Steps to a Naturalistic Account of Human Deontology

Antonella Carassa

Faculty of Communication Sciences, Università della Svizzera italiana, Lugano, Switzerland

Marco Colombetti[1]

Department of Electronics, Information, and Bioengineering, Politecnico di Milano, Milano, Italy

Abstract

In this paper we outline a theory of human deontology from a naturalistic perspective. In doing so we aim to explain how human beings deal with deontic relations (like obligations and rights) thanks to a specialised psychological infrastructure, which evolved to support human cooperation. This infrastructure includes a repertoire of emotions that play a crucial role in evaluating the conformity of actions relative to a deontic relation, in signalling an agent's attitude toward their own actions or those of their deontic partners, and in motivating suitable behavioural responses. Finally we discuss the special case of interpersonal deontology, analysing its properties and relating it to Gilbert's concept of joint commitment.

1. Introduction

Human beings have the capacity to create elements of social reality in ways that are considered exclusive of our species. An important component of most (possibly all) social reality is *deontology*, that is, the complex network of obligations, permissions, rights, and so forth, that people accrue due to a variety of reasons. Part of human deontology, that we label *collective*, derives from large-scale, relatively stable sources like shared moral principles, laws, religious codes, local regulations, and social etiquette. Another part, that we call *interpersonal deontology*, is created by people in their everyday small-scale interactions, examples of this being the obligations that stem from promises, agreements, and the like.

In this paper we sketch a theory of human deontology in general, and then analyse in greater detail the case of interpersonal deontology. We find interpersonal deontology particularly interesting because, in spite of being ubiquitous in human life, it remains somewhat elusive. The main problem is to explain how it is possible for people to create such things as obligations and rights simply by interacting with other people, and to understand what types of interactions may have these effects. Most theories dealing with interpersonal deontology tend to concentrate on promises, analysed from a moral perspective: promissory obligations are thus regarded as moral obligations of some sort (see Habib, 2014, for an overview of the main theories of promissory

---

[1] This research has been carried out at Università della Svizzera italiana, where Marco Colombetti is currently adjoint professor of Communication Technologies.

obligations). But if we consider the type of deontic relations that people bring about and manage in everyday interactions, this view seems to go against intuition. Suppose for example that Ann knocks at Bob's office door asking, "May I talk to you for a minute?", to which Bob answers, "Please, come in." This interaction creates a relation between Ann and Bob, to the effect that Ann is now entitled to walk into Bob's office and talk to him. Now, it seems to us that this entitlement has nothing to do with morality. Certainly, Bob may have moral reasons to accept Ann's request to talk; for example, he may be aware that Ann is facing a personal difficulty on which he would be able to provide help, and therefore feel morally obligated to talk to her; but Ann's entitlement to talk to Bob, and Bob's correlative obligation, derive not from Bob's moral obligation to help Ann, but from their conversational interaction.

If we accept the idea that a significant part of human deontology is created by people in their daily interactions, we are left with the problem of understanding how this may be possible. In this paper we want to give a contribution in this direction. We shall first consider human deontology in general, and then concentrate on interpersonal deontology in the second part of the paper. In sketching our theory we take a naturalistic standpoint, in that we aim to explain how human beings are able to deal with deontology thanks to their psychological infrastructure, which evolved to support certain forms of cooperation.

The article is structured as follows. In Section 2 we introduce our general view of human deontology as consisting of deontic relations. In Section 3 we consider the affective dimension of deontic relations and examine the main features of deontic emotions. In Section 4 we present our view of interpersonal deontology, which takes the move, but partially departs, from Gilbert's plural subject theory and in particular from her concept of joint commitment. Finally, in Section 5 we draw some conclusions and delineate possible future developments of our work.

2. Deontology: a naturalistic view

Human deontology, either interpersonal or collective, is part of social reality; as such, it rests upon representations entertained by human beings. While large-scale collective deontology depends on the representations shared by large communities and is relatively stable, interpersonal deontology is grounded in the representations of a small number of individuals, who create and modify it through their interactions. In both cases, understanding how deontology works calls for an investigation of specific aspects of the human mind. As for any complex phenomenon concerning living organisms, we believe that such an investigation should adopt a naturalistic and evolutionary perspective. This brings in a number of important issues concerning human deontology, more precisely: identifying its function, analysing its structure, understanding the psychological mechanisms involved in its management, and finally, explaining how agents accrue specific deontic positions like obligations, rights, and so on. We shall now submit our proposals concerning the first two issues, leaving the other two to the following sections.

A comprehensive natural history of human deontology has yet to be written; however, there are important contributions to a natural history of morality (Curry, 2004; Tomasello and Vaish, 2013; Tomasello, 2016), which is a substantial part of human deontology, even if it does not exhaust it. In this context, morality is regarded as having the function of supporting human cooperation. To put the matter in the right perspective, however, it is important to construe the

notion of cooperation in a suitable way. More specifically, this concept must not only apply to the kind of *hic and nunc* collaboration that takes place in joint activities, or to cases of helpful and altruistic behaviour, but must be extended to cover all situations in which different agents jointly uphold certain elements of social reality. For example, a rivalry relationship between two agents is a form of cooperation, in the broad sense of the term, because it presupposes that both agents regard each other as rivals.

It is plausible to assume that morality, and deontology in general, has evolved to solve the conflict between the autonomy of human individuals, on the one side, and their interdependence, on the other side. Indeed, the capacity to create and manage obligations, rights, and so on allows people to rely on the fact that their conspecifics will behave in certain predictable ways for suitably extended periods, in spite of their autonomy. A reasonable assumption, therefore, is that the main selective pressure for the evolution of deontology has been the very strong degree of interdependence that humans had to face in certain phases of their evolution (Tomasello et al., 2012). It is difficult to overstate the importance of deontology in human life: most of our social interactions would be impossible without the reliance on the behaviour of others that is licenced by it. Certainly, deontology does not exhaust the solution that evolution has devised to solve the conflict between autonomy and interdependence; for example, affective bonds are another part of the solution to the same problem. However, it would be hard to imagine how we could carry out our everyday life without engaging in a thick network of deontic relations with other people.

In order to cooperate (in the broad sense introduced before), people have to relate to each other; therefore, in view of the fact that the function of deontology is to support cooperation, one may expect that all basic deontic entities, like obligations and rights, have an essential relational component. In the philosophy of law certain scholars, like Wesley Hohfeld (1913), have argued that all legal concepts are relational, in the sense that every obligation is *to someone*, every right is *against someone*, and so on. This contrasts with the view, common in the field of ethics, according to which a moral obligation is the obligation to do what is morally right, or to refrain from doing what is morally wrong, without reference to any agent to which the obligation is directed. We find the idea of a non-directed deontic position, like that of an absolute obligation, very implausible. Clearly, what we have called interpersonal deontology appears to be inherently relational, because interpersonal obligations and rights are jointly created by certain specific agents in order to bind them to each other (Carassa and Colombetti, 2013, 2014). But we believe that all kinds of deontology are relational. Consider a possible candidate for a non-directed deontic position, like the moral obligation to protect the environment. This obligation is not directed to any specific person, but can still be considered to be directed to all our kind, and in particular to the future generations. What we have here is a case of an obligation that is directed to a large collection of generic others, rather than to a small number of well-identified partners, but is still directed. The main difference between an interpersonal obligation and a moral obligation is not that only the former is directed, but rather lies in what we can call the *source* of the obligation. An interpersonal obligation, for instance one deriving from an accepted promise, is directed in that it binds exactly the two agents that respectively made and accepted the promise, that is, the promisor and the promisee. A moral obligation, on the contrary, is not "made" by an agent, but rather recognised as binding on the basis of reasons that may involve values, intersubjective empathy, and other relevant elements. Independently of its source,

however, a moral obligation to perform an action is still directed, in the sense that the action is owed to someone (e.g., to a beneficiary, or to the members of one's moral community, or to a significant moral authority). These considerations lead us to assume that deontology consists of a network of *deontic relations*, which bind agents to each other.

But what does a deontic relation consist in? This question brings us to the second issue, concerning the psychological infrastructure of human deontology. In order to fulfil its function, deontology has to affect human behaviour in the right ways. In certain specific cases, behaviour may be affected without the need of representing deontic relations explicitly; for example, we usually stop our car at a red light automatically, without bringing to consciousness the fact that a red light means a prohibition to go further. However, there are many actions that we perform not automatically, but because we consciously represent them as obligatory (morally, legally, or whatever). Moreover, even if at red lights we stop automatically, we learned to do so because we have been explained that a red light signals a prohibition to go further. It is plausible to assume that deontology could not possible exist without mental representations of deontic relations, which we call *deontic representations*.

We have argued in previous articles (Carassa and Colombetti, 2012, 2014) that deontic representations form a basic dimension of human cognition, together with the epistemic and the volitional dimensions. Deontic representations are strictly related to other types of representations like beliefs, desires, and intentions, but cannot be reduced to them. To see that this is the case, let us consider a relation of obligation; for example, assume that Ann regards herself as obligated to help her friend Bob, who urgently needs economic support. Ann's deontic representation derives from the belief that Bob needs economic support, and may cause Ann's desire to help Bob and thus her intention to lend him some money; however, Ann's deontic representation does not consist of such belief, desire, and intention. In fact, consider that while a belief has mind-to-world direction of fit (Searle, 1983), the mental representation of an obligation has world-to-mind direction of fit, because its function is to bring about a state of affairs that satisfies it. On the other hand, even if the representation of an obligation has the same direction of fit of desires and intentions, it differs from these because it contains an irreducible component of *responsibility to someone* (to Bob, in the example).[2] This is what makes obligations intrinsically relational, contrary to desires and intentions.

The relational nature of deontic representations implies that if *A* considers herself to be deontically related to *B* in some way, then she necessarily regards *B* as reciprocally related to her. In the case of an obligation to do *X*, *A* will regard herself as responsible to *B* to do *X* and, in doing so, will also regard *B* as the agent to whom such responsibility is directed. Borrowing a term from Darwall (2006), we call *authority* the reciprocal of responsibility; therefore, if *A* regards herself as responsible to *B* to do *X*, then *A eo ipso* regards *B* as having authority over her that she does *X*. We advocate that the ability to regard oneself or another agent as being

---

[2] The term "responsibility" may be interpreted in different ways. A crucial distinction is between *historic responsibility*, which concerns something that happened in the past, and *prospective responsibility*, which concerns what an agent will do in the future. For example, saying that someone is responsible for a car accident amounts to attributing historic responsibility; on the contrary, saying that parents are responsible for the safeness of their children amounts to attributing prospective responsibility. In this paper we are only concerned with prospective responsibility.

responsible to someone (or, reciprocally, as having authority over someone) is the core of all deontic relations, in that obligations, rights, and the like can be reduced to relations of responsibility and the reciprocal relations of authority. For example:

- "*A* has an obligation to *B* to do *X*" means that *A* is responsible to *B* to do *X*,

- "*A* has the right against *B* that *B* does *X*" (a claim right) means that *A* has authority over *B* that *B* does *X*,

- "*A* has permission from *B* to do *X*" (a liberty right or privilege) means that *A* has authority over *B* that *B* does not prevent *A* from doing *X*,

- "*A* has the right against *B* to do *X*" (an entitlement) means that *B* is responsible to *A* that *A* can do *X*,

and so on.

The relational nature of deontic representations deserves further analysis. From what we have said so far it is clear that *A* cannot consider herself responsible to *B* to do *X* without at the same time considering *B* as having authority over her that she does *X*: this is a matter of sheer logic, given that *being responsible to* and *having authority over* are reciprocal concepts. However, it may well be that while *A* considers herself responsible to *B* to do *X*, *B* does not consider himself as having authority over *A* that she does *X*: going back to our previous example, Ann may view herself as obligated to help Bob, without Bob viewing himself as having authority over Ann that she helps him. If indeed Bob views himself as having such authority, we shall say that Ann's and Bob's deontic representations are *congruent*. Clearly, whether two agents hold congruent deontic representations is not a matter of logic: representational congruence may or may not take place, depending on the circumstances.[3]

Although there is no guarantee that the deontic representations of different agents are congruent, a lack of congruence may be detrimental to the agents' interactions: think for example of a case where *A* thinks that she has the right against *B* that *B* does something, but *B* does not consider himself to be so obligated. In fact, deontic relations involve limitations to one's behaviour that often run against self-interest, and this may motivate an agent not to align with the deontic representations of others. But as deontology emerged to support cooperation, the congruence of deontic representations is a condition for their function to be fulfilled. We can therefore expect that human beings have evolved mechanisms that work to keep deontic representations aligned.

In the case of collective deontology (i.e., deontology deriving from shared moral principles, laws, regulations, etc.), alignment is supported by common knowledge of the sources of deontic relations and by a variety of enforcing devices (including sanctions for infringements). In the case of interpersonal deontology, which is not supported by external sources, different factors must be at work. To understand these factors, we must bear in mind that the function of deontology is to support cooperation for extended stretches of time, and that interpersonal deontic relations are entirely grounded in the mental representations of the deontic partners. We can therefore expect that the psychological infrastructure on which interpersonal deontology is

---

[3] An analogous situation may arise, for example, with common beliefs. If agent *A* represents *p* as a common belief of *A* and *B*, then she necessarily believes that also *B* represents *p* as a common belief of *A* and *B*; however, this may not be the case.

grounded has evolved specific characteristics, aimed to guarantee the effectiveness of cooperation, to back its stability in time, and to allow agents to regain freedom of movement in case their deontic partners do not live up to their responsibilities. To guarantee the effectiveness of cooperation, it is crucial that agent *A* regards herself to be deontically related to agent *B* only if she believes that *B* regards himself to be reciprocally related to *A*; to achieve stability, agents must conceive of interpersonal deontic relations as persistent in time in absence of renegotiation; and to allow agents to regain freedom of movement in the event of defections, interpersonal deontic relations must be regarded as defeasible by violations. As we shall see in Section 4, these properties are crucial to compare our view of interpersonal deontology to Gilbert's concept of joint commitment.

3. The affective dynamics of deontic relations

To fulfil its function of supporting human cooperation, a deontic relation must be able to affect behaviour. But how does this really happen? A common answer would be that deontology provides reasons for action. However, a similar answer fails to deal with the actual mental processes that drive human action. What we would like to understand is how deontic representations can concretely induce, or at least motivate, an agent to behave in certain ways. A substantial amount of research in this direction has been carried out in the field of moral psychology, and such research is relevant to our goal even if we do not take morality to exhaust all human deontology. A position common among both philosophers and psychologists is that moral behaviour is firmly grounded in human affectivity, and in particular in certain types of emotions, often called *moral emotions*. Jonathan Haidt (2003) classifies such emotions in four families, characterised by different elicitors and action tendencies: the other-condemning family (contempt, anger, and disgust), the self-conscious family (shame, embarrassment, and guilt), the other-suffering family (compassion), and the other-praising family (gratitude and elevation).

But what is an emotion? Even if no generally accepted definition is available (Scherer, 2005), there is significant consensus on some basic points. First, emotions are episodes of subjective experience directed to an object (which may be a thing, an event, a state of affairs, etc.). Second, they present their subject with an affectively charged evaluation of their object. Third, they induce or motivate certain behavioural responses in their subject. Finally, the previous features can be understood in terms of the function that an emotion serves in connection with some significant concern of the subject. For example, an episode of fear is a subjective experience that presents something as dangerous and induces or motivates a fight-or-flight response, with the function of preserving the subject's integrity.

As with all complex biological phenomena, determining the function of an emotion is essential to understand how the emotion works. From the previous considerations we can expect that human beings are capable of experiencing a number of emotions that evolved to support human deontology. Such emotions (that we call *deontic emotions*, to cover also the non-moral part of deontology) contribute to the evaluation of deontically relevant states of affairs, and induce or motivate suitable reactions. Suppose for example that Bob comes late to an important appointment with Ann. Ann will typically be somewhat resented, and will react by rebuking Bob for being late, or by adopting a cold attitude to him, or by displaying her resentment in other ways.

As we have seen, a crucial aspect of emotions is that they present their object in an evaluative mode. But what are the objects of deontic emotions, and what kind of evaluations do they provide? To understand the structure of deontic emotions we must keep in mind that they evolved to support human cooperation through the management of deontic relations. We can therefore expect deontic emotions to be structurally related to deontic relations. To better see this, let us suppose that Bob is bound to Ann by an obligation. What we want to suggest is that Ann and Bob may experience certain emotions *on the background of the obligation*. The object of such emotions will be events that are relevant to the content of the obligation, that is, events that fulfil or violate the obligation (e.g., being punctual or late at the appointment, in the previous example). This far, deontic emotions look similar to emotions that do not concern relational matters: for example, Ann may be annoyed that it has just started to rain, on the background of her intention to take ride with her new bike. But deontic emotions are different from emotions of this type, because deontology is intrinsically relational. In other words, deontic emotions not only have an object, like all other emotions, but are also directed to a recipient, and this reflects the fact that deontic positions are always *to* (or *against*) someone. Therefore, while Ann *is annoyed* by the rain, period, she will *resent Bob* for his delay. For this reason we think that deontic emotions are examples of what Stephen Strawson (1962) calls participant reactive attitudes (see Hurley and Macnamara, 2010, for another interpretation of reactive attitudes as emotions).

Let us summarise what we have said so far. We argued that human deontology is made possible by a specific psychological infrastructure that humans acquired in the course of biological evolution; that such infrastructure includes a repertoire of specific emotions, which we call deontic emotions; and that such emotions are a crucial part of the mechanisms that allow deontology to support human cooperation. We also argued that deontology is intrinsically relational, which implies that deontic emotions are not only about an object, but also directed to a recipient. Now we want to say a few words on how a deontic emotion may serve its function.

Several authors have submitted analyses of different types of emotions, including those we call deontic emotions, from a functional point of view. Fessler and Haley (2003), in particular, analyse a number of emotions whose function is to support human cooperation, namely: romantic love, gratitude, anger, envy, guilt, righteousness, contempt, shame, pride, moral outrage, moral approbation, admiration, elevation, and mirth. Some of these (i.e., guilt, righteousness, shame, moral outrage, and moral approbation) are examples of deontic emotions (or, more specifically, of moral emotions); of these, some (like guilt and righteousness) are self-evaluating, and some (like moral outrage and moral approbation) are other-evaluating.

It is fairly obvious that other-evaluating emotions reflect the relational nature of deontology: for example, Ann's resentment against Bob concerning his being late is other-evaluating, and reflects the relational structure of the deontic relation brought about by an appointment (i.e., the agents' right against one another other that the other agent be punctual). This is because resentment embodies a negative evaluation of an obligor from the standpoint of an unsatisfied obligee. Now let us assume that Bob feels guilty for being late. Contrary to resentment, guilt is self-evaluating; in what ways, then, does guilt reflect the relational structure of deontology? The answer to this question is both simple and interesting. A first remark is that even if guilt is self-evaluating, the evaluation implicitly takes into account the point of view of one's deontic

partner; that is, Bob does not just feel guilty, period, but feels guilty *towards Ann*. Second, part of the specific function of Bob's emotion is to signal to Ann that he recognises his failure to live up to his obligation, which implies that he still takes his obligation to Ann seriously. Signals of this kind are crucial for the maintenance of interpersonal relationships, which is indeed vital for future cooperation (Fessler and Haley, 2003). For example, signals of guilt often induce a reciprocal attitude of forgiveness in the offended partner, thus helping to heal a relational wound. These considerations suggest that the dynamics of personal relationships heavily rely on the partners' capacity to "read" each other's emotions, which is an important component of human intersubjectivity (Morganti, Carassa and Riva, 2008, and Zlatev et al., 2008). This implies that certain intersubjective capacities of humans are a crucial component of the psychological infrastructure of deontology.

4. Interpersonal deontology

As we have also remarked, there is an important difference between collective and interpersonal deontology. To better understand this difference, let us consider an example. Ann, who lives in a country villa in Tuscany, has promised Bob, who lives a few kilometres away, to join him for a drink. As Ann drives out of her villa and enters the public road, she incurs the legal obligation to keep her car on the right side of the road. Ann has not personally contributed to the creation of this obligation: rather, she incurs the obligation, independently of her will, as a consequence of driving on a public road. This contrasts with the obligation that Ann has to Bob because of her promise to join him for a drink: in this case, Ann has intentionally contributed to the creation of the obligation. While Ann's obligation to drive on the right side of the road is an example of collective deontology (deriving from the Italian traffic regulations), her obligation to Bob is an instance of interpersonal deontology, deriving from a previous interaction of Ann and Bob. In general, we call interpersonal deontology the network of deontic relations that certain agents create for themselves;[4] common examples are the deontic relations deriving from promises, agreements, offers, proposals, invitations, and the like. Contrary to interpersonal deontology, collective deontology applies to an agent, under appropriate conditions, even if that agent did not contribute in any way to its creation.

As we have already argued, interpersonal deontology does not pertain to the sphere of morality: for example, Ann's failing to show up at her appointment would not *per se* count as an instance of immoral behaviour.[5] To our knowledge the most articulated treatment of how human interactions may give rise to non-moral deontic relations is offered by Margaret Gilbert with the notion of joint commitment, a concept that lies at the heart of her plural subject theory (Gilbert 1996, 1999, 2000, 2004). According to Gilbert, two or more agents intentionally create a plural subject of doing X by jointly committing to do X "as a body." Once a joint commitment is in place, each party of the joint commitment is obligated to all other parties to give a suitable contribution to the performance of X. Gilbert insists that the obligations of joint commitments

---

[4] By the expression "for themselves" we mean that interpersonal deontic relations bind exactly those agents who jointly create them.

[5] Of course, missing an appointment may be immoral under special circumstances, for example if this is going to cause substantial harm to someone.

are non-moral, and that they are grounded in the joint commitment itself, which in turn has been intentionally co-created by the relevant agents. Therefore, such obligations are examples of what we have called interpersonal deontology, because they do not derive from external sources, but are intentionally brought about by the parties of the joint commitment.

While plural subject theory contains ideas that are crucial to understand interpersonal deontology, we are not completely satisfied with Gilbert's characterisation of joint commitments. In Gilbert's view, the obligations of a joint commitment to do *X* derive from the fact that all parties are jointly committed to emulate, as far as possible, a single agent that does *X*. But this idea seems problematic. Suppose for example that Ann and Bob jointly commit to climb the Matterhorn together; from this, one would like to derive certain deontic relations, like the agents' obligations to help each other when they face a difficult passage. But how can one derive such obligations from the commitment to emulate, as far as possible, a single agent that climbs the Matterhorn? It seems to us that what the two climbers commit to is to contribute, each of them within the limits of their personal capacities, to the success of their common undertaking, which depends on the fact that both of them reach the top of the mountain (and safely return home). Each agent is therefore responsible to the other (and, reciprocally, has authority over the other) for contributing to the success of the joint enterprise: the obligation to help each other to overcome possible difficulties derives from this responsibility. In all this it seems that the concept of emulating a single agent who climbs the Matterhorn plays no role.

We believe that it is possible to give an alternative account of joint commitments, which does not rely on the idea that the parties are trying to act as a single person.[6] In our view, a joint commitment can be understood as a process that results into a network of interpersonal deontic relations (or, to adopt a slightly shorter term, an *interpersonal deontic network*) that binds a set of agents.[7] Before we proceed, it is important to recall that a deontic relation is a relation of responsibility to (or, reciprocally, of authority over) that is grounded in the mental representations of the deontic partners. As we shall see, the deontic relations brought about by joint commitments show certain characteristic interdependencies, which are grounded in the underlying representations.

In the sequel we shall deal with the following questions: How are joint commitments made? What are the properties of the resulting network of deontic relations? Are there alternative ways of bringing about interpersonal deontology, which do not require a joint commitment? Let us start with the first question. Gilbert's view is that a necessary and sufficient condition for a set of agents to jointly commit to do something is that all parties express their readiness to be so committed, in conditions of common knowledge. We believe that this idea is correct, provided

---

[6] It may be objected that if we drop the "as a body" condition the concept of joint commitment is no longer Gilbert's and ought to be named differently. However, we believe that keeping the same term with a partially different definition amounts to recognising that Gilbert, rather than formulating a technical notion that makes sense only within her plural subject theory, has introduced a concept of general scope, of which it is legitimate to submit alternative views.

[7] The term "joint commitment" presents a process-product ambiguity. Ambiguities of this type are common and mostly harmless, but may cause confusion in particular cases. For this reason we prefer to reserve the term "joint commitment" to a process whose product is, in our view, an interpersonal deontic network.

that one gives a suitable definition of what it means to express something. In previous works we have argued that, in this context, an act of expressing should be analysed as a communicative act in the tradition of a Gricean analysis of communication, that is, as an act performed with a reflexive communicative intention (Carassa and Colombetti, 2009, 2014). In a few words, the basic idea is that an agent, by executing a communicative act, manifests not only the intention to commit, but also the intention that such intention be recognised.[8] As several authors have remarked, the effect of a communicative act is to intentionally bring its content "out in the open" (e.g., Airenti et al., 1993; Tomasello 2008). By doing so, an agent makes what we have called a *precommitment* (Carassa and Colombetti, 2009, 2013), which means that the agent is already committed, on condition that the other agents do the same; once all agents have precommitted, a joint commitment is successfully completed, and as a consequence an interpersonal deontic network is produced. In this process, bringing one's readiness to commit out in the open is crucial, because it guarantees that no room is left for denials. That a series of communicative acts is sufficient to produce deontic relationships may look surprising, as the only result of a communicative act is a change in the representations held by the interacting agents; however, this is coherent with the fact that deontic relationships ultimately consist in mental representations. As we shall see shortly, the properties of the deontic network resulting from a joint commitment depend on certain features of such representations.

According to Gilbert, joint commitments entail certain directed obligations. But how do such obligations compare with the deontic network we have been talking about? Gilbert has repeatedly argued (e.g., in Chapter 2 of Gilbert, 2014) that the obligations of joint commitments:

  - come into force simultaneously, when the joint commitment is made (*simultaneity*);

  - cannot be unilaterally rescinded by one of the parties (*lack of unilateral rescindability*); and

  - are binding as long as every party lives up to them, and become voidable if a party violates one of them (*interdependence*).

To carry out our comparison it is important to verify whether such properties are also enjoyed by our interpersonal deontic network. In our view the first property (simultaneity) is entailed by the process of making a joint commitment that we have previously sketched (and described in greater detailed in the already cited papers). In fact, a joint commitment is brought about by a series of communicative acts, each of which contributes certain precommitments. Only when all precommitments are in place a joint commitment is made and the interpersonal deontic network comes into existence. This implies that all the deontic relations that constitute the network are brought about simultaneously.

While simultaneity is a global property of the deontic network that depends on how the process of making a joint commitment is structured, the second property (lack of unilateral rescindability) is a local property enjoyed by every interpersonal deontic relation. As we argued at the end of Section 2, interpersonal deontic relationships are not unilaterally rescindable thanks to the psychological infrastructure of human deontology, in view of their function of supporting human cooperation.

---

[8] Note that a communicative act need not be verbal; for instance, Bob may communicate to Ann that she can come into her office by smiling at her while performing suitable gestures.

We have now to see whether the deontic relations produced by a joint commitment (in our sense) are interdependent in the way clarified by Gilbert. To this purpose, it is important to remember that deontic relations are based on mental representations; therefore, if certain deontic relations are interdependent, this property must be grounded in such representations. In a previous article (Carassa and Colombetti, 2014) we argued that the deontic representations resulting from a joint commitment are interdependent thanks to their distinctive logical form. In a nutshell, every deontic relation of the network conditionally depends on the fact that all the parties of the joint commitment live up to their responsibilities. For example, if Ann and Bob agree that Ann will do *X* and Bob will do *Y*, the resulting deontic network, in case of perfect congruence, will consist in Ann and Bob individually holding the following representations:

Ann: (a1) I am responsible to Bob, to the effect that I do *X*, as long as Bob lives up to (a2);
(a2) Bob is responsible to me, to the effect that he does *Y*, as long as I live up to (a1).

Bob: (b1) I am responsible to Ann, to the effect that I do *Y*, as long as Ann lives up to (b2);
(b2) Ann is responsible to me, to the effect that she does *X*, as long as I live up to (b1).

Now suppose that, for example, Ann does not do *X*, and this is known to both agents: thanks to (a2) Ann's can no longer consider Bob responsible to her to do *Y*; analogously, thanks to (b1) Bob will no longer consider himself responsible to Ann to do *Y*.

We have therefore shown that the interpersonal deontic network that, in our view, is produced by a joint commitment has the same fundamental properties that Gilbert attributes to the obligations of joint commitments.

So far we have examined some consequences of regarding joint commitments as having the effect of creating interpersonal deontic relations. What remains to be understood is whether there are other ways to bring about interpersonal deontology, without making a joint commitment. It seems that we can rule out such a possibility. On the one hand, interpersonal deontology, by definition, has no external sources and is grounded solely in the will of the deontic partners. On the other hand, as we have remarked at the end of Section 2, agent *A* regards herself to be deontically related to agent *B* only if *A* believes that *B* regards himself to be reciprocally related to *A*. This means that interpersonal deontic relations can only be built jointly by the relevant partners. But a process that allows a number of agents to jointly create interpersonal deontic relation is, by our definition, a joint commitment.

To have a more complete view of interpersonal deontology it would be important to investigate it affective dynamics. It is plausible to assume that the emotions related to interpersonal deontology are at least partially different from those related to collective deontology. Strawson (1962) made a distinction between those reactive attitudes that are genuinely personal, like the resentment of an offended party against the offender, and those that are impersonal or vicarious, like the blame that a disinterested observer may cast onto the offender. If we interpret reactive attitudes as emotions, as we have suggested in Section 3, events that significantly affect interpersonal deontic relations can be expected to elicit specific emotions, different from those that are typical of collective deontology. This issue can only be addressed empirically, by developing experimental designs that allow one to clearly distinguish the two types of emotions. To our knowledge, no significant result in this direction has been achieved yet.

5. Conclusions

In this paper we outlined a theory of human deontology from a naturalistic perspective. We placed centre stage the natural function of deontology, that is, to support human cooperation. We argued that deontology is essentially relational, and that human beings have evolved specific mental capacities for dealing with deontic relations, in particular the capacity of representing agents, including themselves, as related to other agents in terms of responsibility or, reciprocally, authority. Moreover, we contended that the psychological infrastructure of deontology includes a repertoire of emotions that support the management of deontic relations in line with the function of promoting human cooperation. Such emotions play crucial roles in evaluating the conformity of an agent's actions relative to existing deontic relations, in signalling an agent's attitude toward their own actions or those of their deontic partners, and in motivating suitable behavioural responses. Finally, we discussed the structure of interpersonal deontology, relating it to Gilbert's concept of joint commitment, which we proposed to revise in certain ways.

If one's goal is to understand the basic mechanisms of human cooperation, deontology is only part of the story. No doubt there are important aspects of human relationships that are not deontic, or at least cannot be completely reduced to the deontic sphere: think for example of romantic love, parent-child attachment, friendship, rivalry, and trust.[9] We are confident, however, that our approach can be extended to non-deontic aspects of human relationships that share with deontology the function of supporting human cooperation. One reason is that all such relationships appear to rely on some kind of reciprocity, either symmetric (as in the case of friendship and rivalry) or complementary (as in the case of parent-child attachment or unilateral trust). Understanding the psychological infrastructure on which such relationships are grounded, including its affective side, is a fascinating challenge for future research.

References

Airenti, G., B.G. Bara and M. Colombetti 1993. "Conversation and behavior games in the pragmatics of dialogue." *Cognitive Science* 17:197-256.

Carassa, A., and M. Colombetti 2009. "Joint meaning." *Journal of Pragmatics* 41: 1837-54.

Carassa, A., and M. Colombetti 2012. "On normative cognition, and why it matters for cognitive pragmatics." *Intercultural Pragmatics*, 9: 271-80.

Carassa, A., and M. Colombetti 2013. "Creating interpersonal reality through conversational interactions." In B. Kobow, B. Schmid, and M. Schmitz, eds., *The background of social reality*, 91-104. Berlin: Springer.

Carassa, A., and M. Colombetti 2014. "Interpersonal responsibilities and communicative intentions." *Phenomenology and the Cognitive Sciences*, 13: 145-59.

Curry, O.S. 2004. *Morality as Natural History: An adaptationist account of ethics*. PhD Thesis, Government Department, London School of Economics.

---

[9] Concerning the non-deontic nature of trust, see Darwall (2017).

Darwall, S. 2006. The second-person standpoint: Morality, respect, and accountability. Cambridge, MA: Harvard University Press.

Darwall, S. 2017. "Trust as a second-personal attitude of the heart." In P. Faulkner and T. Simpson, eds., *The philosophy of trust*, 5-50. Oxford: Oxford University Press.

Fessler, D.M.T., and K.J. Haley 2003. "The strategy of affect: Emotions in human cooperation." In P. Hammerstein, ed., *Genetic and cultural evolution of cooperation*, 7-36. Cambridge, MA: MIT Press.

Gilbert, M.P. 1996. *Living together: Rationality, sociality, and obligation*. Lanham, MD: Rowman and Littlefield.

Gilbert, M.P. 1999. Obligation and joint commitment. *Utilitas*, 11 2, 143-63.

Gilbert, M.P. 2000. Sociality and responsibility: New essays in plural subject theory. Lanham, MD: Rowman and Littlefield.

Gilbert, M.P. 2004. "Scanlon on promissory obligation: The problem of promisees' rights." *The Journal of Philosophy*, 101: 83-109.

Gilbert, M.P. 2014. *Joint commitment: How we make the social world*. Oxford: Oxford University Press.

Gilbert, M.P. 2018. *Rights and demands: A foundational inquiry*. Oxford: Oxford University Press.

Habib, A. 2014. "Promises." In E. N. Zalta, ed., *The Stanford Encyclopedia of Philosophy*. URL = https://plato.stanford.edu/archives/spr2018/entries/promises/.

Haidt, J. 2003. "The moral emotions." In R. J. Davidson, K. R. Scherer, and H. H. Goldsmith, eds., *Handbook of affective sciences*, 852-70. Oxford: Oxford University Press.

Hohfeld, W.N. 1913. "Some fundamental legal conceptions as applied in judicial reasoning." *The Yale Law Journal*, 23: 16–59.

Hurley, E.A., and C. Macnamara 2010. "Beyond belief: Toward a theory of the reactive attitudes." *Philosophical Papers*, 39: 373-99.

Morganti, F., A. Carassa, and G. Riva, eds. 2008. Enacting intersubjectivity: Cognitive and social perspectives to the study of interactions. Amsterdam: IOS Press.

Scherer, K.R. 2005. "What are emotions? And how can they be measured?" *Social Science Information*, 44: 693-727.

Searle, J.R. 1983. *Intentionality: An essay in the philosophy of mind*. Cambridge: Cambridge University Press.

Strawson, P.F. 1962. "Freedom and resentment." In Gary Watson, ed., *Proceedings of the British Academy*, 48, 1-25. Oxford: Oxford University Press.

Tomasello, M. 2008. *Origins of human communication*. Cambridge, MA: MIT Press.

Tomasello, M. 2016. *A natural history of human morality*. Cambridge, MA: Harvard University Press, 2016.

Tomasello, M., A.P. Melis, C. Tennie, E. Wyman, and E. Herrmann 2012. "Two key steps in the evolution of human cooperation: The interdependence hypothesis." *Current Anthropology*, 53: 673-92.

Tomasello, M., and A. Vaish 2013. "Origins of human cooperation and morality." *Annual Review of Psychology*, 64: 231-55.

Zlatev, J., T. Racine, C. Sinha, and E. Itkonen, eds. 2008. *The shared mind: Perspectives on intersubjectivity*. Amsterdam: John Benjamins.