Deontic Modality, Generically

Federico L. G. Faroldi

Abstract

This position paper aims to explore some preliminary suggestions to develop a theory of deontic modalities under a generic understanding. I suggest, for instance, that a sentence such as 'Everyone ought to pay taxes' is true just in case the generic (deontically relevant) individual pays taxes. Different degrees of genericity are explored, without assuming too much about a specific theory of genericity. I argue that such an analysis captures our intuitions about exceptions and the general character of deontic claims better than classical approaches based on possible-world semantics and than defeasibility-based approaches, while remaining within a broadly deductive framework.

[...] law can never issue an injunction binding on all which really embodies what is best for each: it cannot prescribe with perfect accuracy what is good and right for each member of the community at any one time. The differences of human personality, the variety of men's activities and the inevitable unsettlement attending all human experience make it impossible for any art whatsoever to issue unqualified rules holding good on all questions at all times [...] [one] will lay down laws in general form [for the majority, roughly meeting the cases of individuals [...] under average circumstances.

Plato *Politicus*, 294a–b.

[...] all law is universal but about some things it is not possible to make a universal statement which shall be correct. In those cases, then, in which it is necessary to speak universally, but not possible to do so correctly, the law takes the usual case, though it is not ignorant of the possibility of error.

Aristotle *Nicomachean Ethics*, 1137a–b.

1 Introduction and Scope

Notwithstanding early informal ideas in Plato and Aristotle, for the better part of Western Philosophy normative reasoning has been dealt with in a largely deductive framework, alongside other kinds of reasoning. Even since its modern beginnings, as an example, deontic logic has been mostly based on variations of deductive modal logic.

In the past few decades, however, this has been changing. A comprehensive recent resource on legal reasoning and argumentation, for instance, seems to recognize as undisputed the fact that the deductive model (of normative reasoning, i.e. the one which makes use of a deductive logic, formally or informally) "is not applicable to the broad

Filosofiska Notiser, Årgång 6, Nr 1, 2019, 133–155

majority of cases in law, where the arguments used to support or attack a conclusion are defeasible (subject to exceptions) [4, p. ix].” The main dilemma here seems to be that either normative (legal, in this case) reasoning does not have a (well-understood, traditional) logic, or if it does, such a logic does not apply but in a minority of cases. Contemporary research in normative argumentation goes in two directions: it seeks either to apply a non-classical logic, and in particular a defeasible logical approach, or to move within a more pragmatic, dialectical framework, “which uses burdens and standards of proof, along with other devices, to take the context of use of an argument in a specific setting (e.g., in a trial in a specific legal system) into account (p. x).”

In this paper, I offer some reasons to consider a third option: a generic understanding of normative reasoning. It works as a *via media*: it keeps the deductive power of classical logic (or something very close to it), while at the same time taking into account certain pragmatic issues, which, however, are not to be considered as explicit, fixed once-and-for-all, known-in-advance exceptions.

I will argue that such an analysis captures our intuitions about exceptions and the general character of deontic claims better than (i) classical approaches based on possible-world semantics[1] and than (ii) defeasibility-based approaches, while remaining within a deductive framework.[2]

Normative reasoning is usually performed via steps that contain normative, and especially deontic, expressions, or at least it can be reformulated in terms of normative or deontic expressions. I’ll start from these: in this paper I put forward some suggestions to develop a novel theory of (some) deontic modalities under a generic understanding. What is genericity? Suppose you want to show your little niece why the Pythagorean theorem holds (she’s a smart curious kid). You will probably start along the lines of ‘Let ABC be a right triangle. Then...’. What you did was asking her to consider an arbitrary triangle, with the only constraint that it be right: the length of the sides, the area, the possible colors of the lines of its representation, etc. do not matter. From the fact that you can show that the Pythagorean theorem holds for *such an* arbitrary triangle, you can conclude that the Pythagorean theorem holds for *every* right triangle. However, were you to find a right triangle for which the square of the hypotenuse is *not* equal to the sum of the squares of the other two sides, the universally quantified sentence would be false. So far, so good, nothing new.

Things become slightly trickier when we exit the mathematical domain. Let’s now consider generic language. Take the sentence: ‘Dogs have four legs’. We immediately take the generic sentence to be intuitively true, and we are prone to discard obvious counterexamples, e.g. dogs who have fewer than four legs as a result of an accident, dogs who have fewer than four legs because of birth defects, etc. But why? We consider these counterexamples irrelevant to the fact that dogs have four legs. Lots of ink has been spent to come up with theories of generics.[3]

Yet, there is no standard or satisfactory theory of the semantics of generic statements, even when we restrict ourselves to properties predicated of kinds, e.g. K(ind)

[1] For a historical survey and introduction, see [19], and for a state of the art survey, see [15]. For a more applied perspective to the legal domain, see [28].

[2] For an introduction to defeasible reasoning, see [24]. For an introduction and survey of various defeasible approaches to deontic logic, see [30]. For a more applied perspective to the legal domain, see [11].

[3] For a recent survey, see [29].

has P(property). We can group extant proposals in two big families: *frequentist* ones and *conceptual* ones. Here's a very rough sketch.

Frequentist theories hold that a generic is true just in case the property predicated of the kind holds with a high frequency for its members.[4] Such theories have difficulties to account for the typical resilience to exceptions that generics display, and introduce some kind of saliency restriction of the domain of quantification.

Conceptual theories hold that a generic is true just in case in virtue of being of that kind, instances display that property. Leslie for instance maintains that generics are tied to a "default mode of generalization" with the feature of singling out striking properties. This would explain why we say that a generic such as 'Mosquitos carry malaria' is true, while most mosquitos don't carry malaria, i.e. where instances of a true generic are almost all false. This strategy is also used to explain why generics seem to be more sensitive to positive counterinstances ('birds are female' seems odd because being male is a positive counterinstance of being female, something that is not the case with 'birds lay eggs': birds who don't lay eggs don't do anything else, e.g. they don't give birth to live young.) Moreover, some proponents of the conceptual view maintain that there are different types of generics (statistical, characteristic, striking), standing for different kinds of psychological relations between concepts. Such a claim would be substantiated by psychological research on the difference between characteristic properties, statistical properties, and "striking" properties, that can be expressed with generic sentences.[5]

Interestingly enough, some of these examples can be used to argue that generics require (at least) an intensional semantics (if not hyperintensional): 'Ducks lay eggs' is extensionally equivalent to 'Ducks are female', in the sense that all ducks who lay eggs are female (and, let's assume, also viceversa); yet, we understand the former as true and the latter as false. One can even say that 'Ducks lay eggs' is intensionally equivalent to 'Ducks are female', in the sense that necessarily, all ducks who lay eggs are female (and viceversa).

For the purposes of this paper, however, we do not have to choose a specific theory of generics. In fact, all theories, be they frequentist or conceptual, agree on two crucial points: the need to explain the generics' tolerance to exceptions, and the need for some notion or other of relevance. This is enough for our aims.

How does all this relate to deontic modalities? I am going to suggest, very roughly, that a sentence such as 'Everyone ought to pay taxes' is true just in case the generic (deontically relevant) individual pays taxes.

More concretely, I will first analyze generically deontic claims involving a universal clause, like 'ought', rather than deontic claims involving an existential clause, like 'may', at the propositional level. The most used (logical) semantics for such deontic modals is of course a Kripke-style semantics, where 'p ought to be the case' is true at a state just in case p is true in *all* states (with a further restriction of the domain of states, where such a restriction is usually understood in terms of "states that are a deontic alternative to the state of evaluation" or "deontically perfect states"). In formal

[4]The best-developed theory is by Cohen, see e.g. [6, 7].

[5]For a survey and specific references, see [25]. For a representative, see [16].

linguistics, the most common semantics is of course a Kratzer-style semantics, which also employs states or worlds and, under certain assumptions, is equivalent to Kripke semantics.[6]

In Section 2, I will argue that deontic modals, and in particular 'ought', are best understood generically in light of the fact that such an approach fares better in taking care of exceptions and generalizability. In Section 3, I will sketch my proposal for a generic semantics of deontic modals. In Section 4, I will start exploring formal theories of 'ought' understood generically. While I argued elsewhere that possible-world semantics is too coarse-grained for deontic modals,[7] I'll develop proposals along the following lines:

(1) $w \vDash \mathcal{O}\phi$ iff $w^* \vDash \phi$,

in words: ϕ is obligatory at state w, just in case ϕ is true at w^*, the generic obligatory state and ϕ is a sentence in a formal language.

Some questions we'll face are the following: What does it mean for a state to be generic? What is its relation to ordinary states? How does this generic take interact with quantification when ϕ is first-order? What does it take for individuals to exist or to satisfy a formula at a generic obligatory state?

2 Deontic Modals Are Generic

Suppose there is an obligation to ϕ, for instance, the obligation to spend time with family and friends for Thanksgiving. The most common idea to analyze such a modally-flavored statement semantically is to check a subset of the (perhaps contextually) relevant (in some respects) alternative situations to verify that the prejacent (ϕ, in our case "spend time with family and friends") holds at those situations.[8]

"In some respects" here may or may not be crucial to the philosophical and logical significance of such a semantic analysis. It is well-known that in Kripke-style semantics such relevant alternative situations are captured by a binary "accessibility" relation R.[9] Equally well-known is that conditions imposed on R, such as reflexivity, seriality, etc. determine the structure of frames and correspond to certain axioms. Less-well understood is the philosophical and explanatory contribution of such a relation at least for certain modalities, such as deontic modality. In particular it seems that a deontic accessibility relation is meant to capture either "ideal worlds" or, with a rather unsatisfactory tautological flavor, worlds where everything that is obligatory at the point of evaluation is the case.

[6]There are semantics, such as [1] and [14, 13] where 'ought' gets an existential interpretation, whereas 'may' gets a universal interpretation. Mutatis mutandis, the considerations I make will apply to those semantic proposals too.

[7]Cf. [9], [8].

[8]Let's suppose that there are some sort of truthmakers for deontic claims. At this point, no further specification on truthmakers is needed — for the time being, we can take truthmakers to be possible worlds, just to work in a familiar environment. Nothing in the approach depends substantially on the choice of the framework, although of course if the underlying framework is changed, obvious adaptations are needed. In what follows I use interchangeably 'world', 'state', and 'situation'.

[9]Kratzer-style formal semantics follows a different set-up, but the two approaches can be proven equivalent — see e.g. [17], [9].

To go back to the Thanksgiving example, the standard analysis suggests that since in all deontically relevant (ideal?) situations for Thanksgiving time is spent with friends and family, for Thanksgiving it is obligatory to spend time with friends and family.

But suppose you caught a contagious disease. Surely it's better you stay in isolation, even if it is Thanksgiving. In such a situation, you don't spend time with family and friends. Now, from the point of view of this semantic analysis, the crux of the matter is obviously whether the situation where you caught a contagious disease is in the set of situations accessible from the actual one. This can be denied on a lot of grounds. For instance, if one goes along with the ideality analysis, one could argue that a situation where you caught a contagious disease is definitely not ideal, and therefore it is excluded from the set of situations accessible from the actual one, thus bypassing the problem and keeping $\mathcal{O}\phi$ true at the actual situation. Whether such an obvious criticism succeeds is highly doubtful: in fact, such a take would make sense only from a perfectly-good-god-like point of view, where alternative situations are objectively ideal, as it were. However, the job of a good semantical analysis is to provide the truth-conditions or the meaning of an expression — a deontic modal, in this case —, not to tell us when something is substantively good: this is the job of one's background normative theory, not of semantics. More plausibly, in fact, obligations continue to hold even when the situation turns out pretty badly and therefore one cannot conflate alternative and ideal situations in the possible-world semantical analysis. Everyone ought to pay taxes, regardless of whether this year you also took a massive financial hit, the roof of your house collapsed, and your car broke.

While possible-world semantics for deontic modality has been challenged on many grounds,[10] many extant criticisms are irrelevant to some issues that are important but less frequently considered, namely *exceptions* and *generalizability*, to which we now turn.

2.1 Exceptions

What seems to be going on in the Thanksgiving case is that we have found an exception to an obligation we formulated as covering *all* cases. Such an exception is, in this particular case, relevant to the obligation at hand, somewhat restricting its scope.

However, there are exceptions we discard. Consider a moral tenet according to which one ought not to kill, and consider one of those complicated scenarios ethicists like so much, one that is relevant to one's background moral theory and is supposed to create difficult situations or counterexamples. Perhaps it is killing one versus killing many, perhaps it is killing versus letting die. More simply, let's consider a less charged provision, which may be expressed as 'Everyone ought to pay taxes'.[11]

Now, there are no explicit exceptions stated in the provision itself, but it is easy to come up with exceptions. For instance it is conceivable that the king or a foreign ambassador don't have to pay taxes. It is clear that for such simple cases (usually) the legal system has explicit additional provisions in place (which, however, might be of less importance — e.g. our original provision was of constitutional standing, whereas

[10] To recap a few: see [36] and see also Ch. 1 of [9].

[11] This example has a legal background, but the point I wish to make holds more generally.

it might be that the rules regulating the fiscal status of foreign diplomats are of ordinary standing, and therefore less important). However, in less explicitly regulated parts of our normative life it is not so: we understand that the fact that the king or a foreign ambassador don't pay taxes, while being a direct exception to 'Everyone ought to pay taxes', by no means makes 'Everyone ought to pay taxes' less valid or true.

Again, the universal analysis of 'ought' does not seem completely satisfactory. In fact, assuming varying domains of individuals (constant domains are conceptually problematic in a deontic setting), it is unclear why we should exclude worlds where the king or foreign ambassadors are not in the domains of the ideal worlds (in other words, those accessible from the point of evaluation), just for the mere fact that the king or foreign ambassadors don't pay taxes and their existence would make 'Everyone pays taxes' in those worlds false and therefore 'Everyone ought to pay taxes' false at the world of evaluation. This not only would be ad hoc, but also patently inadequate to represent the situation at hand, for the king or foreign ambassadors don't cease to exist in a given world just because they don't pay taxes. On the other hand, it is clearly a fact of our normative lives that we don't go around stating all the time a universal provision and a whole range of exceptions to it.

Of course the problem of exceptions is not new. How to deal with exceptions in deontic reasoning?

A widely used strategy is that of defeasibility (usually defeasibility at the level of inferences, not of object language claims). Let's introduce defeasible approaches in general and then compare them to a generic strategy.

2.1.1 Exceptions: Generic vs defeasible approaches

Defeasible inferences are often described in terms of "typically" or "normally" clauses:[12] Birds generally fly. Tweety is a bird. Therefore (most likely, defeasibly) Tweety flies. This schema is also sometimes called defeasible modus ponens: it is defeasible because there might be exceptions to 'Birds fly', which is indeed not understood as a universally quantified statement. For instance, were Tweety a penguin, she wouldn't fly.[13]

More specifically,

> defeasibility assumes that general propositions are defaults, which are meant to govern most cases or the normal cases. Thus, we can consistently endorse such propositions and deny that they apply to certain cases: the exception serves the rule, or at least it does not compromise the rule.

[12] For such a characterization and an introduction to defeasible reasoning, see e.g. [24].

[13] There is a tendency in the literature to speak indifferently of defeasible and non-monotonic reasoning. Such a carefree attitude seems to me quite misguided, once one tries to give more precise definitions to the terms in question. In particular, but quite roughly, it seems that non-monotonic reasoning is defined with reference to a theory where a conclusion can change once one adds (explicitly, as it were) further premises. However, if one characterizes a defeasible inference as an inference that holds normally, or in most cases, it seems that the exceptions are either already known, or potentially knowable upon discovery but already there, as it were, rather than explicitly added after the fact. Such an account of defeasible inference is best understood as just classical inferences that are *defeated* by the non-normal cases (which one is supposed to know or able to know, given one has a notion of normality in the first place), rather than full-blown *defeasible* inferences.

> To deal with an anomalous case on a defeasibility strategy, we do not abandon the default or change its formulation, but instead we assume that the default's operation is limited on grounds that are different from those that support the use of the default itself. As we saw in the previous example, these grounds may provide an argument that undercuts or rebuts the argument warranted by the default [31, p. 342].

There are several ways a generic approach differs from a defeasible approach. First and most obviously, defeasible premises hold generally or normally, and not, *sensu stricto*, generically. In other words, we expect that, in most cases that property holds, that a significant number of individuals enjoys that property. This is, in a sense, a mere frequentist claim. On the other hand, while there is at least a theory of generic sentences that makes use of a (sophisticated) probabilistic approach understood in frequentist terms, it is quite accepted that there can be true generics of which most or even all instances are false. In this sense, there may be more exceptions than regular cases. The concept of generality is therefore not the same as the concept of genericity: one could even understand 'generally' as a descriptive claim, while reading some evaluativity in 'generically'.

Secondly, let's focus on "exceptions". With defeasible inferences, one establishes a conclusion which, absent evidence to the contrary, holds. If one comes across new evidence which constitutes an exception to what generally holds, then one revises the previously acquired knowledge. In a certain sense, inferences are provisional, because one can always come across new information: exceptions are, in a sense, *post hoc*.[14]

Generic reasoning has, instead, the certainty of deductive reasoning: inferences are not provisional but definitive; exceptions, to the extent that they happen, are in a sense *ante hoc*: they occur before we make an inference, and serve us as a guide to refine the claim we are making.

Despite certain theories of generics that take into account exceptions, one may roughly understand a generic sentence (to the extent it can be understood quantificationally) as a universally quantified sentence on a restricted domain of quantification, where the domain is restricted perhaps by taking into account all the exceptions to the claim in question. As a concrete example, let's take 'Dogs have four legs': it's clear that this generic sentence cannot be understood as a sentence universally quantifying over all dogs. In fact, the generic sentence is true even if all instances aren't. However, we could exclude all the "abnormal" dogs, leg-wise, and afterwards understand the generic sentence as a universally quantified sentence over all the "normal" dogs.

Third, and more structurally, defeasible and non-monotonic approaches usually work at the level of inferences: they provide mechanisms to retract conclusions already established in light of new evidence. Applied to the generic case, this manages

[14]A helpful analysis of defeasibility in the normative domain and its peculiarity is offered by [34]. They distinguish between *factual defeasibility*, which is simply a fact that materially implies the falsity of a default and *overridden defeasibility*, which is split in *strong* and *weak*. Overridden defeasibility has to do with more specific defaults (from 'birds fly' to 'penguins don't fly'), Strong overridden defeasibility has to do with the case when the first default is overridden full stop; weak when the first default can be reinstated in case the second is not appliable (any more). These distinctions, however, do not seem to help with the issues raised in this section.

to deal with exceptions. Defeasible and non-monotonic approaches, however, especially in their default realizations (cf e.g. the work of Horty for the deontic case, such as [23]) are about inferences and rules. But inferences and rules do not have truth-conditions, they are not object-language claims. The overarching evidence, on the contrary, both linguistic and intuitive, is that the kind of generic sentences we employ, both in normative reasoning and outside of it, are (i) object-language claims; and (ii) have truth-conditions. For instance they can be embedded under modals, or iterated: this is strong evidence for the truth-conditional, rather than for the inferential approach.

This is of course an extremely oversimplified account of generic sentences we considered only to illustrate how generic reasoning differs from defeasible reasoning.

However, there is no theoretical reason why a generic and a defeasible approach cannot be combined. The underlying logic can in fact remain defeasible, and one may add the generic component on top of it.[15] On the other hand, one has to be assured that the generic approach cannot be reduced to the defeasible one. In my mind there is a strong theoretical obstacle to such a reductive project. Defeasible reasoning is supposed to work for typical, normal, usual, most cases (cf. e.g. [33, Ch. 1]) but it's quite well-known that there are true generic sentences which are exemplified by a minority, or even no actual cases (and viceversa, false generic sentences that are exemplified by all cases cf. for some examples [29]). Generic reasoning seems to require a notion of relevance which is not reducible to extensional or intensional notions such as commonality or normality.

Another interesting objection has to do with the dynamics of a normative system. One might think that a defeasible approach is better equipped to handle — if we consider legal systems just because in a sense they are more tractable — phenomena like norm derogation, revision, etc. via operations either on the knowledge base or on the additional information or exceptions one may come across. If indeed generic reasoning is to be understood as I roughly characterized it, in terms of deductive reasoning, where all the exceptions are already taken into account from the beginning, then it seems that it is less adequate to model these more dynamic aspects, such as adding new norms, or having a hierarchy of norms, or derogate norms: all phenomena that are genuinely novel and whose content cannot be established in advance. However, broadly speaking, in several prominent accounts of arbitrary reasoning on the market (cf [12], [20]) there is a notion of dependency. Arbitrary objects can depend on other arbitrary objects (subject to certain conditions, such as, for Fine, well-foundedness): for a very simple example, consider when in mathematics one introduces an arbitrary object x and then another arbitrary object y such that $y = f(x)$: clearly y depends on x. In a similar vein this notion of dependency can be used to model the more dynamic aspects of normative

[15] In fact, albeit very loosely, this seems to be done in [3] and [2], which is the only approach I could find who mention generics in a deontic logic context. What's going on here is that they have a non-monotonic logic (common-sense entailment), which the authors argue can handle deontic generics such as "Soldiers should obey". The authors, however, do not argue that deontic statements are to be analyzed generically, but they seem to conflate 'generic' with 'defeasible': "The connective $>$ is a doxastic, nonmoral, generic conditional: $A > B$ means that, where A holds, B normally holds too (p. 165)." Such a normality connective is combined with a deontic operator to form a new connective $>_O$ that is used to express a prima facie conditional obligation, which is evaluated relative to a new accessibility relation which assigns to each world and proposition a set of "good-and-simple" worlds, where worlds are "simple in that moral issues in them are, in a sense, one-dimensional. No moral complications arise; there are no conflicting obligations (p. 165)".

systems.

Even admitting that a generic approach and a defeasible approach are able to deal with exceptions in an equally satisfying manner, one might think that a generic approach is to be preferred because the underlying logic does not have to be revised, and we remain in a deductive framework.[16] Of course one may prefer a defeasible approach exactly because it is non-classical, although it seems that the main reason why some theorists think that deontic reasoning is defeasible is exceptions. But if exceptions can be handled classically, then the main reason to go defeasible disappears.

On this last note we ponder how the generic approach can be justified by considering an important trait of the normative domain: generalizability.

2.2 Generalizability

At least the two major players in the normative domain, ethics and the law, have ambition to be generalizable, in some cases with regard to content (which seems common to all ethical theories except for particularism) and in some cases with regard to application (especially in the law, with regard to claims of equality before the law, etc.).

This is well-captured by an underlying idea in generic semantics, namely that there may be no substantial semantic difference between a sentence such as 'The dog has four legs' and a sentence such as 'Dogs have four legs': in both cases there is a direct reference to the generic dog. The generic dog, for instance, is alive, but plausibly it won't die. And plausibly it is yet to be born, for future dogs will have four legs too. This is quite consistent with e.g. a widely accepted position in the philosophy of law according to which the law is valid and binding for yet-to-be-born citizens, and does not just bind the people alive at the time of its coming into existence. Such an idea seems not impossible, but clearly problematic in a possible-world semantic framework: in fact it would require exact information about which individuals there are; knowing that such and such law holds regardless of which individuals there are is not enough.

The generic approach takes care of this aspect, without being committed to a theory where modals are reduced to quantifiers over worlds.

The deductive reasoning proper of generic reasoning, as compared to defeasible approaches, seems to offer several advantages in the normative domain. It is straightforward, more predictable (for future cases, thus enhancing its action-guiding role) and replicable in a way that reasoning riddled with exceptions is not and (it may be argued) it is more suitable to the neutrality required by normative fields like ethics and the law, where by neutrality one refers to the widely held idea that e.g. everyone is equal before the law, and exceptions cannot be *ad personam*, but rather be grounded in general principles.

Objections If one thinks of generics as claims admitting of exceptions and is considering the normative domains, it is quite plausible to connect to generics *prima facie* obligations. *Prima facie* obligations are obligations that, at least under one theory, hold "other things being equal", i.e. if there aren't more important obligations, *prima facie* ones become actual. They can tolerate exceptions and normally they can conflict in a

[16] See [12] and [27] for technical details.

way that all-things-considered, or *pro toto*, obligations, can't. For instance [2] gave an account of *prima facie* obligations based on commonsense entailment, which is itself a way to give an account of generics in terms of a normalcy-based conditional and a non-monotonic notion of consequence. One, however, can object to the thesis of the present paper that *prima facie* obligations are the *only* kind of obligations that can be understood generically: all-things-considered, or *pro toto*, obligations, can't, exactly because they cannot conflict and they can't, if understood properly, tolerate exceptions. However, it is very implausible to hold that only *prima facie* obligations are to be understood generically, as in Asher and Bonevac's approach, because when they meet with an exception, they stop being obligations: they do not tolerate exceptions, after all.

Obviously, if there is just one kind of obligation that can be *prima facie* or *pro toto* depending on contextual factors, then those who think that *prima facie* obligations are basically generics (such as [2]) owe us an account of why obligations stop behaving like generics when they become *pro toto*, given that in such a theory context shift has substantial impact and cannot just be abstracted away.

On the contrary, if one admitted that there are two separate entities, *prima facie* and *pro toto* obligations, and only the former are generics but the latter cannot be, the obvious retort would be to deny that there are such things as two separate entities for a whole range of reasons: lack of linguistic or legal evidence for the dualist thesis, theoretical economy, etc.

However, there is a further consideration that is worth spelling out: that the thesis according to which "deontic modals are generic" is itself a generic: if it is true, some exceptions (i.e. that some (uses of) deontic modals are not generic) may very well be acceptable. We set aside this potentially problematic strategy for the reminder of this paper.

3 A general (generic) proposal

In previous work I suggested that the semantics of generic claims of the form 'K has P', where K is a kind, and P is a property, is the following:[17]

(2) 'K has P' is true *in virtue of* Pk, for an *arbitrary* $k \in K$.

An initial remark. Such a clause, via the use of *in virtue of*, is hyperintensional. This takes care, among others, of examples such as 'Ducks lay eggs' vs 'Ducks are female'. I will ignore this part in what follows to work in the familiar environment of possible worlds, which is intensional.

Given the scope of this paper, I am furthermore assuming that the *in virtue of* condition in (2) takes care of issues of relevance, as one would expect from an acceptable theory of generics.

[17] Cf. [10]. As it is usual in the recent literature, I'm considering only so-called characterizing generics, where the property P can be (potentially) ascribed to an individual of that kind, and not only to the kind itself (such as, for instance, something like 'Dinosaurs are extinct').

3.1 Obligation

We now want to give a generic semantics for obligation (to start with). Obviously not all obligations, when expressed in linguistic terms, take the form of a generic statement.

There are two immediately plausible ways to give obligation a generic semantics: first, we can try to understand whether a sentence of the form $\mathcal{O}\phi$ (where $\mathcal{O}$ is a sentential operator and ϕ is a sentence) can be analyzed as or at least translated to a generic statement of the form 'K has P', which is in turn analyzed via (2), in a more "covert" fashion.

Second, we can try to apply (2) directly (modulo minor adaptations) to a sentence of the form $\mathcal{O}\phi$, in a more "overt" fashion.

In the first case, suppose we try to analyze a sentence of the form $\mathcal{O}\phi$ via kinds. It is not obvious that this strategy works, unless the obligation is already expressed generically, as it were, such as in the following sentence 'Citizens ought to pay taxes'. Now according to (2) 'Citizens ought to pay taxes' is true in virtue of an arbitrary citizen having to pay taxes. Plausibly this strategy also works for universally quantified obligations such as "All citizens ought to pay taxes".[18]

However, it is not trivial to analyze all obligations in such a fashion. For instance, take 'Peter ought to pay taxes', and suppose this is an individual obligation. Let's also assume, for the sake of the argument, that genuine individual obligations can exist, rather than being always derivative of universal obligations via instantiation. A clear example of individual obligations are certain court pronouncements that, in civil law systems, order someone to do such and such, e.g. to pay a certain sum, or undertake certain work, etc. A way to analyze 'Peter ought to pay taxes' via kinds, i.e. via 'K has P', would be to understand K as being the situations, contexts, or worlds as a silent component. If we want to make it explicit, we would get something of the form 'Normally, ϕ', or 'In all situations, ϕ'. Such an addition not only sounds unnatural, but it may introduce extraneous elements in the semantic analysis. On the other hand, it is implausible to analyze 'Peter ought to pay taxes' by saying that an arbitrary Peter has to pay taxes, for it is not even clear we can make sense of 'an arbitrary Peter', for it cannot be an arbitrary individual that just contingently happens to be Peter.

However, we could still adopt an adapted version of (2):

(3) 'Peter ought to pay taxes' is true in virtue of an arbitrary (deontically relevant) world verifying 'Peter pays taxes'.

How an arbitrary deontically relevant world verifies a sentence such as 'Peter pays taxes' obviously depends on the the background theory of arbitrariness. For instance, the underlying formal theory might very well allow that individual obligations such as 'Peter ought to pay taxes' are analyzed classically, i.e. via universal quantification.

The "covert" generic analysis therefore is more general than the "overt" generic analysis of obligations.

[18] The difference between a generic statement that admits of exceptions and a universally quantified statement that plausibly doesn't, at least if interpreted strictly, is relegated to the background theory of arbitrary objects or arbitrary reference, which is beyond the scope of the present paper. A very simple account of this, however, would maintain that there is a (possibly pragmatic) restriction in the domain of K over which an arbitrary k may range.

There are at least two potential problems worth mentioning at this point. First, it is unclear whether such a covert analysis covers all the cases of arbitrary reasoning, or, in other terms, whether this is the right semantical approach to all sentences interpreted generically. Second, it is not clear that the inherited universal quantification is appropriate. In fact, consider the generic (non-deontic) statement ‘Dogs have four legs’. Now suppose that we have our domain $D = \{d_1, ..., d_n\}$ of ordinary dogs and we adjoin a domain $D^* = \{d_1{}^*, ..., d_n{}^*\}$ of arbitrary dogs, and we have a (possibly partial) function $\delta : D^* \rightarrow \wp D$ which associates to an arbitrary dog a set of ordinary dogs (i.e. a function from arbitrary dogs to sets of ordinary dogs). Now, one may think that if, given an arbitrary dog d^*, for which δ is defined, $\delta(d^*) = D$. But if this were the case, then it would be plausible to think that, if there is a three-legged dog in D, the generic reading of ‘Dogs have four legs’ will be false, given that not all values of an arbitrary dog have four legs. However, the function δ has to be clearly highly discriminating, and it has to be able to select ordinary individuals according to relevance. Without going much further than what is possible within this work, I suggest that the δ function encodes relevance criteria that are contextually determined — it is quite plausible to think, for instance, that even within a scientific discourse, different disciplines have different standards of how many and how relevant instances are needed to establish certain conclusions. Such built-in contextual features are not new for modal semantics, as the currently accepted formal semantics Kratzer-style analysis of modals shows.[19]

3.2 Permission

How to deal with permission? Can we analyze generically statements of the form “Everyone may purchase a ticket”, and variations of a deontically interpreted ‘may’, such as ‘allowed to’, ‘can’, perhaps ‘might’ (if there is a deontic reading of it), etc.?

In classical (possible-world based) semantic theories ‘may’ (and its sibilings) is understood existentially, i.e. as an existential quantifier on the set of deontically relevant worlds accessible from the point of evaluation. In other terms, permission is understood as being the dual of obligation.[20]

There are several well-known arguments and paradoxes pointing out that such a weak or implicit permission (as lack of obligation to the contrary, as it were) is only adequate for a rather small subset of real-life permissions. A more adequate (i.e. taking into account how permissions are used in everyday life) understanding of permission should be one of strong or explicit permission, and should be defined as a primitive, alongside an operator for obligation.[21]

I don’t take a stance here, but I merely suggest how to expand the technical proposals above to cover permission.

[19] There still seems to be an additional problem: how to account for statements that are generically true, but whose instances are all false? This is not a concern for the present purposes, because such problems are usually generated by the fact that one only considers actual or contingent instances. In our settings, however, it is clear that we have non-actual possibilities in mind too.

[20] For a survey of these proposals, see [19] and [26]. Notice there are proposals where permission is understood universally, cf. e.g. [1], [14].

[21] For a survey on the varieties of permission, cf. [18].

Implicit or weak permission is not a problem, insofar as one can define it as the dual of obligation, understood generically.[22]

Explicit or strong permission is slightly more problematic. However, one can see how, in a possible-world framework, it is possible to define a different accessibility relation, which defines a set of strongly permissible worlds within which arbitrary worlds can take their values.

An immediate objection has to do with the arguably different logical properties that obligation and permission operators have, which in the classical case are inherited from their respective quantifiers. Again, this is not a problem for implicit or weak permission; for explicit or strong permission, instead, this is not a problem introduced by the arbitrary framework, but is rather inherited from a primitive understanding of a permission operator. It is therefore to proponents of explicit or strong permission that this objection should be put forward.[23]

3.3 Some generics already have normative force

Interestingly enough, some generics, usually about social kinds, carry with them a certain normative force. Typical examples include: ‘Parents don’t abandon their children’, ‘Friends take care of each other’. Clearly, these don’t describe only what’s typical of the kind, but are also used to convey expectations or obligations, without being explicitly issuing obligations or commands. How does this play out with having a generic interpretation of obligations? Is this a further piece of (linguistic) evidence in favor of the generic analysis of deontic modality, or it is rather evidence for the opposite thesis, namely that generics need to be analyzed (partly) in terms of deontic claims? The latter suggestion is almost certainly false: there are clear counterexamples, i.e. generics (especially generics not about social kinds, it seems) which can hardly be interpreted in such an overt deontic way (e.g. ‘Mosquitos carry malaria’). With regard to the former claim, it seems more plausible to explain it with reference to a common *stipulative* or performative origin that some generics/arbitrary claims (‘Let *a* be a prime number...’) and some deontic claims have.

[22] This is already done, for instance, for the ϵ-operator, where one can introduce the quantifiers defining them in the following way: $\exists x A(x) =_{df} A(\epsilon x A)$, $\forall x A(x) =_{df} A(\epsilon x(\neg A))$. For an introduction to the ϵ-calculus, see e.g. [32].

[23] Somewhat relevant at this point is Moltmann’s recent work on an objectual semantics for modal claims. While the exact details of her account are irrelevant to the present endeavor, the general idea is that modal claims are made true by modal objects: sentential claims are predicated of them. Beside using some ideas from truthmaker semantics, Moltmann (in *Variable Objects and Truth-Making*, ms, 2016) understands these modal objects as “variable objects”. Variable objects are an extension of Fine’s notion of variable embodiments, i.e. things that have different manifestations at different times. Moltmann’s notion applies to things that have different manifestations as different objects at different times (and in different worlds or situations). While there’s some connection, here, with Fine’s notion of arbitrary objects understood as the semantic values of variables, there is definitely a more robust connection with Horsten’s conception of arbitrary objects as variable objects (cf. [21, 20, 22]), which are (roughly) understood as individual concepts (i.e. functions from worlds to ordinary objects). Moltmann is not really happy with identifying her variable objects with individual concepts for philosophical and linguistic reasons. Thus her idea is that variable objects are indeed objects (i.e. of syntactic type *e*, i.e. individuals), but they are associated with functions from situations to objects.

4 Formal accounts: some preliminary steps

In this section I offer some technical remarks on how to go about a generic semantics for deontic modals based on the informal proposal of the last section.

In 4.1, I lay down two arbitrary-theoretic approaches to propositional obligations: the first I label the "more objects" approach. It adds arbitrary worlds to the domain or ordinary possible worlds. The second I dub the "arbitrary selection" approach: it comes with a function that selects arbitrarily among the ordinary worlds: in this second approach there aren't distinct arbitrary worlds, but an arbitrarily selected ordinary world at which deontic sentences are evaluated.

In 4.2, I move to the first-order, considering obligations that are satisfied by arbitrary individuals, either within the "more objects" approach or the "arbitrary selection" approach.

In 4.3, I consider first-order obligations that are satisfied by arbitrary individuals in arbitrary worlds.

In 4.4, I discuss these formal approaches and deal with some objections.

4.1 The propositional case

There are several options available in the literature on generics (or reasoning with arbitrary objects) which can serve as starting points to investigate the deontic case. I am going to mention two, without any pretense of exhaustivity: the first, which adds more objects to the domain of ordinary possible worlds, is inspired by Fine's work on arbitrary objects;[24] the second does not add more objects to the domain, but selects objects arbitrarily among those already available in the domain.[25] While not exhaustive, these approaches are representative of two opposing philosophical stances.

"More objects" approach We start from a classical deontic frame $F = (W, R)$, where W is a non-empty set of points and R is a binary serial relation on W. We expand it to an arbitrary deontic frame $F^* = (W, A, R, \Delta)$, where W and R are as before, A is a (non-empty) set of arbitrary points such that $W \cap A = \varnothing$, and $\delta_i \in \Delta : A \to \wp(W)$ is a family of partial functions indexed to a set I.[26]

Transparently, the underlying idea is that A is a set of arbitrary worlds, which we adjoin to ordinary worlds. Then, given a standard valuation function $V : At \to \wp(W)$ added to the frame, we have a function which connects each arbitrary world to a set of ordinary worlds as its values, encoding the idea that an arbitrary world makes true ϕ just in case ϕ is true at all its values, i.e. the ordinary worlds it is mapped to. The connecting principle is the following (hereafter reference to a model is suppressed to avoid notational clutter):

[24] Cf. [12].

[25] And it is inspired, in spirit (but not technically), by [5].

[26] It is quite plausible that δ has to respect certain structural conditions in order to avoid problems and have an interesting logic, such as those discussed by [12]. There are some philosophical concerns one might have about introducing arbitrary objects as distinct to ordinary objects — e.g. what are their identity conditions? [12] for instance argued that two (independent) arbitrary objects are identical iff they have the same range of values, a position that generated some controversy and that Fine qualified and mitigated in unpublished work (Fine, p.c.).

Table 1: $\mathcal{M}_1$

δ_i	a_1	a_2	a_3
δ_1	$\{w_1, w_2\}$	$\varnothing$	$\{w_3\}$
δ_2	$\varnothing$	$\{w_2\}$	$\{w_2, w_3\}$

(4) $w^* \vDash \phi$ iff for all $w \in \delta_i(w^*), w \vDash \phi$, for all $i \in I$.

In words, an arbitrary world w^* makes true ϕ just in case ϕ is true at all its values, i.e. the ordinary worlds $w \in W$ it is mapped to by all assignment functions $\delta \in \Delta$. Let's see an example.

Example 1 Let's take a model $\mathcal{M}_1 = (W_1, A_1, \Delta_i, V)$, such that $W = \{w_1, w_2, w_3\}, A = \{a_1, a_2, a_3\}, \Delta_1 = \{\delta_1, \delta_2\}, V(p) = \{w_2, w_3\}$, and Δ_1 is further specified according to table 1.

In this model, for instance, $a_3 \vDash p$ but $a_2 \nvDash p$ because (with a slight abuse of notation) for all $w \in \delta_1(a_3), w \vDash p$ and for all $w \in \delta_2(a_3), w \vDash p$, but for all $w \in \delta_1(a_2), w \nvDash p$.

In case ϕ is not atomic, one just employs the usual inductive clauses at (all) ordinary worlds (selected by the δ functions). There is now a discussion to be had about bivalency. Functions in Δ are not required to be complete. In Example 1, for instance, $\delta_1(a_2) = \varnothing$ and $\delta_2(a_1) = \varnothing$. Close in spirit to the supervaluationist, arbitrary worlds which do not get associated to ordinary worlds won't be verifying anything (some words on this connection with supervaluationism will come later). However, not verifying something does not imply that its negation is verified, e.g. it does not mean that $a_2 \nvDash p$ iff $a_2 \vDash \neg p$.

We are now ready to tackle the deontic clause. In the standard possible-world semantics, the clause is the following:

(5) $w \vDash \mathcal{O}\phi$ iff for all w' s.t. $wRw', w' \vDash \phi$.

In words, ϕ is obligatory at w just in case ϕ is true at all worlds that are deontically ideal from w's perspective (deontic R can be interpreted in many ways, most of which have elicited philosophical criticism. This issue is not relevant for the point at hand, so I am sidestepping it until Section 4.4.)

The proposal of the previous section was along the line that an obligation is true, at a world, just in case it is true at an arbitrary deontically ideal world.

But how to cash out the notion of arbitrary deontically ideal world in formal terms?

There are two immediate ideas which come to mind: first, one can just introduce a deontic accessibility relation R^* between arbitrary worlds. In this first case, the deontic clause would be the following:

(6) $w^* \vDash \mathcal{O}\phi$ iff for all $w^{**} s.t. w^* R^* w^{**}, w^{**} \vDash \phi$.

In words, roughly, ϕ is obligatory at an arbitrary world just in case all accessible arbitrary worlds verify ϕ, where an arbitrary world verifies ϕ according to (4), i.e. an arbitrary world makes true ϕ just in case ϕ is true at all its values, i.e. the ordinary worlds the arbitrary world in question is mapped to.

Example 2 Let's take model $\mathcal{M}_1$ and add an accessibility relation to obtain model $\mathcal{M}_2 = (W_1, A_1, R^*, \delta_i, V)$, which is like model $\mathcal{M}_1$ except for $a_2 R^* a_3$. Now for instance $\mathcal{M}_2, a_2 \vDash \mathcal{O}p$ because $a_3 \vDash p$ (see Example 1).

Second, one might introduce a deontic accessibility relation R^m between ordinary worlds and arbitrary worlds. In this second case, the deontic clause would be the following:

(7) $w \vDash \mathcal{O}\phi$ iff for all w^* s.t. $wR^m w^*, w^* \vDash \phi$.

Much more relevant is, to my mind, the thought that a lot of deontically relevant work is being done by the δ function itself. When we consider models based on arbitrary frames, it is important to consider what happens with regard to the interpretation function.

"Arbitrary selection" approach Another approach to the propositional case, although quite different from a philosophical point of view, is the one I label "arbitrary selection" approach: instead of adding new, "arbitrary" worlds to the domain of ordinary worlds, which might be suspicious on metaphysical grounds, we select an ordinary world *arbitrarily*. Philosophically, this approach can be thought to move along the lines of arbitrary reference (cf. e.g. [5]).

In order to get concrete about this, we suppose that, given a set A of objects and a collection of subsets of $\wp A$, there is a choice function χ (in the usual mathematical sense) such that the function picks one object of A for each subset. In our case the choice function will select among the sets of worlds. The attentive reader will notice that there is a clear affinity to the ϵ calculus, where given $A = \{x : \phi x\}$, $\epsilon x \phi x$ denotes χA.[27] At this stage the details are not important; we can just understand χ as a choice-function.

For the propositional case, a model is a tuple of the form $\mathcal{M} = (W, R, \chi, V)$. Now for the arbitrary selection clause for an ought operator:

(8) $w \vDash \mathcal{O}\phi$ iff $\chi(R(w)) \vDash \phi$

where $R(w)$ is shorthand for $\{w' : wRw'\}$. The choice function selects among the sets of worlds that are accessible from the point of evaluation a world, arbitrarily, thus making formally precise the philosophical content of this second proposal.

[27] For details, historic, proof-theoretic, and model-theoretic, see at least [27], who explores the connections of the ϵ calculus to the arbitrary object theory of [12]. In particular, it is shown that the extensional version of ϵ calculus is the logic of representable arbitrary objects which obey identity (in Fine's sense).

Example 3 Let the model be $\mathcal{M}_3 = (W, R, \chi, V)$ such that $W = \{w_1, w_2, w_3\}$, $w_1 R w_2$, $w_1 R w_3$, $w_3 R w_2$, $w_3 R w_1$, $\chi(R w_1) = w_3$, $\chi(R w_2) = \varnothing$, $\chi(R w_3) = w_1$, and $V(p) = \{w_2, w_3\}$. Then $w_1 \vDash \mathcal{O}p$ but $w_3 \nvDash \mathcal{O}p$.

Since the set of worlds is stipulated to be non-empty, assuming that there are worlds that are accessible from the point of evaluation, we are guaranteed to always have an arbitrarily selected ideal world, given that we are using a standard choice function.

Intermediate approach There is (at least) a natural intermediate approach between the two just mentioned: to have an arbitrary selection among the adjoint arbitrary worlds, so that the connecting principle should be along the following lines:

(9) $w^* \vDash \phi$ iff $\chi\{w : w \in \delta(w^*)\} \vDash \phi$,

i.e. an arbitrary world w^* verifies a sentence ϕ just in case an arbitrarily selected world out of the ordinary worlds that w^* takes as values verifies ϕ. Now for the deontic clause, one adapts naturally (6) and (2), depending on whether the accessibility relation is defined among arbitrary worlds or between ordinary worlds and arbitrary worlds.

Discussion I sketched three approaches (still within the purview of possible-world semantics) to make a generic semantics for propositional obligations more concrete. The "more objects" approach has slightly higher ontological costs, for it requires that there are arbitrary objects disjoint from ordinary objects. However, the "arbitrary selection" approach has a higher explanatory burden, for it needs to be able to explain what this new arbitrary selection is and whether and how it differs from ordinary or random selection, if it wants to capture and preserve the peculiarities of generic reasoning.

Besides metaphysical preferences, before these semantics are axiomatized, there are no obvious intrinsically logical reasons to prefer one approach over the others.

4.2 The first-order case I: arbitrary individuals

It is now quite natural to consider quantified statements, where quantification is understood generically, but obligations get analyzed classically.

Therefore we will keep the familiar semantic clause for obligations, but we need a mechanism to analyze quantification generically.

The "more objects" approach We start from a classical deontic frame $F = (W, I, R)$, where W is a non-empty set of points, I is a non-empty set of ordinary individuals and R is a binary serial relation on W. We expand it to an arbitrary deontic frame $F^* = (W, I, A, R, \Delta)$, where W, I and R are as before, A is a (non-empty) set of arbitrary individuals such that $I \cap A = \varnothing$, and $\delta_j : A \to \wp(I)$ is a family of partial functions such that $\delta_j \in \Delta$, $j \in J$ is an index.

The underlying idea is that A is a set of arbitrary individuals, which we adjoin to ordinary individuals. Then we have a principle which connects arbitrary to ordinary individuals, encoding the idea that a first-order sentence $\phi(c)$ is true, with c an arbitrary individual just in case $\phi(a_j)$ is true at all values c takes, i.e. the ordinary individuals a_j c is mapped to. The connecting principle is the following:

(10) $w \vDash \phi(c)$ iff for all $a_j \in \delta_j(c)$ and for all $j \in J, w \vDash \phi(a_j)$.

It will be interesting to consider the interaction between quantifiers and modal operators. Given the mostly non-technical nature of this work, I will map out the possibilities here and leave the details to future work.

In the non-modal case, arbitrary objects are mapped to ordinary objects. In the modal case, we can have one domain of ordinary objects for all worlds, or one domain for each world.

Option 1 is to have one set of arbitrary objects mapped to one domain of ordinary objects for all worlds.

Option 2 is to have one set of arbitrary objects mapped to a different domain of ordinary objects for each world (and therefore to have a δ that is relativized to each world).

But now it's clear that what might change, from world to world, are also the arbitrary objects. Thus more options need to be considered.

Option 3 is to have a different set of arbitrary objects for each world mapped to one domain of ordinary objects for all worlds.

Option 4 is to have a different set of arbitrary objects for each world mapped to a different domain of ordinary objects for each world.

Let's just consider as an illustration the universal quantifier, for instance in a sentence of the form

(11) Everyone ought to ϕ.

Let's fix a point of evaluation w. According to the above analysis, there must be an arbitrary object, c, such that for all its values a_i, $w \vDash \mathcal{O}\phi(a_i)$ is the case. Since we are interpreting obligations classically, this boils down to check whether for all worlds accessible from $w, \phi(a_i)$ is the case. Suppose that ϕ is of the form $P(x)$ (this can be extended to relational symbols in the usual way). It is now clear that it is crucial whether we have constant domains or not.

"Arbitrary selection" approach As in the propositional case, it makes sense to consider the case where we do not add new arbitrary individuals, rather, we sort of refer arbitrarily to the individuals that already exist. One way to do so is via a choice function as before. We define one on the set of individuals I such that, assuming constant domains:

(12) $w \vDash \mathcal{O}\phi(c)$ iff for all w' s.t. $wRw', w' \vDash \phi(\chi(I))$.

In words, ϕ is obligatory of an arbitrary individual at world w just in case ϕ is true of an arbitrarily selected *ordinary* individual at all worlds accessible from the world of evaluation, w.

Intermediate approach As in the propositional case, one can formulate a clause whereby one introduces an arbitrarily selected individual among the adjoined arbitrary individuals. On this approach, ϕ is obligatory of an arbitrary individual at world w just in case ϕ is true of an arbitrarily selected *arbitrary* individual at all worlds accessible from the world of evaluation, w.

4.3 The first-order case II: arbitrary obligations, arbitrary individuals

As a third step, we can now consider to have both arbitrary obligations and arbitrary individuals in our semantics.

Since the precise technical steps are somewhat involuted but modular, I am focusing on the underlying substantial ideas in this section.

We start from a classical deontic frame $F = (W, I, R)$, where W is a non-empty set of points, I is a non-empty set of ordinary individuals and R is a binary serial relation on W. We expand it to an arbitrary deontic frame $F^* = (W, W^*, I, A, R, \Delta_w, \Delta_i)$, where W, I and R are as before, W^* is a (non-empty) set of arbitrary points such that $W \cap W^* = \emptyset$, A is a (non-empty) set of arbitrary individuals such that $I \cap A = \emptyset$. The intuitions behind this arbitrary deontic frame are exactly as before.

How do arbitrary individuals and arbitrary worlds interact? There are several options.

The easiest is to have a family of partial functions $\delta_w \in \Delta_w$ such that $\delta_w : W^* \to \wp(W)$ and a family of partial functions $\delta_i \in \Delta_i$ such that $\delta_i : A \to \wp(I)$. We keep the previous intuitions with regard to the fact that we adjoin disjoint sets of arbitrary worlds and arbitrary individuals, and each arbitrary world and (respectively) each arbitrary individual may get mapped to a set of ordinary worlds and (respectively) a set of ordinary individuals.

But what about their interaction? The most straightforward idea is the following:

(13) $w^* \vDash \mathcal{O}\phi(c)$ iff for all $a_i \in \delta_i(c)$ and for all $w \in \delta_w(w^*), w \vDash \mathcal{O}\phi(a_i)$, for all $\delta_w \in \Delta_w$ and $\delta_i \in \Delta_i$.

In words, $\mathcal{O}\phi$ is true of an arbitrary individual at an arbitrary world w^* just in case $\mathcal{O}\phi(a_j)$ is true at all values c takes, i.e. the ordinary individuals a_j c is mapped to, in all ordinary worlds w to which w^* is mapped to by all δ functions.

One can combine the "more objects" approach with the "arbitrary selection" approach with regard to worlds and individuals, or adopt two intermediate approaches. Such an investigation, which is not only a matter of different philosophical motivations, but also of formal details, is left to future more technical work.

4.4 General discussion

In this section we have sketched some options to give a generic semantics to ought claims within the familiar possible world framework. It is clear that not only in order to be useful, but also in order to evaluate each semantic proposal, the logic details need to be worked out properly.

Let me make some considerations on points still untouched by the previous discussion. First, we sketched a generic approach to sentences. However, much as defeasible approaches are often concerned about inferences, we can define a notion of generic truth and consequence too, which agrees with the classical notion when there is no reference to arbitrary objects. Such a notion is not too conceptually dissimilar from the notion of supertruth in (at least some) supervaluationist accounts.[28] In the "more objects" approach for instance we can define local (or truth-to-truth) and global (case-to-case, relative to assignments) notions of validity (cf. [12, §6], [27]).

Second, even without the full-blown logical machinery developed, we can already grasp some logical features of such a generic approach by considering classical examples. Let's focus on (10) and let a be an arbitrary prime number and $\phi(x)$ stand for $E(x) \vee O(x)$ (x is even or x is odd). Now it is clear that $Ea \vee Oa$ is true, while neither disjunct is (for an analogy, consider a simple supervaluationist semantics for vague propositions A and $\neg A$. A may be true on some precisifications, and $\neg A$ may be true on others. So while neither A nor $\neg A$ are true on all precisifications, $A \vee \neg A$ is). Similarly, neither Ea nor $\neg Ea$ is true. Classical semantical rules for disjunction and negation, for one, fail.

Third, in the "more objects" approaches we required that *all* value assignments satisfied the condition. As we highlighted in the informal part of the paper, this might well be adequate for certain domains, say, certain parts of mathematics, but less adequate to other domains, like everyday language or normative reasoning. This is the case for the reasons we mentioned: a kind, for instance, may possess properties not possessed by any of its instances (or none at all!) and, conversely, all instances of a kind may possess (contingent, for instance) properties we do not want to ascribe to the kind (for a similar discussion, see [12, p. 43]).

Finally, at the beginning of Section 2 we recapped some of the traditional criticisms to the traditional formal analysis of deontic modalities in a possible-world approach, namely aimed at the accessibility relation which is supposed to capture ideality or deontic relevance. Now one might think that the generic approaches we sketched in this section do not make any progress at all in these respect, given that they employ an accessibility relation R as orthodox possible-world approaches. This is surely true, but let me note that the generic accounts were set up in the possible-world framework on purpose to start off with a familiar basis and focus on the additional generic features. Moreover, one could dispense with the accessibility relation altogether and have the δ function perform all the relevant work. Would this reintroduce the problems generated from the use of an accessibility relation? No. In fact the δ function is decidedly less metaphysical in flavor than an accessibility relation and has relevance requirements built-in. In the end we used a possible-worlds framework out of familiarity and convenience. Nothing in the arbitrary superstructure conceptually depends on the use of it.

[28] For an introduction to supervaluationist approaches, see [35].

5 Conclusion

In this preliminary position paper I suggested a generic analysis of deontic modals via arbitrary means. I argued that such an analysis captures our real-world intuitions about deontic reasoning (and reasoning which employs deontic claims) better than the most well-known semantics.

The deductive reasoning proper of generic reasoning, as compared to defeasible approaches, seems to offer several advantages in the normative domain. It is straightforward, more predictable (for future cases, thus enhancing its action-guiding role) and replicable in a way that reasoning riddled with exceptions is not and (it may be argued) it is more suitable to the neutrality and generalizability required by normative fields like ethics and the law.

A couple of immediate problems and objections to a generic analysis of deontic modals are obvious. First, it is not clear at all that all kinds of normative reasoning can be understood within such a framework. Several refinements will be needed. Second, and perhaps relatedly, it is doubtful that, even such a theory is a good theory for deontic modals, we can reduce or understand normative reasoning in terms of deontic modals only: practical reasons, rules etc. may play an essential part too. However, it seems a good starting point, especially if we adopt an instrumental viewpoint. As a further caveat, it is rather obvious that not all deontic modals are used deontically. Our claims are to be understood as restricted to such usages.

It remains to be seen whether, once the technical details are worked out, such an analysis also offers solutions to the outstanding *logical* problems traditional deontic logics have, modulo the well-known ones generated by the possible worlds framework.

Acknowledgements Many thanks to Kit Fine, Joke Meheus, and the students of my Logic seminar at Ghent University who have discussed some of these ideas with me, and to an anonymous referee for this journal.

References

[1] Anglberger, Albert, Federico L. G. Faroldi, and Johannes Korbmacher (2016), An Exact Truthmaker Semantics for Obligation and Permission, in Deontic Logic and Normative Systems, pp. 16–31.

[2] Asher, Nicholas and Daniel Bonevac (1996), Prima facie obligation, *Studia Logica*, 57, 19.

[3] Asher, Nicholas and Daniel Bonevac (1997), Common Sense Obligation, in Nute (1997), pp. 159–203.

[4] Bongiovanni, Giorgio, Gerald Postema, Antonino Rotolo, Giovanni Sartor, Chiara Valentini, and Douglas Walton (eds.) (2018), *Handbook of Legal Reasoning and Argumentation*, Springer, Dordrecht.

[5] Breckenridge, Wylie and Ofra Magidor (2012), Arbitrary Reference, *Philosophical Studies*, 158, 3, pp. 377–400.

[6] Cohen, Ariel (2004a), Existential Generics, *Linguistics and Philosophy*, 27, 2, pp. 137–168.

[7] Cohen, Ariel, Generics and Mental Representations, Linguistics and Philosophy, 27, 5, pp. 529–556.

[8] Faroldi, Federico L. G. (2019a), Deontic Modals and Hyperintensionality, *Logic Journal of the IGPL*.

[9] Faroldi, Federico L. G. (2019b), *Hyperintensionality and Normativity*, Springer.

[10] Faroldi, Federico L. G. (2019c), The Content of Generics, under review.

[11] Ferrer Beltrn, Jordi and Giovanni Battista. Ratti (eds.) (2014), *The Logic of Legal Requirements: Essays On Defeasibility*, Oxford University Press, Oxford.

[12] Fine, Kit (1985), *Reasoning with Arbitrary Objects*, Aristotelian Society.

[13] Fine, Kit (2015a), Compliance and Command, I, ms.

[14] Fine, Kit (2015b), Compliance and Command, II, ms.

[15] Gabbay, Dov, J. F. Horty, X. Parent, R. van der Meyden, and L. van der Torre (eds.) (2013), *Handbook of Deontic Logic and Normative Systems*, College Publications, London, vol. 1.

[16] Gelman, S.A. (2010), Generics as a Window onto Young Childrens Concepts, in *Kinds, Things, and Stuff: The Cognitive Side of Generics and Mass Terms*, ed. by F.J. Pelletier, Oxford University Press, New York, pp. 100–122.

[17] Goble, Lou (2013), Notes on Kratzer Semantics for Modality with Application to Simple Deontic Logic, ms.

[18] Hansson, Sven Ove (2013), Varieties of Permission, in Gabbay et al. (2013), vol. 1.

[19] Hilpinen, Risto and Paul McNamara (2013), Deontic Logic: A Historical Survey and Introduction, in Gabbay et al. (2013), vol. 1, pp. 3–136.

[20] Horsten, Leon (2018a), Generic structures, *Philosophia Mathematica* (July 2018), doi: 10.1093/philmat/nky015.

[21] Horsten, Leon (2018b), *Generic structures*, English, Cambridge University Press, United Kingdom.

[22] Horsten, Leon and Stanislav Speranski (2018), Reasoning about arbitrary natural numbers from a Carnapian perspective, *Journal of Philosophical Logic*.

[23] Horty, John F. (2012), *Reasons as Defaults*, Oxford University Press.

[24] Koons, Robert (2017), Defeasible Reasoning, in *The Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta, Winter 2017, Metaphysics Research Lab, Stanford University.

[25] Leslie, Sarah-Jane and Adam Lerner (2016), Generic Generalizations, in *The Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta, Winter 2016, Metaphysics Research Lab, Stanford University.

[26] McNamara, Paul (2014), Deontic Logic, in *The Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta, Spring 2014.

[27] Meyer Viol, W. P. M. (1995), *Instantial Logic*, PhD thesis, University of Utrecht.

[28] Navarro, Pablo E. and Jorge L. Rodríguez (2014), *Deontic Logic and Legal Systems*, Cambridge University Press, Cambridge.

[29] Nickel, Bernhard (2016), *Between Logic and the World. An Integrated Theory of Generics*, Oxford University Press, Oxford.

[30] Nute, Donald (ed.) (1997), *Defeasible Deontic Logic*, Springer, Dordrecht.

[31] Sartor, Giovanni (2018), Defeasibility in Law, in Bongiovanni et al. (2018), pp. 315–364.

[32] Slater, Barry Hartley (2019), Epsilon Calculi, in *The Internet Encyclopedia of Philosophy*, https://www.iep.utm.edu/.

[33] Straer, Christian (2014), *Adaptive Logics for Defeasible Reasoning*, Springer.

[34] van der Torre, Leon and Y.-H Tan (1997), The Many Faces of Defeasibility, in Nute (1997), pp. 79–121.

[35] Varzi, Achille C. (forthcoming), Indeterminate Identities, Supervaluationism, and Quantifiers, *Analytic Philosophy*.

[36] Von Fintel, Kai (2012), The Best We Can (Expect to) Get? Challenges to the Classic Semantics for Deontic Modals., Paper presented in a session on Deontic Modals at the Central APA, February 17, 2012.

Federico L. G. Faroldi
federico.faroldi@ugent.be

Lise Meitner Fellow,
Department of Philosophy
University of Salzburg, Austria

Research Foundation Flanders (FWO)
Centre for Logic and Philosophy of Science
Ghent University, Belgium

International Fellow,
JQY Academy
Goethe University, Frankfurt