again, of the conflicting kind. It leaves the implication defeated but not rebutted. From a linguistic perspective, the result is indeterminacy.

Some readers may resist this last conclusion. People may think that if we just knew enough about the circumstances of the legislation and the intentions of those involved in its production, we would know the right answer. But this is wrong. First, remember that it is an objectified sense of uptake that we must consider here. The question is not directly about the intentions or purposes of various agents involved in the process of legislation; it is a question about pragmatic commitments. Knowing all the relevant contextual background, what would a reasonable hearer infer? The simple point is that, even if all the relevant background is known, the inference to some particular aspects of the communicated content might remain genuinely indeterminate.²⁷ Suppose, for example, that the truth of the matter about the legislation of the act in *Casey* is that the drafters just made a mistake; they forgot that other pieces of similar legislation refer explicitly to expert fees. But suppose that they actually intended to include them. Still, this would not necessarily settle the issue. There is no contradiction in suggesting that, by saying that X in circumstances C, the speaker intended to convey that Y, but failed to do so.

More generally, it is very important not to conflate two distinct questions here: What are the *criteria of success* for an act of communication? And what is the *content that has actually been communicated* on an occasion of speech? It is very plausible to assume that an act of communication (fully) succeeds when, and only when, the hearer has fully grasped the content that the speaker intended to convey by his or her utterance. In this sense, it is clearly true that the criteria of success for an act of communication are reducible to the relation between the content that the speaker intended to convey and the hearer's actual uptake. Therefore, even in case of doubt, we can say that there is a truth of

²⁷ My argument in the text relies on the distinction that I have defended on several occasions, between questions pertaining to what the law says or implicates, and questions about the interpretation of the law when it is not clear enough what the law says or, not clear how what the law says determines the issue at hand. Legislative intent may well be relevant to modes of statutory interpretation, I am not denying that.

the matter about success or failure, determined by facts about communication intentions and hearer's uptake. But things look different when you focus on the question of what the speaker has actually managed to communicate or, as we discussed earlier, the content that the speaker is committed to, given what she said in the context of the expression. Here, the question itself presupposes some objective standpoint. We are simply not asking what the speaker intended to convey, but what she actually did convey. What was said and/or implicated by the speaker's utterance in a given context of speech? In answer to this question, we cannot rely exclusively on the intentions of the speaker, because speakers may fail to communicate what they had intended to. And we cannot rely on the actual uptake of the hearer for the same reason, that hearers may fail, on particular occasions, to register what they should have inferred from the utterance in the relevant context. In other words, the question itself presupposes that we have in mind some objectified sense of uptake. What is actually said and/or implicated by an utterance in a given context of speech is determined by reasonable uptake: What is the content that a reasonable hearer, knowing all the relevant contextual background of the utterance, would have inferred from what was uttered in that context? Under normal circumstances, if the communication was successful, then a reasonable hearer would have inferred exactly what the speaker intended to convey. But this is only the limiting case; things might turn out differently on actual occasions. What was intended to be said is not necessarily what has been said.²⁸

Given this objective framework of pragmatic inferences, I hope it becomes clearer why conflicting defeats result in genuine indeterminacy and not some factual uncertainty. When we face some uncertainty about a result we would be warranted in assuming that more information might remove the doubt. If I am not quite sure about what *you* meant, I can ask you and you can clarify your intention. But if the relevant question is not about what you meant, if it is not about your intentions but about what you actually managed to convey, or the content you are committed to given what you said, then a quest for additional clarification is misplaced. Indeterminacy is inquiry resistant. To be sure, I am not claiming that in the legal case indeterminacy entails that there is no right answer to

²⁸ I defended an objective conception of asserted content in my *The Language of Law*, chs 1 and 5.

the legal problem that the court needs to resolve. Many factors can bear on the legally correct solution to a given case, including, potentially, deference to the subjective intentions or purposes of the law makers. The only warranted conclusion here is that the judicial decision aiming to resolve the issue is not going to be based on what the law says or implicates. In the cases we have been discussing, what the law says or implicates is legally indeterminate, and thus any judicial decision is going to amount to a modification of the law, perhaps creating new law if the decision is followed as a precedent. And this, of course, happens on countless occasions.²⁹

²⁹ I am grateful to the participants of the conference on Defeasibility in Law, held at Goethe University in Frankfurt (March 2015) for helpful comments on a draft of this paper, particularly to my commentator, Klaus Gunther.