Episteme, 9, 2 (2012) 171–188 © Cambridge University Press
doi:10.1017/epi.2012.7

ON KNOWLEDGE AND OBLIGATION

CAN BAŞKENT, LOES OLDE LOOHUIS AND ROHIT PARIKH
CBaskent@gc.cuny.edu

ABSTRACT

This article provides a brief overview of several formal frameworks concerning the relation between knowledge (or belief) on the one hand, and obligation (or permission) on the other. We discuss the paradox of the knower, knowledge based obligation, knowingly doing, deontic dynamic epistemology, descriptive obligations, and responsibilities as dynamic epistemology.

1. INTRODUCTION

This article provides a brief overview of several issues concerning the relation between knowledge (or belief) on the one hand, and obligation (or permission) on the other, from a slightly technically oriented view point. What we are trying to achieve is to introduce the technical aspects of the field to philosophically oriented readers, and hopefully help build some communication between the two usually separate communities of philosophers and computer scientists/logicians. While aiming to introduce several of the logical advancements in the field, we try not to go into too many technical details in order to avoid overwhelming the reader with unnecessary formalisms.

The relation between knowledge and obligation is a complicated one. Questions like

> When does an agent know that she ought to do action *a*?
> If she does *a*, without knowing that she should, does she still comply with her obligation?
> When does she have enough knowledge about the situation to be allowed to make a public announcement?
> Is it possible that she cannot be blamed for not complying with an actual obligation?
> What happens to her obligations if another agent (perhaps knowingly) fails to satisfy their joint obligation?

all deal with knowledge and obligation, but they do not seem to touch upon similar issues. Many of the studies aiming at understanding or formalizing the interaction between knowledge and obligations only address a particular question or problem related to this interaction. Hence, it is difficult, if not impossible, to provide the reader with an overview of the field without seeming to 'hop' from one topic to the next. Instead, we hope that the following 'roadmap' of this paper will help the reader keep track of the different aspects or problems we will discuss.

We start by briefly considering some purely philosophical issues relating the notions of knowledge and obligation. We motivate the general discussion by mentioning a well-known case from the ethics literature: *Singer's Pond Example*. The example illustrates that obligations are influenced by the knowledge one has, and serves as a starting point

https://doi.org/10.1017/epi.2012.7 Published online by Cambridge University Press

to discuss the idea of *normativity* from an epistemic perspective. We mention several aspects of epistemic normativity, and start with the connection between beliefs and norms. The relationship between the two brings us to an important question: 'Are our beliefs voluntary?' This question is of central importance as it directly relates to the deontology of beliefs: to what degree can holding a belief be permissible or obligatory? Another issue we point to is the connection between epistemic normativity and actions.

After these initial remarks of mostly philosophical nature, we will look at the interface of knowledge and obligation from a more formal perspective. We will start by discussing a well known paradox, *The Paradox of the Knower*, to highlight why simply combining logics of knowledge with deontic logic does not yield a system that descibes the interactions between the two sub-fields well. Then, in section 5, we discuss some frameworks that formalize how knowledge of the situation, and knowledge of the outcomes of actions, influence an agent's obligation to perform an action. It is made precise how knowledge of the circumstances is required for obligations to arise. In section 6, the focus shifts from defining when an obligation arises, to the epistemic requirements for *fulfilling* one's obligations. It is argued that simply performing the right action is not enough, if an agent doesn't know what he is doing. In section 7, the dual of obligation, namely permission, is taken as a primitive operator. Here, we discuss a framework that captures which epistemic conditions make a public announcement permissible. After discussing a formal distinction between two kinds of obligations (prescriptive and descriptive) in section 8, we present a framework that formally connects obligations on the one hand, and responsibilities on the other. It is often assumed that obligation and responsibility are two sides of the same coin: if an agent is obliged to ensure that a certain outcome will hold, then he is to be held responsible if this outcome fails to be the case. However, as we will see in this section, this need not always be the case, and depends on the agent's knowledge.

We conclude with several remarks which we believe are important to develop in the future.

We summarize what we said above by giving a list of topics that we will address in this paper.

- Philosophical remarks
- The paradox of the knower
- Knowledge based obligation
- Knowingly doing
- Deontic dynamic epistemology
- Descriptive obligations
- Responsibilities as dynamic epistemology

2. PHILOSOPHICAL REMARKS

Before describing existing formalisms that connect knowledge and obligation, let us spend some time discussing the subject of deontic epistemology and epistemic normativity from a purely philosophical point of view. Even if it is not our main goal here, we will briefly mention some philosophical works that fall within the domain of our inquiry. We are aware that the field is very broad, and the topics we have selected can very well be criticized for falling short of describing the field fully. For that reason, we advise the reader not to take our work as a full survey of the field, but rather as a subjective exposition of the

https://doi.org/10.1017/epi.2012.7 Published online by Cambridge University Press

field by working logicians. We focus our attention on those philosophical issues that raise some important questions for working logicians who may try to give a formal account of the topics in question.

In this section we first point out that knowledge and obligation appear in popular ethical treatises by pointing out a well-known example from the modern ethical literature. Then, we underline some aspects of obligations that bear some logical significance, such as normativity. This carries our discussion to the topic of beliefs, which is a central issue in epistemic and doxastic logic. Within the notion of belief, we concentrate on 'doxastic voluntarism'. This debate is interesting as it carries some meta-logical value that most logicians have not yet broadly addressed.

We start by recalling a well-known ethical debate. In a famous thought-experiment, Peter Singer discusses the following scenario (Singer 1993). Assume that on your way to your department at university, there is a small pond, and one day you see that a small child has fallen in it and is drowning. Clearly, no one can deny that you have an obligation to save the child even if it gets your clothes muddy and delays you from work. The principle behind such actions, Singer maintains, is this: 'If it is in our power to prevent something very bad from happening, without thereby sacrificing anything of comparable moral significance, we ought to do it'. Even though this principle looks very plausible and reasonable, its generalizations are not always that easy to follow. For example, following Singer's example, extreme poverty is bad and it is in our power as affluent people to reduce it 'without sacrificing anything of comparable moral significance'. Yet, we do not usually seem to spend sufficient effort to reduce the world's poverty even though, to some extent, it is in our power to reduce it.

Among many other objections to this thought-experiment discussed by Singer (1997) himself, one is of epistemic nature. Namely, in the pond example, we *know* that the child is about to drown. Thus, the knowledge of this situation triggers our action by imposing the obligation. On the other hand, once, say, poverty in Africa is considered, we usually do *not know* that a specific person in a specific town is in extreme poverty even though we consider it *epistemically likely* or hepistemically possible that, say, someone in Sudan is malnourished. Therefore, these two situations impose different types of obligations. The first obligation follows from a definitive knowledge of the event whereas in the second obligation, we do notpossess (full) knowledge of the situation. It is beyond our capacity to possess full knowledge of the problem of world's poverty. Thus, it may be claimed that we do have less obligation in the absence of the relevant knowledge. Another difference betweensaving the drowing child and saving some life in a poor country in Africa is that we *know* that by jumping into the pond we can save the child, whereas in the case of donating to a charity, we *believe* that it is highly likely that our contribution will help end suffering, yet we do not *know*. Peter Unger (1996) raises the very related issue that lack of full knowledge of the *outcome* of the action may create some hesitation in the implementation of the action.

Thus, our knowledge state influences which obligations we have. Another well-studied epistemic-deontic relation is the one between beliefs and permission. In (Pollock 1987), epistemic norms are defined as those norms which describe if it is epistemically permissible to hold various beliefs. A number of philosophers discussed different properties of norms. Some defined them in terms of whether they are rule based or not (Boghossian 2008; Engel 2007; Horwich 2008). Some identified themselves as *internalist* or *externalist* theorists (BonJour and Sosa 2003; Feldman and Conee 2001; Kornblith 2001; Pollock 1987;

https://doi.org/10.1017/epi.2012.7 Published online by Cambridge University Press

Sosa 1999; Wright 2008). Some philosophers take a rather radical position by maintaining that there cannot be any philosophical grounds for our familiar epistemic norms (Horwich 2008).

But before we adopt a certain theory of what epistemic norms are, it is important to ask oneself the question: *To what extent are our beliefs voluntary?* (Alston 1988). In this respect, doxastic voluntarism is the view that our beliefs are voluntary and that we can control them (Feldman 1988).

Let us briefly discuss a well-known argument from voluntarism. The so-called 'Voluntarism Argument' of Feldman has been endorsed by several philosophers like Plantinga and Alston (Alston 1988; Feldman 1988; Plantinga 1988).

1. Doxastic voluntarism is false.
2. If doxastic voluntarism is false, then no one has epistemic obligations.
3. Therefore, no one has epistemic obligations.

At first sight, this argument seems sound and plausible: I cannot control believing that what I see in front of me is my computer screen, so I cannot be held responsible for this belief. However, several objections can be raised to this argument and some have been summarized in Feldman (1988). For example, one could simply maintain that beliefs *are* voluntary and therefore disagree with (1). A more subtle reason for rejecting (1) has been proposed by Heil (1983) who claims that even though we may not be able to control our beliefs directly, we can control them *indirectly* by performing actions that lead us to alter the way we form beliefs. We agree with Feldman that there is a distinction between epistemic obligations to believe something (right now) and obligations to take actions to obtain evidence. Apart from rejecting (1), also (2) can and has been criticized. As Feldman (1988) mentions, one could reject (2) by maintaining that obligations pertain to actions rather than belief states. Feldman himself endorses a different objection to the voluntarism argument. Namely, he rejects (2) by claiming that we can have obligations concerning involuntary beliefs in the same way as we can have a legal obligation to pay our mortgage even if our financial situation no longer allows for this. Thus, according to Feldman, the famous 'Ought' does not imply 'Can' any more, at least in the case of epistemic obligations.

In a fairly recent article, Weatherson (2007) discusses the voluntarism argument. Contrary to Feldman, Weatherson agrees with (2) but attempts to refute the argument by rejecting (1). He claims that there are two kinds of beliefs: perceptual and inferential ones. The perceptual beliefs are spontaneous and involuntary while the inferential beliefs, the ones that involve *reasoning*, are voluntary 'in that we have the capacity to check them by paying greater heed to counter-possibilities' [ibid]. Involuntary or perceptual beliefs, Weatherson argues, are best evaluated using externalist standards like reliability. Beliefs that involve reasoning on the other hand, are justified only when 'well supported by reasons'.

Parikh (2008) distinguishes lower level perceptual animal beliefs and higher level inferential beliefs and furthermore proposes a formal system to reason about such beliefs. Another significant separation of beliefs comes from Gendler. In her insightful exposition, Gendler (2008) distinguishes two kinds of beliefs: *belief* and *alief*. The author cites several psychological experiments that she uses to address this distinction: people who are hesitant to drink a glass of juice with a completely sterilized dead cockroach in it or people who hesitate to wear a laundered shirt previously worn by a person they dislike.

https://doi.org/10.1017/epi.2012.7 Published online by Cambridge University Press

In such cases, as Gendler points out, 'they *believe* that the items in question are harmless, they also *alieve* something very different' (emphasis is hers), namely, that those objects are 'filthy' and those people felt that they should 'stay away' from them. This distinction accounts for the belief–action mismatch in some situations. In other words, we believe that the aforementioned glass of juice is harmless, but we alieve that it is filthy. The action of avoiding that glass of juice is caused by the alief in this case, not by the belief. This issue indirectly brings along the concept of alief revision. If we alieve p, and then learn not-p, what happens then? Gendler addresses this question, but does not provide a full answer akin to the famous discussions on belief revision.

An important aspect of epistemic normativity that we already briefly touched upon is its connection with actions. This debate revolves around issues such as the rationality criterion for actions, the relation between different levels of beliefs and actions and, moreover, the interaction between actions and knowledge. Alston (1988), when discussing the connection between deontology and epistemic justification, indicates that justification of *actions* is that they should not violate any relevant rules, laws, regulations etc. On the other hand, an action being permitted does not necessarily imply that 'it was required or obligatory' as it only means that 'its negation was not required or obligatory' which is indeed an application of the duality of the logical definition of the deontic modality. However, notice that the justification of beliefs differs heavily from that of actions. As argued by Alston, believing a proposition at a given time means that 'the relevant rules or principles do not forbid [the agent's] believing' that proposition at that given time. It should be remarked carefully that this does not entail any deontic obligation to believe. You may justify your belief in a god, but it does not mean that everyone is obliged to believe in a god.

In their article, Hawthorne and Stanley (2008) discuss the relation between actions and epistemology. They propose a 'reason-knowledge principle' which states that 'where one's choice is p-dependent, it is appropriate to treat the proposition that p as a reason for acting [if and only if] you know that [p]'. This principle does not solely apply to actions but does apply also to beliefs in such a way that we can treat the proposition p as a reason for believing in the proposition q only if we know that p, and p implies q.

As we will see, most of the existing formal frameworks deal with obligations to perform actions, rather than obligations or permissions to believe something. Formalizing these notions would give us a better understanding of their similarities and differences.

3. FORMALIZING KNOWLEDGE AND OBLIGATION

In the rest of this paper we focus on several formal frameworks that combine some epistemic and deontic logics. But before we go into details, we must pause for a moment and ask the question: why is it important to develop logics that combine agents' knowledge (or belief) on the on hand, and agents' obligations (or permissions) on the other?

Motivations for developing these formalisms are at least twofold. First, they provide us with a way of developing a deeper understanding of social situations and interactions. Several examples of important questions, some of which we already briefly touched upon in the introduction, are the following. When does agent i know that she ought to do action a? If she does not know that she ought to do a, is she obliged to do it anyway? Are obligations and blameworthiness always two sides of the same coin? Or can it be that

https://doi.org/10.1017/epi.2012.7 Published online by Cambridge University Press

because of my current beliefs, I cannot be blamed for not complying with an actual obligation? Also, in what cases ought she to know that she ought to do *a*? That is, in what cases is she obliged to obtain enough knowledge about her obligations? If agent *i* does *a* without knowing that it is her obligation to do so, does she fulfill her obligation nevertheless? How does direct interaction influence this picture: Suppose agents *i* and *j* are obliged to both perform *a* and *j* fails to satisfy his obligation, what happens to agent *i*'s part of the obligation? In general, how can obligations and permissions be defined in situations where agents may (perhaps knowingly) fail to satisfy their obligations?

These and many other questions arise when we try to pinpoint the exact interactions between knowledge (and beliefs) on the one hand, and obligations (and permissions, blameworthiness) on the other. The different interactions between knowledge and obligations and their properties can be, and are to some extent, made precise in philosophical logic, and can thus provide us with a better understanding of the subject.

A second motivation for developing combined logical systems, is that they can help explicate important issues for both *computer science applications* and *law*. In *distributed algorithms* for example, it is useful to formally work with notions such as what an agent is allowed to know in a certain state and what action she is obliged to perform given her current knowledge. An example of a formal system in which exactly these issues are addressed is Lomuscio and Sergot (2003). In this work the authors develop so-called *Deontic Interpreted Systems*, a deontic extension of interpreted systems. Interpreted systems provide a general framework for reasoning about knowledge and communication in a multi-agent setting, and were introduced in Halpern and Moses (1990). They have been used in computer science applications to describe agents's knowledge and their communication.

Lomuscio and Sergot's goal in their paper is to provide a grounded semantics for the deontic notions of the *ideal functioning behavior* of an agent, the *knowledge that an agent is permitted to have*, and that of the knowledge an agent *has on the assumption that the components of the system are functioning correctly according to their protocols*.

The main tool used to formalize these concepts is (local and global) states of violation and compliance. The notion of a local state in standard interpreted systems is used to represent the information available to the agent. In *deontic* interpreted systems, the set of local states is divided into allowed or correct (green) and disallowed or incorrect (red) states for each agent. Global states are tuples of local states for each agent. In this framework, the formula $\mathrm{OK}_i\ \varphi$, 'agent *i* ought to know that φ', intuitively means the agent knows that φ in all the local states he is allowed to be in. More precisely, it means that an agent *i* should know φ, if in all global states in which his local state is a green state, he knows φ.

Also when defining *computer security policies*, combined notions of knowledge and obligations are important. In Glasgow et al. 1992, for example, a framework is developed to 'specify and reason about security policies and to verify that the system adheres to such policies,' which introduces epistemic, temporal and deontic modalities in an interactive fashion.

In *law*, it is also extremely useful to formally connect the notions of knowledge and obligations. If agent *i* unknowingly commits the crime *b*, that is, performs an act that is not permissible, is she still to be blamed for her act? How does this situation change if *b* is an act that is explicitly ruled out by a *law*? Isn't agent *i* supposed to know the law? In notions like *knowingly*, *purposefully* and *recklessly* performing a criminal act as well as the notion of *liability* are formalized.

https://doi.org/10.1017/epi.2012.7 Published online by Cambridge University Press

Thus, there are many interesting interactions between knowledge and obligation that are worth formalizing using deontic-epistemic logics. Both the logic of knowledge and deontic logics have been extensively studied, and solid frameworks are well established. For example, the logic called S5 is a commonly used modal logic to represent knowledge. In this system knowledge is assumed to be *reflexive* (if an agent knows something, then it must be true), satisfy *positive introspection* (if an agent knows that something is true, he also knows that he knows it), and *negative introspection* (agents know what they don't know). A standard system of deontic logic is so-called 'Standard Deontic logic' or simply SDL. It is a standard modal logic satisfying, among other contraints, *normalcy* (if both a conditional and its antecedent are is obligatory, then so is its consequent) and *necessitation* (if something is derivable, then the claim that that thing is obligatory is also derivable).

The formal interaction between the two systems, however, has not been studied extensively. The most straightforward way to combine epistemic and deontic (modal) logic would be to simply add all the axioms of knowledge or belief ($S5$, or a weaker version-without negative introspection called $S4$) to the standard deontic axioms of standard deontic logic (SDL). This yields a strong logic but as it turns out, also a logic that fails to capture many of the subtleties that relate knowledge and obligation. A striking example showing that simply stacking the two logics together results in apparent contradictions is Åqvist's paradox of the knower. We will discuss the paradox and some of its implications in the next section.

Our aim in this paper to highlight some of the main ideas and high-level connections between different approaches. Therefore, we do not go into much technical detail but rather introduce formalisms only to clarify concepts. For more details on the formal systems, e.g., their syntax, semantics and axiomatization, we refer the reader to the original works.

4. THE PARADOX OF THE KNOWER

An important paradox of deontic logic is Åqvist's paradox. In his seminal work, Åqvist (1967) discussed his famous paradox of epistemic obligation. This paradox, also called the paradox of the knower, can be summarized as follows. Consider the following three sentences (from Hawthorne and Stanley 2008):

(1) If there is a fire, Gladys (a firefighter) ought to know it. ($p \rightarrow \mathrm{OK}_p$)
(2) There is a fire. (p)
(3) There ought not to be a fire. ($O\neg p$)

If it is assumed that, on the one hand, knowledge implies truth ($Kp \rightarrow p$), and moreover, that the standard epistemic modality ought (O) is normal ($O(p \rightarrow q) \rightarrow (Op \rightarrow Oq)$) and satisfies necessitation (If $\vdash p$ then $\vdash Op$), then both Op and $O\neg p$ can be derived from (1 – 3). The straightforward derivation of the paradox is as follows.

(4) $Kp \rightarrow p$ (axiom)
(5) OKp (from 1 and 2)
(6) $OKp \rightarrow Op$ (from 4, normality and necessitation of O)

https://doi.org/10.1017/epi.2012.7 Published online by Cambridge University Press

(7) Op (from 5 and 6)
(8) $Op \wedge \wedge O\neg p$ (from 3 and 7)

Thus, there ought to be a fire and there ought not to be a fire, which, according to Åqvist, is absurd. The problem is that $OKp \rightarrow Op$ is valid. This means that, if I ought to know something, it ought to be the case. Clearly, Åqvist and many others argue, this formula cannot be accepted as a valid principle of a formal account of 'ought to know'.

The paradox of the knower has sparked a lot of research on epistemic obligations and several systems of deontic logic were developed to account for this problem (e.g. Feldman 1990; Goble 1990; Hulstijn 2008), each focussing on different aspects of the paradox and discussing different solutions.

For example, Feldman (1990) points out the temporal aspect of the paradox. If p is settled true at time t, then $O\neg p$ can no longer be true at the same time (assuming that ought implies can). Thus, Feldman argues, making the temporal aspect explicit in the framework solves the paradox. Moreover, Feldman argues that obligations are personal and that facts about one person's obligations cannot entail facts about another's. In the current example: How can it be possible that from Gladys's obligation to know about the fire, we derive Nature's obligation for there to be a fire? Feldman makes these ideas explicit by introducing obligation operators with two subscripts: one for time and one for agents. Given these insights we can no longer derive Op and the paradox seems to be circumvented.

In a different paper on the paradox of the knower, a formal notion of knowledge-whether is distinguished from knowledge-that (Hulstijn 2008). The main point of the framework is that if, in the paradox, 'know that' (K_t) is replaced with 'know whether' (K_w), the paradox does not arise. It still holds that 'if there is a fire and Gladys knows whether there is a fire then Gladys knows that there is a fire'.

$$p \wedge K_w p \rightarrow K_t p$$

and by the same reasoning as before we can derive that

$$Op \wedge OK_w p \rightarrow OK_t p$$

but because we do not have Op we cannot derive OK_{tp} and thus Op and according to the reasoning in the paper the paradox is 'stopped'. This is for the simple reason that knowing whether φ does not imply that φ is true.

In the next section we discuss several frameworks that do not try to formalize what it means to have an obligation to know, but rather, how knowledge influences which obligations arise.

5. KNOWLEDGE BASED OBLIGATION

One approach to formalizing obligations under uncertainty was developed by Horty (2001). In line with Unger's argument mentioned in the introduction, Horty's deontic theory of agency models personal obligations to do something in situations in which *outcomes* of actions are uncertain. Basically, an agent has an obligation to perform a

https://doi.org/10.1017/epi.2012.7 Published online by Cambridge University Press

particular action *a* if, no matter what the other agents choose to do, all the possible results of performing *a* will be at least as good as the results of performing any other action.

In a recent work, Horty's decision theoretic approach has been generalized by using some standard techniques from game theory (Olde Loohuis 2009). The basic idea in that paper is that obligations should not be defined irrespective of *all* the possible moves of other players, but rather, that obligations should be defined in a world in which other agents are assumed to act *responsibly* or *rationally*. Of course, assuming rationality of other players is only one way of taking other agents into account. The framework presented can be extended to include all kinds of 'types' of players, not necessarily responsible ones.

In their (2006) paper Pacuit, Parikh and Cogan (PPC henceforth) also formalize obligations to act under uncertainty. Even though their model accounts for uncertainty of outcomes of actions, their main focus is on uncertainty of the *circumstances*. In order to choose responsibly, they argue, one needs sufficient knowledge about the circumstances. A motivating example that is used is the following.

> ***Example 1*** (*PPC*). Uma is a physician whose neighbor is ill. Uma does not know and has not been informed. Uma has no obligation (as yet) to treat her neighbor.

A formal system combining epistemic and deontic logic is proposed that can be used to study situations like the example above: situations in which obligations are circumstantial and depend on an agent's knowledge. This system builds on the history based semantics of Parikh and Ramanujam (2003) which can be summarized as follows.

Moments in time form a tree-like Kripke model: the past is deterministic, the future branches. Each (infinite) branch of the tree is called a history (H) and each history is assigned a real numbered value, a utility to the world.

In this model events happen but each agent observes only some events. This is made precise by a distinction between *global histories* and *local histories*. A global history H includes all the (relevant) events that have taken place. An agent i's local history h, on the other hand, contains only those events from H that agent i has actually observed. For events that the agent does not observe, he only notices the passing of time, modeled using non-informative clock ticks.

We denote the global histories up to and including moment m by H_m. At moment m, two histories H, H' are said to be *indistinguishable* for agent i, $H_m \sim_i H'_m$, if and only if the i's corresponding local histories h_m and h'_m are the same.

In the PPC framework, actions are also events. An action a can be performed at a finite history and yields a set $a(H)$ of *global extensions* of H. Formally,

$$a(H) = \{H' | Ha \sqsubseteq H'\}$$

where $\sqsubseteq$ is the initial segment relation and Ha denotes the finite history H appended with event a. Thus, *a(H)* is the set of infinite histories that start with the finite history H, followed by action a, and then followed by some other (infinite) sequence of actions. Also, note that in PPC's framework agents move sequentially.

In order to formalize *obligations*, a notion of H-good histories *G(H)* is introduced. *G(H)* is defined as the set of extensions of the finite history H with the highest possible value. Thus, the set of H-good histories is actually the set of *H-optimal* ones. Given

https://doi.org/10.1017/epi.2012.7 Published online by Cambridge University Press

this notion of H-goodness, an action a is defined to be *good* if and only if $G(H) \sqsubseteq a(H)$, i.e. every H-good history involves performing a. Note that this does *not* imply that performing a guarantees an H-good future.

Finally, the PPC notion of obligation is as follows:

> ***Definition 2 (PPC)***. An agent i is *obliged* to perform an action a at global history H and moment m iff a is an action which i (only) can perform, and i *knows* that it is good to perform a. Formally, $(\forall H')(H_m \sim_i H'_m$ and $H' \in G(H'_m)$ implies $H' \in a(H'_m))$.

Thus, at moment m, agent i is obliged to perform a iff a is a good action and i knows this.

It is clear that this notion of knowledge based obligation eliminates Uma's obligation to treat Sam in Example 1. The reason is that Uma does not know that Sam is ill. Assuming that treating Sam when he is ill is optimal, it follows since Sam is ill, treating him is good. However, Uma does not know this and hence has no obligation (as yet) to treat him.

More formally, there is a finite history H'_m that Uma cannot distinguish from the actual history H_m in which Sam is not ill. In extensions of this history, treating Sam is not optimal, i.e. for $H' \in \mathit{treat}(H'_m)$, $H' \notin G(H'_m)$. Thus, $(\forall H')$ $(H_m \sim_i H'_m$ and $H' \in G(H'_m)$ implies $H' \notin \mathit{treat}(H'_m))$ is not satisfied and Uma does not have the obligation to treat Sam.

The PPC-framework is also applied to an actual situation called the Kitty Genovese Murder, where a New York City woman was stabbed to death in 1964. Thirty-eight people saw the murder but only after thirty-five minutes did somebody call the police. Why? The informal analysis PPC give is the following. 'The people who saw Kitty being killed did not have default knowledge that they had the obligation to help her. They all knew that the good histories were ones in which someone called the police, but not all these histories were ones where *they themselves* were the caller. Someone else could be the caller'.

Apart from the *absolute* notion of knowledge based obligation, also a weaker concept of *default obligation* is introduced in PPC. Intuitively, an agent has a default obligation to perform action a if all maximal histories that an agent considers plausible are ones in which a is performed.

In a (2007) follow up work by Pacuit, history based structures and temporal logics have been compared and connected.[1]

6. KNOWINGLY DOING

While PPC focus on the way knowledge states *imply* certain obligations and how changes in an agent's information state *lead* to changes in his obligations, Broersen focuses on the epistemic conditions in which an agent can *comply with* or *violate* an obligation. Broersen (2008) presents a logical study of the interaction between 'ought-to-do' and 'knowingly'

1 Similar to the histories of the PPC system, temporal logical formulas are interpreted on *runs* - a concept familiar from automata theory. Pacuit investigates the relation between the two and develops a connection. Based on his results, it seems possible to translate the PPC framework to standard temporal deontic logic with the obvious interpretation on runs.

https://doi.org/10.1017/epi.2012.7 Published online by Cambridge University Press

or 'consciously' doing. Broersen introduces epistemic modalities within a deontic 'Seeing To It That' or STIT logic. STIT logic, a modal logic of action reports, was originally proposed in a series of papers by Belnap, Bartha and Horty (see Belnap et al. 2001 for an overview). The central concept of the logical system is that of 'seeing to it that', or, ensuring that something holds by performing an action. Technically, the main modality of this logic, $[i\ stit : \varphi]$ denotes that *the agent i sees to it that φ is true*.

Broersen (2008) uses a STIT framework of actions with, in addition, an indistinguishability equivalence relation for each agent between history/moment indices. Actions, in Broersen's model, take effect at the next moment in time. The concept of *knowingly* doing is defined as follows. An agent knowingly does something if its action 'holds' for all the history/state indices in the epistemic equivalence set containing the actual history/state index. Within this epistemic model action model, Broersen introduces the following STIT operator $[i\ xstit]\ \varphi$ with the intuitive interpretation that agent i sees to it that in the next state φ holds. A personal epistemic ought operator $O[i\ xstit]\varphi$ is defined as follows:

$$O[i\ xstit]\varphi = \Box(\neg[i\ xstit]\varphi \rightarrow [i\ xstit]V)$$

where V is a 'violation constant denoting that a violation occurs'. Also, $\neg[i\ xstit]\varphi$ expresses that i does not see to it that φ and thus allows a choice with possible outcome $\neg\varphi$. The intuition behind the above definition is that agent i ought to see to it that φ if and only if, by not seeing to it, he allows a violation.

According to Broersen, simply seeing to it that φ is not enough to comply with the obligation to see to it that φ. Instead, Broersen proposes, an agent should perform the action *knowingly* in order to avoid violation. Formally:

$$OK[i\ xstit]\varphi = \Box(\neg K_i[i\ xstit]\varphi \rightarrow [i\ xstit]V)$$

Finally, Broersen discusses a third variant of the obligation operator, one that involves not only knowledge about the action, but also knowledge about the obligation. Awareness of an obligation is associated with awareness of the act of bringing about a violation in case the agent does not comply in the following way:

$$KOK[i\ xstit]\varphi = \Box(\neg K_i[i\ xstit]\varphi \rightarrow K_i[i\ xstit]V)$$

7. DEONTIC DYNAMIC EPISTEMOLOGY

So far, we have discussed various frameworks that capture the idea that agents's obligations depend on their knowledge of the circumstances. Therefore, based on this assumption, if an agent's knowledge changes, his obligations can change as well. This brings us to a dynamic epistemic approach to deontic logic where we discuss how knowledge updates effect permissions.

In a nutshell, dynamic epistemic logic is a field within epistemic logic which formalizes knowledge changes and updates in multi-agent settings. Within the field of dynamic epistemic logic, there are several different methods to achieve such updates. In this section, we focus on 'public announcement logic' – a logic where agents' knowledge changes by

https://doi.org/10.1017/epi.2012.7 Published online by Cambridge University Press

external truthful announcements (Plaza 1989; Gerbrandy 1999; van Ditmarsch et al. 2007). In the framework of public announcement logic (PAL, henceforth), a truthful external agent makesan announcement, and all the agents hear the announcement, believe it, and consequently update their knowledge based on this announcement. PAL framework has successfully been used to analyze games (especially the backward induction scheme) (Van Benthem and Gheerbrant 2010; Baskent forthcoming; Fagin et al. 1995).

The language of PAL introduces a special modality $[\bullet]$ (and its dual $\langle\bullet\rangle$) for public announcements. PAL updates the epistemic model by 'state elimination'. Namely, after an announcement of, say, φ the possible worlds that do not satisfy φ are eliminated from the model. Consequently, the epistemic accessibility relation and the valuation are updated as well.

A recent framework purposed by Balbiani, van Ditmarsch and Seban captures not only the obligations to act in general, but rather the more subtle notion of 'permissions to say' within the context of PAL (2009a, 2009b). They extend the basic language of PAL by adding a new permissibility modality $P(\psi, \varphi)$ which reads 'after ψ has been publicly announced, it is permitted to say φ'. We call their system of public announcements with permissions PPAL.

Card games present an intuitive case study for PPAL (Balbiani et al. 2009a, 2009b). In card games, say, poker or the French game 'la belote', players make announcements during the game play. However, they are not *allowed* to make *any* announcement. Rules of the game *allow* or *disallow* the players to make some certain announcements.

Let us now discuss the deontic permissibility modality $P(\psi, \varphi)$ which satisfies the monotonicity principle:

$$\text{If } M\vDash(\psi \leftrightarrow \psi') \wedge (\langle\psi\rangle\varphi \leftrightarrow \langle\psi'\rangle\varphi'), \text{ then } M\vDash P(\psi, \varphi) \leftrightarrow P(\psi', \varphi')$$

The monotonicity principle intuitively says that after 'two logically equivalent announcements, the same formulas are permitted to be announced, and moreover if two propositions are logically equivalent, they are permitted to be said in the same way' (Balbiani 2009b).

We refrain from going into more technical details, and refer the reader to the aforementioned work to see the formal aspects of PPAL such as its full axiomatization and completeness. However, here we mention one major similarity of PPAL to other public announcement logics. As common in the variations of PAL, the announcement operator with the permission modality in its scope is reducible to the basic language of epistemic logic by the reduction axioms using the following scheme:

$$[\psi]P(\psi', \varphi) \leftrightarrow (\psi \rightarrow P(\langle\psi\rangle\psi', \varphi))$$

The given reduction axiom simplifies the complex formula with public announcement and permission operators, and reduces it to a formula with simpler syntax. By this method, every PPAL formula can be re-written as a PAL formula which we know to be complete.

In conclusion, PPAL presents a precise, intuitive and effective framework to discuss a certain class of games where rules prohibit or allow some moves, announcements or declarations.

https://doi.org/10.1017/epi.2012.7 Published online by Cambridge University Press

8. DESCRIPTIVE OBLIGATIONS

Another approach to (dynamic) epistemic deontic logic focuses on prescriptive and descriptive obligations to analyze the 'distinction and relation between prescriptive and descriptive obligations' (Aucher et al. 2010). An example to clarify the distinction between two such obligations can be helpful.

> ***Example 3*** (Aucher *et al.* 2010)[2]. John is driving on a highway with speed limit 130 km/h. He does not know whether he is speeding, because his speed controller is defective. But it is obligatory by the law that he knows whether he is speeding (epistemic norm 1). Besides, if he drives too fast, he should slow down. He should also know that if he drives too fast he has to slow down (epistemic norm 2). On the other hand, if he does not drive too fast, he is still permitted to speed up (and thus not to slow down). In this situation, there are two kinds of normative events.
> Prescriptive event: He comes upon road works and there is a sign announcing that he should slow down. This event can be modeled by the communicative act [*slow!!*].
> Descriptive event: A police car behind him tells him to slow down. This event can be modeled by the communicative act [*Oslow!*]. As a result he learns that he is speeding.

The logical system Aucher et al. introduced is based on the system that Castañeda developed, and reflects his idea of a two-sorted system with distinguished atoms for grounds and foci [12]. For Castañeda, grounds are the 'circumstances that originate, or surround, obligations' and are propositions that can be true or false. Foci, on the other hand, are neither true or false, but are 'the core contents of commands'.

The language has two sets of atoms: one for those propositions which cannot on their own be the foci of deontic operators and one for those propositions which can. The first set of atoms is called propositions while the second is called practitions. We refer the reader to the original work for the details of the system, including a rather involved axiomatization (Aucher et al. 2010).

The given language can be extended by epistemic and deontic dynamic operators, yielding a rather complex syntax and semantics. In such extended systems, moreover, authors discuss several cases where norms can be changed, and how such a system would work in a multi-agent setting.

This work by Aucher is one of the instances where philosophical concerns are addressed within formal logic, and are presented by rather involved technical language with axiomatization. Their system is significantly more complex than what Castañeda has envisioned, making it more expressive and but at the same time more complex.

9. RESPONSIBILITIES AS DYNAMIC EPISTEMOLOGY

Lima et al. (2010) develop a logic for reasoning about responsibility that is closely related in spirit to Broersen's 'knowingly doing'. Lima et al. consider the problem of task allocation in multi-agent systems. One way to allocate tasks to agents is by assigning

2 The example comes from an earlier non-published version.

https://doi.org/10.1017/epi.2012.7 Published online by Cambridge University Press

obligations to them. From his obligations an agent can then infer which actions are forbidden, permitted or obligatory. For example, from 'it is obligatory for agent i that φ' the agent i can infer that 'it is obligatoryfor i to execute action a'. However, Lima et al. argue that just assigning obligations is not enough to adequately guide an agent's decisions. They give the following illustrative example:

> ***Example 4*** (Lima *et al.* 2010). Consider a certain company with two bank accounts, 1 and 2. Bob, one of its employees, will pay a bill for the company by using bank account 1. Alice, another employee from the company, knows that Bob will withdraw money from one of the bank accounts to pay a bill, but she does not know which bank account Bob will use for that. It is Alice's task to keep the balances of these two bank accounts non-negative. [...] Now suppose that both accounts have not enough money to cover the payment to be performed by Bob, but that the total amount of money in these two bank accounts together is enough to cover the payment. Then, Alice can fulfil her obligation by transferring some money from account 2 to account 1. However, because her knowledge about the situation is incomplete, she also considers it possible that she should do the opposite.[3]

In this example Alice does not have enough information to decide what to do to fulfill her obligations. And thus, Lima et al. argue, she cannot be held *responsible* for keeping the balances nonnegative. However, if Alice has the possibility of asking Bob which bank account he is going to use she becomes responsible again. Note that contrary to Pacuit *et al.* (2006) she still has the *obligation* to keep the balances positive but cannot act in a responsible way. Thus, in line with Lima's argument, obligation and responsibility are not necessarily two sides of the same coin: it may be that an agent has the obligation to perform an action, or to achieve a state, but because of lack of knowledge she cannot be held responsible for failing to satisfy her obligation. Lima et al. make this distinction explicit by defining two distinct types of *responsibilities*.

In their paper, Lima et al. propose an extension of PDL that they call Coalitional Epistemic Dynamic Logic. They introduce group actions $\delta|_G$ meaning roughly that group G jointly performs action δ. Given $\delta|_G$, $[\delta|_G : \textbf{ensures}]\ \varphi$ means that if a group G executes δ, they ensure that φ will be the case, and $\langle G\rangle\varphi$ that group G is able to ensure φ by some action. Given these operators the authors are able to define obligations and furthermore responsibility: it is obligatory for an agent i that φ holds if the occurrence of $\neg\varphi$ may result in a violation. Technically, $O_i\,\varphi$ if the empty group of agents ($\langle\emptyset\rangle$) has an action that can ensure that a violation occurs if $\neg\varphi$.

$$O_i\varphi := \langle\emptyset\rangle(\neg\varphi \rightarrow vio_i)$$

Here, vio_i is a special atom with the meaning that the agent i is in violation.

Using these concepts, the authors define two types of responsibilities: *forward-looking* and *backward-looking responsibilities*. A forward-looking responsibility, also called task responsibility, means the following. An agent i is forward-looking responsible for an

3 Some linguistic errors in the original have been corrected to improve readability.

https://doi.org/10.1017/epi.2012.7 Published online by Cambridge University Press

outcome φ if she is held responsible for accomplishing the task of achieving φ. In their words: 'φ must be achieved by i, and agent i will be held accountable if it does not happen.' (Lima et al. 2010). Formally, they propose the following inductive definition of forward-looking responsibility: $R_i^1\ \varphi := O_i\ \varphi$, and

$$R_i^n := O_i(\bigvee_{\delta \in \Delta} K_i[\delta|_i : \textbf{ensures}]\varphi)^{n-1} \wedge \langle\emptyset\rangle R_i^{n-1}\varphi,$$

which means that agent i is forward-looking responsible for φ ($R_i^{\omega}\varphi$) if and only if it is obligatory for i that φ holds ($O_i\varphi$) for every reachable state. Moreover, the agent also needs to know how to ensure $\varphi(\vee_{\delta\in\Delta} K_i\ [\delta|_i : \textbf{ensures}]\varphi)$, and cannot lose this knowledge in the future. Thus, it is obligatory for i to keep the knowledge of how to ensure φ.

An agent i is backward-looking responsible for the (bad) outcome φ if he is to be blamed for this outcome. This, according to Lima et al., can only be the case if the agent has had enough knowledge about the situation at hand. Going back to example 4, Alice does not have enough knowledge to be backward-looking responsible for not transferring money from account 2 to 1. We refer to the original paper for a formal treatment of this concept.

The work of Lima et al. is significant in the sense that they formalize task allocation issues by using an extended version of propositional dynamic logic and then use the very system to express responsibilities in a natural way. Their complex operators enable us to represent various deontic situations in this system.

Before concluding, we would like make a remark on the use of 'violations' as atomic propositions that is common in the literature (see Lima et al. 2010: example 7). Whether we are talking about history based models, STIT frames, plain Kripke models or interpreted systems, violations are very useful for defining obligation and responsibilities. However, at the same time they 'hide' many of the complex relations between knowledge and obligations. Violations depend not only on the propositional truth values of a state but depend on the structure of the mode, for example uncertainty relations, as well. When does a violation hold at a moment in history, or at a state in a Kripke model? If a violation is a consequence of an action, how can it be a property of a state as well?

10. CONCLUSION

In this paper, we have provided a brief overview of several of the logical systems formalizing different aspects of the relation between knowledge and obligation. As we have seen, most of these frameworks extend different existing frameworks and address different issues, which makes comparing them hard. However, we hope to have provided the reader with a clear picture of the (types of) issues that are addressed and the (types of) solutions that are proposed, without going into too much technical detail.

We have discussed several of these issues captured by different frameworks. For example, the paradox of the knower has taught us that one cannot simply put together a logic of knowledge and a logic of obligation to obtain a epistemic-deontic logic, but that one has to be aware of many of the subtle interactions between the concepts. Also, we have seen how knowledge (of the circumstances, and of the outcomes of our actions) influences our obligations, how knowledge can influence whether or not an agent satisfies our obligations,

https://doi.org/10.1017/epi.2012.7 Published online by Cambridge University Press

and how obligation and blameworthiness are not necessarily two sides of the same coin, instead forward-looking and backward-looking responsibilities can be distinguished.

However, most of the works in the current literature are only first attempts to formalize the respective notions and many other interesting issues still remain unexplored.

Apart from the work that needs to be done to solidify and expand on existing frameworks discussed in this paper – many frameworks do not have complete axiom systems or complexity results, both of which are important for applications in computer science and law, there are many other issues yet untouched that we believe would be fruitful to develop in the future. Here we mention only a few of them.

First, it would be valuable to model *interaction between players* more explicitly using various techniques from *game theory*. For example, how do the available actions of other players influence the obligations of an agent? What if an agent does not know what other available actions other agents have? How do *costs* of actions enter the picture? What happens if other agents fail to behave according to the norm, when and how are obligations of others influenced? What if obligations can only be satisfied by cooperation? These and many other questions are very relevant and are worth exploring.

Second, we think it would be fruitful to develop a *paraconsistent* deontic-epistemic logic to express conflicting obligations and/or conflicting epistemologies which may lead to conflicting obligations. Apart from being philosophically intriguing, this would be useful for applications in, for example, distributed algorithms and database theory.

Third, as mentioned in the previous section, it would be interesting to have a formal account of *violations* of a system or an agent. When do they arise? How do they depend on the local history, actions, and knowledge state of an agent? How can different types of violations be formalized? Some offences are felonies, some are misdemeanors, and others are not crimes but merely punished by fines.

Finally, to the best of our knowledge, most formal systems deal with obligations to perform *actions* under epistemic circumstances, rather than permissions to have *beliefs*. As we have mentioned in the introduction, a rich philosophical literature exists discussing notions like doxastic voluntarism. These and other notions could be formalized using logical frameworks to provide a clearer picture of the relation between obligation and permission to believe.

ACKNOWLEDGEMENTS

We would like to thank Eric Pacuit for his comments on earlier versions of this paper, and his insightful suggestions.

REFERENCES

Alston, William P. 1988. 'The Deontological Conception of Epistemic Justification.' *Philosophical Perspectives*, 2: 257–99.

Åqvist, Lennart. 1967. 'Good Samaritans, Contrary-to-Duty Imperatives, and Epistemic Obligations.' *Noûs*, 1(4): 361–79.

Aucher, Guillaume, Boella, Guido, and van der Torre, Leendert. 2010. 'Prescriptive and Descriptive Obligations in Dynamic Epistemic Deontic Logic.' In *AI Approach to the COmplexity of Legal Systems (AICOL 2010)*, LNAI vol. 6237, pp. 150–61. Berlin: Springer.

https://doi.org/10.1017/epi.2012.7 Published online by Cambridge University Press

Balbiani, Philippe, van Ditmarsch, Hans, and Seban, Pablo. 2011. 'Reasoning about Permitted Announcements.' *Journal of Philosophical Logic*, 40(4): 445–72.
2009. 'Raisonnements sur la permission de dire.' In *Cinquièmes journées francophones modèles formels de l'interaction*. Rennes: L'Université de Rennes.
Başkent, Can. Forthcoming. 'Public Announcement Logic in Geometric Frameworks'. *Fundamenta Infomaticae*.
Belnap, Nuel, Perloff, Michael, and Xu, Ming. 2001. *Facing the Future*. Oxford: Oxford University Press.
Boghossian, Paul A. 2008. 'Epistemic Rule.' *Journal of Philosophy*, 105(9): 472–500.
BonJour, Laurence, and Sosa, Ernest. 2003. *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*. Oxford: Blackwell.
Broersen, Jan. 2008. 'A Logical Analysis of the Interaction between "Obligation–to-Do" and "Knowingly Doing".' *Lecture Notes in Computer Science*, 5076: 140–54.
Castañeda, Hector-Neri. 1988. 'Knowledge and Epistemic Obligation.' *Philosophical Perspectives*, 2: 211–33.
Engel, Pascal. 2007. 'Epistemic Norms and Rationality.' In W. Strawinski (ed.), *Festschrift for Jacek Jadacki*, pp. 359–74. Amsterdam: Springer.
Fagin, Ronald, Halpern, Joseph Y., Moses, Yoram, and Vardi, Moshe Y. 1996. *Reasoning about Knowledge*. Cambridge, MA: MIT Press.
Feldman, Fred. 1990. 'A Simpler Solution to the Paradoxes of Deontic Logic.' *Philosophical Perspectives*, 4: 309–41.
Feldman, Richard. 1988. 'Epistemic Obligations.' *Philosophical Perspectives*, 2: 235.
Feldman, Richard, and Conee, Earl. 2001. 'Internalism Defended.' In H. Kornblith (ed.), *Epistemology: Internalism and Externalism*, pp. 231–60. Malden, MA: Blackwell.
Gendler, Tamar Szabó. 2008. 'Alief and Belief.' *Journal of Philosophy*, 105(10): 634–63.
Gerbrandy, Jelle. 1999. *Bisimulations on Planet Kripke*. PhD thesis, Institute of Logic, Language and Computation, Universiteit van Amsterdam.
Glasgow, Janice, Macewen, Glenn, and Panangaden. 1992. 'A Logic for Reasoning about Security.' *ACM Transactions on Computer Systems*, 10(3): 226–64.
Goble, L. 1990. 'A Logic of Good, Should, and Would.' *Journal of Philosophical Logic*, 105(10): 570–90.
Halpern, Joseph Y., and Moses, Yoram. 1990. 'Knowledge and Common Knowledge in a Distributed Environment.' *Journal of ACM*, 37(3): 549–87.
Hawthorne, John, and Stanley, Jason. 2008. 'Knowledge and Action.' *Journal of Philosophy*, 105 (10): 570–90.
Heil, John. 1983. 'Doxastic Agency.' *Philosophical Studies*, 43(3): 355–64.
Hipinem, Risto. 2001. 'Deontic Logic.' In Lou Gable (ed.), *The Blackwell Guide to Philosophical Logic*, vol. 4, pp. 159–82. Oxford: Blackwell.
Horty, John F. 2001. *Agency and Deontic Logic*. Oxford: Oxford University Press.
Horwich, Paul. 2008. 'Ungrounded Reason.' *Journal of Philosophy*, 105(9): 453–71.
Hulstijn, Joris. 2008. 'Need to Know: Questions and the Paradox of Epistemic Obligation.' *Lecture Notes in Computer Science*, 5076: 125–39.
Kornblith, Hilary. 2001. *Epistemology: Internalism and Externalism*. Oxford: Blackwell.
Lima, Tiago de, Royakkers, Lambér, and Dignum, Frank. 2010. 'A Logic for Reasoning about Responsibility.' *Logic Journal of IGPL*, 18(1): 99–107.
Lomuscio, Alessio, and Sergot, Marek. 2003. 'Deontic Interpreted Systems.' *Studia Logica*, 75(1): 63–92.
Loohuis, Loes Olde. 2009. 'Obligations in a Responsible World.' In X. He, John F. Horty and Eric Pacuit (eds), *Logic, Rationality and interaction (LORI)*, LNCS vol. 5834: 251–62. Berlin: Springer Verlag.
Pacuit, Eric. 2007. 'Some Comments on History Based Structures.' *Journal of Applied Logic*, 5(4): 613–24.
Pacuit, Eric, Parikh, Rohit, and Cogan, Eva. 2006. 'The Logic of Knowledge Based Obligation.' *Synthese*, 149(2): 311–41.
Parikh, Rohit. 2008. 'Sentences, Belief and Logical Omniscience, or What does Induction Tell us?' *Review of Symbolic Logic*, 1(4): 459–76.
Parikh, Rohit, and Ramanujam, Ram. 2003. 'A Knowledge Based Semantics of Messages.' *Journal of Logic, Language and Information*, 12(4): 453–76.

https://doi.org/10.1017/epi.2012.7 Published online by Cambridge University Press

Plantinga, Alvin. 1988. 'Chisholmian Internalism.' In D. F. Austin (ed.) *Philosophical Analysis*, pp. 27–51. Dordrecht: Kluwer.
Plaza, Jan A. 1989. 'Logic of Public Communication.' In M. L. Emrich, M. S. Pfeifer, M. Hadzikadic, and Z. W. Ras (eds), *4th International Symposium on Methodologies for Intelligent Systems*, pp. 201–16. Oakridge: Oakridge National Laboratory, ORNL/DSRD-24.
Pollock, John. 1987. 'Epistemic Norms.' *Synthese*, 71(1): 61–95.
Singer, Peter. 1993. *Practical Ethics*. Cambridge: Cambridge University Press.
1997. 'The Drowning Child and the Expanding Circle.' *New Internationalist*, 289 (April): 28–30.
Sosa, Ernest. 1999. 'Skepticism and the Internal/External Divide.' In J. Greco and E. Sosa (eds), *The Blackwell Guide to Epistemology*, pp. 145–57. Malden, MA: Blackwell.
Unger, Peter. 1996. *Living High and Letting Die*. Oxford: Oxford University Press.
van Benthem, Johan, and Gheerbrant, Amelie. 2010. 'Game Solution, Epistemic Dynamics and Fixed-Point Logics.' *Fundamenta Infomaticae*, 100(1–4): 19–41.
van Ditmarsch, Hans, van der Hoek, Wiebe, and Kooi, Barteld. 2007. *Dynamic Epistemic Logic*. Amsterdam: Springer.
Weatherson, Brian. 2007. 'Deontology and Descartes' Demon.' *Journal of Philosophy*, 105(9): 540–69.
Wright, Crispin. 2008. 'Internal-External: Doxastic Norms and the Defusing of Skeptical Paradox.' *Journal of Philosophy*, 105(9): 501–17.

CAN BAŞKENT is a PhD student in computer science at the City University in New York. His research interests include modal and paraconsistent logics, and their applications in game theory.

LOES OLDE LOOHUIS is a PhD student in computer science at the City University in New York. Her research interests include game theory and epistemic logic, and their applications in cancer biology.

ROHIT PARIKH is a Distinguished Professor of Computer Science, Mathematics and Philosophy at the Graduate Center and Brooklyn College of the City University of New York. His research interests include logic, game theory and social software – a research program he introduced that uses mathematical and logical tools to analyze social procedures.

https://doi.org/10.1017/epi.2012.7 Published online by Cambridge University Press