
Reasoning about constitutive norms in BDI agents

N. CRIADO*, *University of Bolton, Deane Rd, Bolton, Greater Manchester
BL3 5AB*

E. ARGENTE†, *Departament de Sistemes Informàtics y Computació, Universitat
Politécnica de València, Camino de Vera s/n. 46022 Valencia (Spain)*

P. NORIEGA‡, *Institut d'Investigació en Intel·ligència Artificial, Consejo
Superior de Investigaciones Científicas, Campus de la UAB, Bellaterra, Catalonia
(Spain)*

V. BOTTI§, *Departament de Sistemes Informàtics y Computació, Universitat
Politécnica de València, Camino de Vera s/n. 46022 Valencia (Spain)*

Abstract

Software agents can be members of different institutions along their life; they might even belong to different institutions simultaneously. For these reasons, agents need capabilities that allow them to determine the repercussion that their actions would have within the different institutions. This association between the physical world, in which agents' interactions and actions take place, and the institutional world is defined by means of *constitutive* norms. Currently, the problem of how agents reason about constitutive norms has been tackled from a theoretical perspective only. Thus, there is a lack of more practical proposals that allow the development of software agents capable of reasoning about constitutive norms. In this article we propose an information model, knowledge representation and an inference mechanism to enable Belief-Desire-Intention agents to reason about the consequences of their actions on the institutions and making decisions accordingly. Specifically, the information model, knowledge representation and inference mechanism proposed in this article allows agents to keep track of the institutional state given that they have a physical presence in some real-world environment. Agents have a limited and not fully believable knowledge of the physical world (i.e. they are placed in an uncertain environment). Therefore, our proposal also deals with the uncertainty of the environment.

Keywords: Constitutive norms, institutions, BDI agents, multi-agent systems.

1 Introduction

The term *norm* has been traditionally used in the multi-agent system (MAS) area to refer to regulative norms [7] that define patterns of behaviour aimed at regulating the actions of software agents and the interactions among them.¹ However, norms can be employed in MAS not only for regulation

*E-mail: n.criado@bolton.ac.uk

†E-mail: eargente@dsic.upv.es

‡E-mail: pablo@iia.csic.es

§E-mail: vbotti@dsic.upv.es

¹The *norm* concept has been ambiguously employed inside the MAS area as a synonym of law, guideline, criterion, social expectation, etc. For a review of the different definitions given to the norm concept see [14].

issues, but also for establishing social institutions that give rise to new types of facts. These facts are called *institutional facts* since they only make sense within institutions [33]. This type of norms is known as *constitutive norms* since they create the institutional reality; i.e. they regulate the creation of institutional facts. ‘*Raising your arm in an auction counts as bidding*’ is a well-known example of constitutive norm.

Traditionally, constitutive norms have been used as bricks for building the ontology of institutions [20]. These contextual ontologies define a link between abstract concepts in which regulative norms are defined to the brute facts that take place in the application domain. Thus, constraints aimed at achieving the desired behaviour (i.e. the regulative norms) are specified at a higher abstract level (i.e. in terms of institutional facts) in order to allow different situations to be controlled through a reduced set of norms [2, 35].

We claim that constitutive norms are not only simple bricks for building institutional ontologies used on the definition of regulative norms, but they also allow agents to infer the consequences of their actions over the institutional state (e.g. constitutive norms allow agents to know that they cannot raise their arm in an auction if they do not want to make a bid). As a consequence, agents need to consider constitutive norms not only for translating abstract regulative norms into specific ones, but also for selecting the most suitable actions according to their goals and the institutional repercussions [20].

For example, a virtual assistant agent that imitates the behaviour of a human seller must be endowed with capabilities that allow it to participate in different and, even unknown, e-markets such as eBay² or Amazon.³ Each e-market is an institution that has constitutive norms that define the protocols and ontology that is used within this institution. For example, an actual e-market might have its own constitutive norms that define the process by which purchase contracts are formalised, the mechanism that must be used to open a new auction, etc. As a consequence, the virtual assistant agent requires capabilities for reasoning about constitutive norms and knowing the specific mechanisms and protocols by which the trading operations are performed in each market. These capabilities are also required in scenarios controlled by constitutive norms such as: social simulation scenarios, environments in which humans and agents interact in a realistic way, scenarios in which humans delegate tasks to personal software agents, and so on. In these scenarios, agents should represent constitutive norms explicitly to be able to determine if they are relevant and take these norms into account to make decisions accordingly.

In the literature, several proposals have been made in order to define agents endowed with capabilities for reasoning about regulative norms [10, 13, 26]. However, the role of constitutive norms in agent reasoning has not been considered in depth and there is a lack of agent decision-making procedures that consider constitutive norms.

This article answers one main question: ‘Is it possible to develop agents that consider the institutional state by reasoning about constitutive norms?’ This question entails the development of agents that are capable of reasoning about constitutive norms and determining the changes that occur in the institutional state. A suitable solution to this problem must take into account the fact that agents are situated in the real world. Thus, agents have uncertain knowledge about the current state of their environment.

In this article, we propose to endow BDI agents with an information model, knowledge representation and an inference mechanism for reasoning about constitutive norms. Thus, our proposal brings agents the possibility of reasoning about the possible interpretations of their actions within

²<http://www.ebay.com>

³<http://www.amazon.com>

different institutions and making decisions accordingly. Finally, we performed some experiments to evaluate the usefulness of our approach when it is used to reason about constitutive norms in uncertain environments. Specifically, these experiments are aimed at determining to what extent our proposal allows agents situated in uncertain environments to keep track of the institutional state.

This article is structured as follows: Section 2 describes some related work; Section 3 provides the basic definitions used in this article; Section 4 describes the information model, knowledge representation and inference mechanisms for reasoning about constitutive norms; Section 5 describes the main agent types that can be defined with our proposal; Section 6 contains an example that illustrates how our proposal enables BDI agents to represent and reason about constitutive norms; Section 7 describes the different experiments that we carried out; and Section 8 concludes the article and discusses some possible directions for future work.

2 Related work

The notion of constitutive norm was defined by philosophy in [29]. In this work, Rawls introduces the notion of constitutive norms (or constitutive rules⁴ according to Rawls's terminology) as rules that define the very institutions. However, the most well-known and referred proposal on constitutive norms was made by Searle in [31–33]. In these papers Searle proposes a classification of norms into regulative and constitutive ones. Regulative norms regulate forms of behaviour that exist independently and antecedently to them. Constitutive norms create or define these forms of behaviour controlled by regulative norms. In particular, constitutive norms define the *count-as* relationship that defines how the institutional world (i.e. the institutional facts) is built in terms of actions or state of affairs occurring in the physical world (i.e. brute facts). Therefore, constitutive norms specify under which conditions a brute fact *count-as* an institutional fact.

In the Artificial Intelligence field, the modelling of the *count-as* relationship has been introduced by Jones and Sergot in [24]. In this work, they propose the formalization of the *count-as* relationship in any action logic by means of the \Rightarrow_s operator. In this sense, the expression $X \Rightarrow_s Y$ means that ‘within the context s occurrence of X *count-as* Y ’. In the existing literature, several variations of the \Rightarrow_s operator have been proposed [18, 20, 21]. For example, in [21] Grossi and Dignum have proposed an alternative definition of the *count-as* connective for dealing with non-monotonicity. Mainly, they propose to redefine $X \Rightarrow_s F$ as ‘in context s , X *count-as* Y if it is not inconsistent’. Governatori et al. propose in [18] to model both the *count-as* connective and other normative links by means of a unique non-monotonic conditional. In [22], Grossi *et al.* used modal logic to capture distinct meanings of the *count-as* relationship.

In the MAS field, constitutive norms are used as bricks for building the ontology of institutions [20]. These contextual ontologies define a link between abstract concepts in which regulative norms are defined to the brute facts that take place in the application domain. According to this view of constitutive norms as an abstraction mechanism for allowing the definition of abstract regulative norms, in [2, 35] proposals on the specification and implementation of norms inside electronic institutions are described. Thus these approaches face up with the implementation of constitutive norms from an institutional perspective. In particular, they propose that the institution should translate abstract regulative norms into specific ones making use of the ontology defined by constitutive norms. These specific regulative norms are expressed in terms of specific and precise facts that

⁴Note that constitutive norms have been also called constitutive rules. For example, Rawls and Searle refer to them by using the term constitutive rules. Within the MAS field the term constitutive norm has been widely used [7, 30, 34]. Given that our proposal belongs to the MAS field, we have adopted the term constitutive norm in this article.

are controllable by the institution infrastructure. Similarly, in [1] an implementation of constitutive norms to relate abstract organizational specifications and norms to specific situations that took place in the physical world was proposed.

A noteworthy work on constitutive norms is the proposal of Boella *et al.* in [7]. In this work, they define a formal model of *Normative MAS* (NMA) in which the coordination and cooperation is achieved by means of constitutive and regulative norms. In addition, they use the metaphor of NMA as agents, thus the NMA have mental attitudes. In this sense, constitutive norms are not modelled as operative constraints of an institution but as beliefs of the normative agent, whereas regulative norms are the goals of the NMA. In this proposal, Boella *et al.* use constitutive norms for describing the legal consequences of actions in the normative system [8]. Thus, *metanorms* that define legal procedures for the definition of the normative system (i.e. the norm change procedures) are also constitutive.⁵ The work described in [6] details how reasoning about constitutive norms can be done from an institutional perspective. In particular, [6] proposes a mechanism for analysing and characterizing the notions of redundancy and equivalence of normative systems formed by both constitutive and regulative norms.

The problem of how agents reason about constitutive norms has only been tackled from a logical and formal perspective. However, no one has yet provided means for software agents to take constitutive norms into account in their practical reasoning; i.e. for inferring the consequences of their actions on institutions and keeping track of the institutional state. As claimed by Grossi *et al.* in [20], there is a need for proposals to allow software agents to consider constitutive norms. Similarly, in [1] it is pointed out that constitutive norms may be used by software agents to determine normative consequences of actions and determine their future actions according to norms. Our thesis is that constitutive norms are not just simple bricks for building institutional ontologies used on the definition of regulative norms. As a consequence, agents need not only to consider constitutive norms for translating abstract regulative norms into specific ones, but also they must have an explicit and subjective representation of constitutive norms. Thus, they would be able to reason about the consequences that their behaviour should have on the institutional state.

3 Preliminaries

The purpose of this article is not to propose, compare or improve existing norm or agent definitions, but to make use of these definitions for proposing an information model, knowledge representation and inference mechanism to allow agents to reason about constitutive norms. The aim of this section is to provide the reader with the basic notions of constitutive norm and normative agent used in this article.

3.1 Normative definitions

We make use of a first-order predicate language \mathcal{L} whose alphabet includes: the logical connectives $\{\wedge, \vee, \neg\}$; parentheses, brackets, and other punctuation symbols; an infinite set of variables and predicate, constant and function symbols. For simplicity, variables are implicitly universally quantified.⁶ In this article variables are written as any sequence of alphanumeric characters beginning with a capital letter. Predicate, constant and function symbols will be written as any sequence

⁵Notice that our proposal does not consider this dimension of constitutive norms as metanorms.

⁶Note that the appropriate use of Skolem functions [27] allows all existential quantifiers to be removed without loss of expressivity.

of alphanumeric characters beginning with a lower case letter. Let us assume the standard definition for well-formed formulas (*wffs*). We will make use of the standard notion of substitution of variables in a *wff*; i.e. σ is a finite and possibly empty set of pairs Y/y where Y is a variable and y is a term [17]. The set of predicate symbols is formed by *action* predicates (\mathbb{X}), which describe the actions that can be performed by agents; and *state* predicates (\mathbb{P}), which describe properties of the physical and institutional world. For example, the predicate $raise(hand)$ is an action predicate that represents the action of raising the *hand*. Similarly, the predicate $handRaised(agent)$ is an example of a state predicate that represents a property of the physical world: i.e. it represents that the *agent* has raised its hand. The institutional predicates (\mathbb{I}) are subset of the state predicates ($\mathbb{I} \subseteq \mathbb{P}$) describing the state of the institutional state. Brute predicates (\mathbb{B}) are the subset of state predicates ($\mathbb{B} = \mathbb{P} \setminus \mathbb{I}$) describing the state of the physical world. For example, the predicate $handRaised(agent)$ is a brute predicate that represents a state of the world in which the *agent* has raised its hand. Similarly, $bid(agent)$ is an example of an institutional predicate that represents a state of the world in which the *agent* has bid for an object. Specifically, this bid in the institutional world has been inferred from a state of the world in which the *agent* has raised its hand. This relationship between the physical world and the institutional world is defined in terms of constitutive norms as we explain below.

3.1.1 Constitutive norm

Constitutive norms define how actions or state of affairs taking place in the physical world (i.e. brute facts) modify facts on the institutional state (i.e. institutional facts).⁷ Thus, they do not define restrictions on the behaviours, but they introduce new classifications of facts, called institutional facts [33]. Institutional facts have been traditionally used for the definition of general regulative norms [1, 2, 35]. For example *cheating* is an institutional fact that can be defined by means of the following constitutive norm: *looking at a book in an exam* or *looking at others' cards in a card game* count-as *cheating*. The notion of *cheating* can be used in order to express in a single regulative norm that *all forms of cheating are forbidden*. Thus, constitutive norms can be used for defining the ontology used by the institution in the expression of regulative norms. Besides that, constitutive norms might allow agents to know their capabilities for modifying the institutional state. Our proposal entails agents to reason about this. Next, the formal definition of constitutive norm is provided.

DEFINITION 3.1 (Constitutive Norm)

A constitutive norm is a tuple $\langle BF, IF, C \rangle$ where:

- BF is an atomic formula built from the set \mathbb{B} , such that it represents the brute fact involved in the constitutive norm,
- IF is an atomic formula built from the set \mathbb{I} , such that it represents the institutional fact defined by the constitutive norm,
- C is *wff* of \mathcal{L} that determines the context or type of situation in which the constitutive norm is pertinent.

⁷Constitutive norms can also define how institutional facts bring about higher-level institutional facts. In this article, we will focus on constitutive norms that define the relationship between brute facts and institutional facts. However, the information model, knowledge representation and inference mechanism proposed in this article can also be used for reasoning about constitutive norms that define relationships between institutional facts.

For example, in most countries like Spain, to drive exceeding the speed limits inside the town boundaries count-as a driving offence. In Spain this limit is defined as 50 Km/h. This fact is represented by the following constitutive norm:

$$\langle \text{speed}(\text{Vehicle}), \text{drivingOffence}(\text{Vehicle}), \text{inTown}(\text{Town}, \text{Vehicle}) \wedge \text{speed}(\text{Vehicle}) > 50 \rangle$$

(Driving Offence Norm)

Once the context of a constitutive norm holds it becomes relevant and must be considered by agents. In our proposal, agents use relevant constitutive norms to extend their mental state. To ensure that the new formulas that are inserted into the agents' mental state are grounded, we define the notion of well-formed constitutive norm as follows:

DEFINITION 3.2 (Well-Formed Constitutive Norm)

A constitutive norm $\langle BF, IF, C \rangle$ is a well-formed norm iff $v_{BF} \subseteq (v_C \cup v_{IF})$ and $v_{IF} \subseteq (v_C \cup v_{BF})$; where v_X is the set of variables occurring in any formula X .

Specifically, in our proposal agents use constitutive norms to infer beliefs about the institutional state. For this reason, when an agent believes that a constitutive norm $\langle BF, IF, C \rangle$ is relevant (i.e. when there is a substitution σ such that $\sigma(C)$ holds in the agent's beliefs), then a belief about the physical world (i.e. a belief about $\sigma(BF)$) must be used to infer a belief about the institutional world (i.e. a belief about $\sigma(IF)$). For this reason, we define that $v_{IF} \subseteq (v_C \cup v_{BF})$. Moreover, we propose that agents use constitutive norms and abstract desires (i.e. desires about the institutional world) to infer concrete desires (i.e. desires about the physical world). For this reason, we define that $v_{BF} \subseteq (v_C \cup v_{IF})$.

3.2 Normative agent definition

A normative agent in this article is defined as a practical reasoning agent [9] whose actions are directed towards its goals. Specifically, this article focuses on how an agent considers constitutive norms in its decisions; i.e. how an agent considers the institutional repercussions of its acts. To make such kind of reasoning an agent considers its current circumstances (i.e. the beliefs about the world in which it is placed); and its objectives or situations that the agent would like to accomplish or bring about (i.e. its desires); and the norms that are in force in its environment. Besides that, we want that our agents can perform practical reasoning in a dynamic and uncertain world. For these reasons, the Graded BDI architecture, which has been proposed and described in detail in [11], is used in this article. Specifically, this architecture has an explicit representation of graded mental attitudes, such as graded beliefs and desires, and fits perfectly the purpose of our article.

DEFINITION 3.3 (Normative BDI Agent)

A Normative BDI agent is defined as a tuple $\langle B, D, I, N \rangle$, where:

- B, D, I are the sets of graded beliefs, desires and intentions of the agent. These sets are composed of $M(\gamma, \rho)$ expressions, where: $M \in \{\text{belief}, \text{desire}, \text{intention}\}$ is a graded modality used for representing graded beliefs, desires or intentions, respectively; γ is a grounded formula of \mathcal{L} ; and $\rho \in [0, 1]$ represents the degree associated with this mental proposition. ρ represents a

certainty degree in case of belief, a desirability degree in case of desires,⁸ and an intentionality degree in case of intentions.⁹

- N is a set formed by $norm(n)$ expressions, where n is a constitutive norm.

Thus, the sets B, D and I contain the cognitive elements, whereas the set N contains the normative elements. We have decided to represent norms separately from beliefs, desires and intentions due to two main reasons. Firstly, we consider that representing norms independently of other mental attitudes allows us to explain the norm reasoning process with more clarity: i.e. we are able to define explicitly the relationships among norms and beliefs and desires. Secondly, norms are different from beliefs. Norms have their own dynamics, semantics and are considered in different steps of the agent reasoning.

Definition 3.3 describes the information model of the Normative BDI agent used in this article. Figure 1 explains how this model is used in the agent's life-cycle. This workflow is a simplification of the workflow we proposed in [16]. For the purpose of this article it is only necessary to know that in Normative BDI agents the information flows from perception to action according to three main steps. Firstly, the agent *perceives* the environment and updates its beliefs and norms¹⁰ (see Figure 1(a)). Secondly, in the *deliberation* step, the desire set is revised (see Figure 1(b)). For example, new desires may be created from the agent requirements. Similarly, desires that have been achieved must be dropped. At this step the agent considers the set of norms that are relevant and considers them to extend the cognitive elements (this reasoning process is labelled as *norm reasoning* in Figure 1(b)). Finally, in the *decision making* step, desires help the agent to select the most suitable plan to be intended (see Figure 1(c)). Algorithm 1 illustrates the pseudo-code of the workflow performed by Normative BDI agents, showing how the components of the information model are updated.

In this article we only focus on how a Normative BDI agent takes constitutive norms into account. Other problems such as the generation of intentions from graded mental propositions [11], or the acquisition of norms [3] have been addressed in other proposals.

4 Reasoning about constitutive norms: norm reasoning rules

In our proposal, constitutive norms are used to extend the agent beliefs and desires.¹¹ Specifically, this extension is carried out by a set of *norm reasoning* rules. Figure 2 illustrates the norm reasoning process that we propose. Thereby, agents are able to determine the effect that their actions would have on the institutional state. Algorithm 2 illustrates the pseudo-code of the norm reasoning rules executed by Normative BDI agents.

Next, the specific norm reasoning rules for creating beliefs and desires from constitutive norms are provided. These rules are operational rules [28] that define a transition relation between configurations of Normative BDI agents. In our proposal a configuration of a Normative BDI agent is a tuple $Conf = \langle B, D, I, N \rangle$, where: B, D, I are the sets of graded beliefs, desires and intentions; and N

⁸As defined in [11], the desirability of a proposition γ represents to what extent an agent wants to achieve a situation in which γ holds.

⁹According to [11], in our proposal intentions are not considered as a basic attitude. Thus, the intentions of Normative BDI agents are generated online from the agents' beliefs and desires. The intentionality degree of a proposition γ is the consequence of finding a best feasible plan that permits a state of the world where γ holds to be achieved.

¹⁰Norm acquisition rules are part of the perception step since they are responsible for organizing, identifying and interpreting perceptions to represent and understand the norms that regulate the agent environment.

¹¹Beliefs and desires generated from constitutive norms are added to the belief and desire sets of agents. These two sets are later used by agents to generate and drop intentions online. Thus, there is not a direct link between constitutive norms and intentions. As a consequence, intentions are not considered in our proposal.

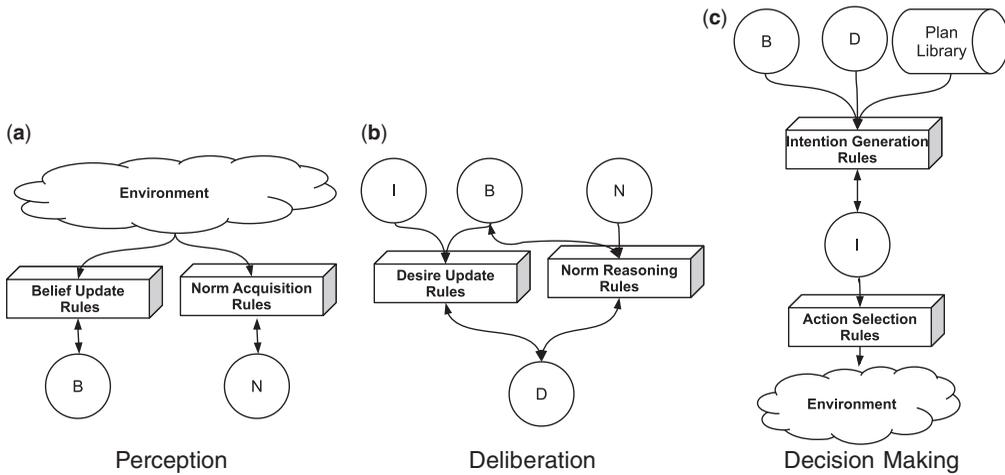


FIG. 1. Reasoning Phases in a Normative BDI Agent. The sets that contain the cognitive and normative elements (i.e. the sets B, D, I, N) are represented as circles. The reasoning processes are represented as boxes where: the input links represent the information used by the reasoning process, and the output links represent the information updated by the reasoning process.

Algorithm 1 Pseudocode of algorithm executed by Normative BDI agents

```

function REASONING( $B, D, I$ )
  while true do
     $p \leftarrow perceiveEnvironmentRules()$                                 ▷ Perception Phase
     $B \leftarrow BeliefUpdateRules(B, p)$ 
     $N \leftarrow NormAcquisitionRules(N, p)$ 
     $D \leftarrow DesireUpdateRules(D, B, I)$                                 ▷ Deliberation Phase
     $B, D \leftarrow NormReasoningRules(B, D, N)$ 
     $I \leftarrow IntentionGenerationRules(B, D, I, PlanLibrary)$           ▷ Decision Making Phase
     $a \leftarrow ActionSelectionRules(I)$ 
     $execute(a)$ 
  end while
end function

```

is the set of norms (see Definition 3.3). Thus, our norm reasoning rules define transitions between configurations as follows:

$$\frac{preCond}{Conf \rightarrow Conf'}$$

where the top of the rule—represented by the expression $preCond$ —is a boolean expression that represents the precondition of the rule, and the bottom of the rule—represented by the expression $Conf \rightarrow Conf'$ —defines the transitions between configurations: i.e. how the initial configuration—represented by the expression $Conf$ —changes once the rule is applied—represented by the expression $Conf'$.

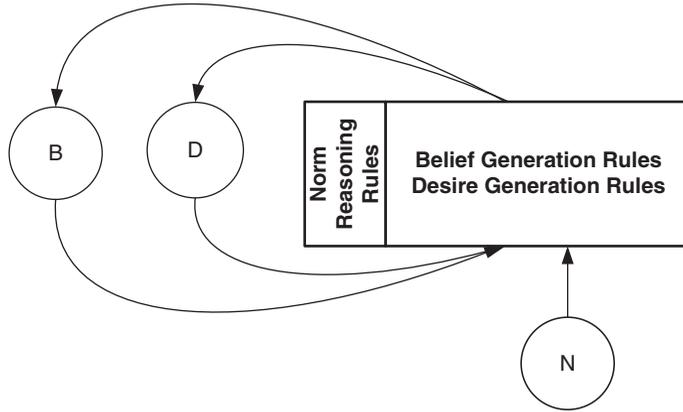


FIG. 2. Reasoning about constitutive norms in normative BDI agents. The sets that contain the cognitive and normative elements (i.e. the sets B, D, N) are represented as circles. The reasoning processes are represented as boxes where: the input links represent the information used by the reasoning process, and the output links represent the information updated by the reasoning process.

Algorithm 2 Pseudocode of NormReasoning algorithm

```

function NORMREASONINGRULES( $B, D, N$ )
   $B \leftarrow$  BeliefGenerationRules( $B, N$ )
   $D \leftarrow$  DesireGenerationRules( $D, N$ )
  return ( $B, D$ )
end function

```

4.1 Belief generation operational rules

As aforementioned, constitutive norms are used by agents for inferring beliefs about the institutional world. Next, we describe the specific rules that are used by agents to infer beliefs from constitutive norms.

Believing institutional facts being true. Informally, a constitutive norm is a rule that determines in which circumstances a brute fact *counts-as* institutional fact. Thus, when an agent believes that a constitutive norm is relevant to its current situation and the brute fact involved in the norm is true, then it should believe that the institutional fact is also true. We propose that agents must also have an explicit and subjective knowledge (i.e. beliefs) about institutional facts. To this aim, we propose in this article the following rule that models this reasoning process as follows:

$$\frac{\exists \text{norm}(\langle BF, IF, C \rangle) \in N \wedge \exists \sigma : (\text{belief}(\sigma(C), \rho_{\sigma(C)}) \in B \wedge \text{belief}(\sigma(BF), \rho_{\sigma(BF)}) \in B)}{\langle B, D, I, N \rangle \rightarrow \langle B', D, I, N \rangle}$$

where $B' = f_{\text{Realism}}(B, \text{belief}(\sigma(IF), f_{\text{Expansion}}(\rho_{\sigma(BF)}, \rho_{\sigma(C)})))$
(Believing institutional facts being true)

This rule can be applied when (i) there is a constitutive norm—i.e. exists a norm $\text{norm}(\langle BF, IF, C \rangle)$ in N —; (ii) the agent believes that this constitutive norm is relevant—i.e. there is a substitution σ such that the expression $(\text{belief}(\sigma(C), \rho_{\sigma(C)}))$ is in B ; where $\sigma(C)$ denotes the result of applying σ

to C , and $\rho_{\sigma(C)}$ is a real number within the $[0, 1]$ representing the certainty about the context of the constitutive norm (i.e., the relevance or pertinence of the norm to the current situation)—; and (iii) the basic fact involved in the constitutive norm holds—i.e. the expression $\text{belief}(\sigma(BF), \rho_{\sigma(BF)})$ is in B ; where $\rho_{\sigma(BF)}$ is a real number within the $[0, 1]$ that represents the certainty about the brute fact. When the three preconditions are true, then a new belief will be inferred corresponding to the new institutional fact—i.e. a belief as $\text{belief}(\sigma(IF), f_{\text{Expansion}}(\rho_{\sigma(BF)}, \rho_{\sigma(C)}))$ is inferred.

The certainty degree assigned to the new belief represents the certainty about the institutional fact. It is defined by the $f_{\text{Expansion}}$ function, which combines the relevance of the norm—i.e. the value of $\rho_{\sigma(C)}$ —and the certainty in which the brute fact holds—i.e. the value of $\rho_{\sigma(BF)}$ —as a real value within the $[0, 1]$ interval. Both conditions, the relevance of the norm and the brute fact, are required for creating a new belief. For example, if a norm is not relevant and it has no longer effect (i.e. $\rho_{\sigma(C)}$ is 0), then the certainty of any institutional fact inferred from this norm must be 0 regardless of the certainty of the brute fact. Moreover, the higher the relevance or the certainty of the brute facts is, the higher the certainty of the institutional fact must be. Therefore, $f_{\text{Expansion}}$ is defined as a numerical fusion operator¹² that can be given different definitions depending on the properties that are required in each concrete application. In this article, we consider that the conditions that are necessary to create a belief about an institutional fact (i.e. the existence of a relevant constitutive norm and the existence of a belief about the brute fact involved in this constitutive norm) are independent (e.g. the consideration of a constitutive norm as relevant does not imply that the agent believes that the brute fact involved in the norm holds). As a consequence, the combination among the uncertain values that cause the creation of a belief about an institutional fact is defined as follows:

$$f_{\text{Expansion}}(\rho_{\sigma(BF)}, \rho_{\sigma(C)}) = \rho_{\sigma(BF)} * \rho_{\sigma(C)}$$

f_{Realism} is a function that takes as input a set of beliefs and a new belief inferred from a constitutive norm and combines them. The way in which the belief set and the new belief are combined determines the agent personality. In the next section, we describe some of the main agent types according to the definition given to this function.

Believing institutional facts being false. When an agent believes that a constitutive norm is relevant to its current situation and the brute fact involved in the norm is false, then it should believe that the institutional fact is also false.¹³ To model this reasoning process, we propose in this article the following rule:

$$\frac{\exists \text{norm}(\langle BF, IF, C \rangle) \in N \wedge \exists \sigma : (\text{belief}(\sigma(C), \rho_{\sigma(C)}) \in B \wedge \text{belief}(\neg\sigma(BF), \rho_{\sigma(BF)}) \in B)}{\langle B, D, I, N \rangle \rightarrow \langle B', D, I, N \rangle}$$

where $B' = f_{\text{Realism}}(B, \text{belief}(\neg\sigma(IF), f_{\text{Expansion}}(\rho_{\sigma(BF)}, \rho_{\sigma(C)})))$
(Believing institutional facts being false)

If an agent believes that a constitutive norm is relevant—i.e. $\text{belief}(\sigma(C), \rho_{\sigma(C)})$ is in B —and the basic fact involved in the constitutive norm does not hold—i.e. $\text{belief}(\neg\sigma(BF), \rho_{\sigma(BF)})$ belongs to B —, then a new belief will be inferred corresponding to the negation of the institutional fact—i.e. $\text{belief}(\neg\sigma(IF), f_{\text{Expansion}}(\rho_{\sigma(BF)}, \rho_{\sigma(C)}))$ is generated.

¹²For a review and classification of data fusion operators see [5].

¹³Note that our agents do not use a close world assumption where everything unknown is false. In contrast, our agents assume that everything unknown is uncertain.

4.2 Desire generation operational rules

Constitutive norms are also used for inferring desires. Specifically, constitutive norms are used to infer desires to bring about states of affairs in which brute facts hold or not. Agents' motivations are the basis for determining which actions will be carried out. Since agents have no capabilities for altering the institutional state directly, then constitutive norms define how abstract desires (which are related to institutional facts) can be redefined in terms of brute facts that can be modified by agents.

Desiring to bring about brute facts. When an agent believes that a constitutive norm is relevant to its current situation and it desires to achieve a state of affairs in which the institutional fact defined by the norm holds, then the agent should also desire to achieve a state of affairs in which the related brute fact holds. To model this reasoning process, we propose in this article the following rule:

$$\frac{\exists \text{norm}(\langle BF, IF, C \rangle) \in N \wedge \exists \sigma : (\text{belief}(\sigma(C), \rho_{\sigma(C)}) \in B \wedge \text{desire}(\sigma(IF), \rho_{\sigma(IF)}) \in D)}{\langle B, D, I, N \rangle \rightarrow \langle B, D', I, N \rangle}$$

where $D' = f_{\text{Orientation}}(D, \text{desire}(\sigma(BF), f_{\text{Expansion}}(\rho_{\sigma(IF)}, \rho_{\sigma(C)})))$
(Desiring to bring about brute facts)

This rule can be applied when (i) there is a constitutive norm—i.e. if exists a norm $\text{norm}(\langle BF, IF, C \rangle)$ in N —; (ii) the agent believes that this constitutive norm is relevant—i.e. if there is a substitution σ such that the expression $(\text{belief}(\sigma(C), \rho_{\sigma(C)}))$ is in B —; and the agent desires to bring about a state of affairs in which the institutional fact defined by the norm holds—i.e. $\text{desire}(\sigma(IF), \rho_{\sigma(IF)})$ is in D . When the three preconditions are true, then a new desire will be inferred corresponding to the brute fact—i.e. $\text{desire}(\sigma(BF), f_{\text{Expansion}}(\rho_{\sigma(IF)}, \rho_{\sigma(C)}))$ is inferred.

As before, the desirability degree assigned to the new desire is defined by the $f_{\text{Expansion}}$ functions:

$$f_{\text{Expansion}}(\rho_{\sigma(IF)}, \rho_{\sigma(C)}) = \rho_{\sigma(IF)} * \rho_{\sigma(C)}$$

$f_{\text{Orientation}}$ is a function that takes as input a set of desires and a new desire inferred from a constitutive norm and combines them. In the next section, we describe some of the main agent types according to the definition given to this function.

Desiring not to bring about brute facts. When an agent believes that a constitutive norm is relevant to its current situation and it desires to achieve a state of affairs in which the institutional fact defined by the norm does not hold, then the agent should also desire to achieve a state of affairs in which the brute fact does not hold. To model this reasoning process, we propose the following rule in this article :

$$\frac{\exists \text{norm}(\langle BF, IF, C \rangle) \in N \wedge \exists \sigma : (\text{belief}(\sigma(C), \rho_{\sigma(C)}) \in B \wedge \text{desire}(\neg\sigma(IF), \rho_{\sigma(IF)}) \in D)}{\langle B, D, I, N \rangle \rightarrow \langle B, D', I, N \rangle}$$

where $D' = f_{\text{Orientation}}(D, \text{desire}(\neg\sigma(BF), f_{\text{Expansion}}(\rho_{\sigma(IF)}, \rho_{\sigma(C)})))$
(Desiring not to bring about brute facts)

The main difference between the implementation of constitutive and regulative norms is that regulative norms are motivational [12] (i.e. they create new desires to comply with them) whereas constitutive norms are a special kind of inference rules for extending the belief and desire theories. However, constitutive norms do not affect directly the agents' behaviour. Constitutive norms do

not create desires to bring about states of affairs different from the ones already desired. Actually, constitutive norms create new desires about the physical world that allow existing desires about the institutional world to be achieved. For this reason, our agents have no motivations for rejecting the desires inferred from constitutive norms. As a consequence, the rules described in this section do not include any condition that checks whether the agent accepts the formulas that are inferred from constitutive norms.

5 Agent types with respect to constitutive norms

Classically, agent types are characterized by stating conflict resolution types in terms of orders between mental attitudes [10]. For example, an agent is social when obligations are stronger than the other motivational components [19], etc.

According to the *norm reasoning* rules defined in Section 4, constitutive norms are taken into account by agents to extend their belief and desire set. It is possible that the formulas generated by the norm reasoning rules are inconsistent¹⁴ with the existing beliefs or desires of the agent. For example, an agent may observe some evidence that sustains a belief about an institutional fact while being able to infer a belief about the same institutional fact with a different degree.

In this article, we propose that the $f_{Realism}$ function (versus the $f_{Orientation}$ function) defines a prevalence order among the existing set of beliefs (versus desires) and the new belief (versus desire) inferred from a constitutive norm. Specifically, the $f_{Realism}$ function (versus the $f_{Orientation}$ function) determines which must be the degree of a belief about an institutional fact (versus a desire about a brute fact).¹⁵ Therefore, we can define different agent types according to the priority that they give to the formulas inferred from norms and the existing beliefs and desires.

5.1 Realism function: priority among beliefs

According to the priority relationship that can be defined between the existing beliefs¹⁶ and the beliefs that can be inferred from constitutive norms at some point, we classify agents into:

- *Observation-Realistic* agents are those ones that in case of inconsistency between a belief that is inferred from a constitutive norm and an existing belief determine that the existing belief prevails. This means that if an observation-realistic agent is able to observe an institutional fact (e.g. it is informed by a third party that its civil status is married), then this belief is more plausible than the inferences that it can make, at this specific moment, according to the constitutive norms that it knows. More formally, in an observation-realistic agent the $f_{Realism}$ function is defined as follows:

$$f_{Realism}(B, belief(BF, \rho)) = \begin{cases} B \cup \{belief(BF, \rho)\} & \text{if } \nexists \rho' : belief(BF, \rho') \in B \\ B & \text{if } \exists \rho' : belief(BF, \rho') \in B \end{cases}$$

¹⁴In this article, we define that two formulas $M(\gamma, \rho)$ and $M'(\gamma', \rho')$ are inconsistent when $M = M'$, $\gamma = \gamma'$, and $\rho \neq \rho'$; i.e. when the two graded formulas agree on their logical content but not in the degrees.

¹⁵Notice that this question is different from the problem of maintaining consistency (e.g. taking control about replicated or inconsistent formulas) in the mental sets (B, D, I) . This problem is out of the scope of this article and has been addressed by other proposals. For example, in [25] Joseph proposes a coherence-based mechanism for solving inconsistencies in graded BDI agents.

¹⁶Note that the existing beliefs might contain also beliefs that have been inferred from constitutive norms at some point in the past.

- *Norm-Realistic* agents are those ones that in case of inconsistency between a belief that is inferred from a constitutive norm and an existing belief determine that the belief inferred from the constitutive norm prevails. This means that if a norm-realistic agent is able to infer an institutional fact (e.g. it knows it has formalised a marriage contract and it is able to infer that it is married), then this conclusion is more plausible than the information that it already believes. More formally, in a norm-realistic agent the $f_{Realism}$ function is defined as follows:

$$f_{Realism}(B, belief(BF, \rho)) = \begin{cases} B \cup \{belief(BF, \rho)\} & \text{if } \nexists \rho' : belief(BF, \rho') \in B \\ B \setminus \{belief(BF, \rho')\} \cup \{belief(BF, \rho)\} & \text{if } \exists \rho' : belief(BF, \rho') \in B \end{cases}$$

5.2 Orientation function: priority among desires

According to the priority relationship that can be defined between the existing desires¹⁷ and the desires that can be inferred from constitutive norms at some point, we classify agents into:

- *Desire-Oriented* agents are those ones that in case of inconsistency between a desire that is inferred from a constitutive norm and an existing desire determine that the existing desire prevails. This means that if a desire-oriented agent determines that it desires to achieve a state of affairs in which a brute fact holds (e.g. it does not want to exceed 50 Km/h to minimize petrol consumption), then this consequence is more important than the inferences he can make, at this specific moment, according to the constitutive norms that it knows. More formally, in a desire-oriented agent the $f_{Orientation}$ function is defined as follows:

$$f_{Orientation}(D, desire(IF, \rho)) = \begin{cases} D \cup \{desire(IF, \rho)\} & \text{if } \nexists \rho' : desire(IF, \rho') \in D \\ D & \text{if } \exists \rho' : desire(IF, \rho') \in D \end{cases}$$

- *Norm-Oriented* agents are those ones that in case of inconsistency between a desire that is inferred from a constitutive norm and an existing desire, the desire inferred from the constitutive norm prevails. This means that if a norm-oriented agent is able to derive a desire to achieve a brute fact (e.g. it desires to avoid driving offences and it desires not to exceed 50 Km/h since this counts-as a driving offence inside a town), then this conclusion is more important than the desires that it already has. More formally, in a norm-oriented agent the $f_{Orientation}$ function is defined as follows:

$$f_{Orientation}(D, desire(IF, \rho)) = \begin{cases} D \cup \{desire(IF, \rho)\} & \text{if } \nexists \rho' : desire(IF, \rho') \in D \\ D \setminus \{desire(IF, \rho')\} \cup \{desire(IF, \rho)\} & \text{if } \exists \rho' : desire(IF, \rho') \in D \end{cases}$$

In this section, we have provided an agent typology and a descriptive characterization of each agent type. Section 7 describes a set of experiments that have been carried out to evaluate the performance of our proposal for reasoning about constitutive norms and providing an experimental characterization of the different agent types. The next section illustrates how a Normative BDI agent reasons about constitutive norms with our proposal.

¹⁷Note that the existing desires might contain also desires that have been inferred from constitutive norms at some point in the past.

6 Illustrative example

This example shows how an agent employs our proposal for reasoning about constitutive norms. Specifically, this example shows how an agent uses constitutive norms for extending its mental theory and how constitutive norms affect the decision making process.

6.1 Initial situation

Let us suppose that there are two agents a and b which are ‘a couple’. In this example, we will focus our attention in agent a , which is a Normative BDI agent that makes use of our proposal for reasoning about constitutive norms. Agent a considers that a couple are two agents that are in love and that live together. According to these conditions, agents a and b are a couple. Thus, a has a belief corresponding to being a couple with b with a certainty degree equal to 1—i.e. $\text{belief}(\text{couple}(a, b), 1)$ belongs to B . Regarding motivations of agent a , let us suppose that it wants to be married with agent b with the highest intensity—i.e. $\text{desire}(\text{married}(a, b), 1)$ belongs to D . Moreover, it also wants to gain a mortgage with agent b with the highest intensity—i.e. $\text{desire}(\text{mortgage}(a, b), 1)$ belongs to D . Figure 3 shows how the different formulas are generated in the normative and cognitive elements of the Normative BDI agent.¹⁸ Next, the different reasoning phases of this agent are described.

6.2 Reasoning process

Agent a executes Algorithm 1. Thus, it starts the reasoning cycle by perceiving the environment and updating the beliefs and norms. In the following we illustrate how formulas are generated by each reasoning process.¹⁹

Perception Phase: Norm Acquisition Rules. Marriage is an institutional fact that is defined by a constitutive norm that claims that if any pair of agents X, Y , which are a couple²⁰—represented by the expression $\text{couple}(X, Y)$ —, formalise a marriage contract—represented by the expression $\text{formalised}(\text{marriage}, X, Y)$, then it *counts-as* as they are married—represented by the expression $\text{married}(X, Y)$. This norm is formally defined as follows:

$$\langle \text{formalised}(\text{marriage}, X, Y), \text{married}(X, Y), \text{couple}(X, Y) \rangle$$

Similarly, gaining a mortgage is an institutional fact that is defined by a constitutive norm that claims that if any pair of agents X, Y , which are married²¹—represented by the expression $\text{married}(X, Y)$ —, formalise a deed contract—represented by the expression $\text{formalised}(\text{deed}, X, Y)$ —, then it *counts-as* as they gain a mortgage—represented by the expression $\text{mortgage}(X, Y)$. This norm is formally defined as follows:

$$\langle \text{formalised}(\text{deed}, X, Y), \text{mortgage}(X, Y), \text{married}(X, Y) \rangle$$

¹⁸The generation of intentions from desires is not relevant to this case study and, as a consequence, the intentions have not been included in Figure 3.

¹⁹For simplicity we do not include the explanation of the reasoning processes when they do not alter the mental state of agent a ; i.e., when the set of formulas remains unchanged.

²⁰Here, we assume that marriages of convenience or arranged marriages between strangers are not accepted. Therefore, agents that do not maintain a relationship (i.e. are not a couple) cannot be married even if they formalise a marriage contract.

²¹Here, we assume agents need to be married to gain a mortgage. Of course, in real-life, people are not required to be married for signing mortgages.

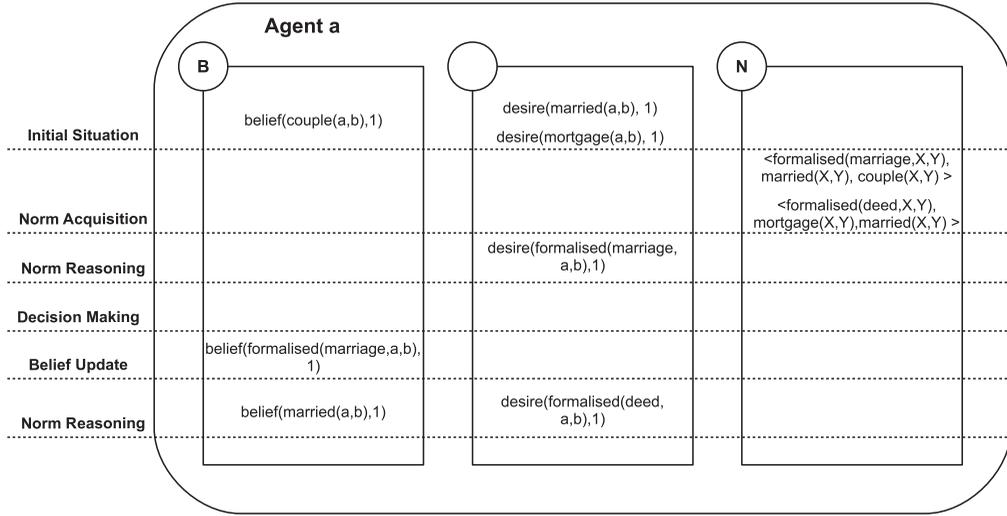


FIG. 3. Belief, desire and normative sets of agent *a*. These sets are represented as rectangles. The horizontal dashed lines determine the formulas that are inserted in the different reasoning phases.

Let us assume that agent *a* acquires these constitutive norms by consulting a public norm repository (e.g. the OMS in the THOMAS framework [15]) or artifact (e.g. the NormativeBoard in the ORA4MAS framework [23]). Thus, agent *a* updates its normative set accordingly.

Deliberation Phase: Norm Reasoning Rules. The marriage norm is used by the agent to determine how the desire of being married can be translated in terms of properties of the environment that can be modified by the agent. Specifically, the ‘Desiring to bring about brute facts’ rule is executed and a new desire to formalise a marriage contract is created as follows:

$$\frac{\exists norm(\langle BF, IF, C \rangle) \in N \wedge \exists \sigma : desire(\sigma(IF), \rho_{\sigma(IF)}) \in D \wedge belief(\sigma(C), \rho_{\sigma(C)}) \in B}{\langle B, D, I, N \rangle \rightarrow \langle B, D', I, N \rangle}$$

where $BF = formalised(marriage, X, Y)$; $IF = married(X, Y)$; $C = couple(X, Y)$;
 $\sigma = \{X/a, Y/b\}$; $\rho_{married(a,b)} = 1$; $\rho_{couple(a,b)} = 1$;
 $D' = f_{Orientation}(D, desire(formalised(marriage, a, b), f_{Expansion}(1, 1)))$

Considering the proposed definition of $f_{Expansion}$ (see Equation 4.2):

$$f_{Expansion}(1, 1) = 1 \times 1 = 1$$

Since the agent has not any desire about the expression $formalised(marriage, a, b)$, then the new desire inferred from the constitutive norm is inserted into the desire set.²² Thus, a new desire is generated inside the desire set—i.e. a desire such as $desire(formalised(marriage, a, b), 1)$ is inserted into D . This implies that the agent has been able to create a more specific desire about formalising a contract that will allow it to achieve a more abstract desire about being married.

²²Note that there is not an inconsistency between the existing desires and the desire inferred from the constitutive norm and, in this case, both norm-oriented and desire-oriented agents behave equally.

Decision-Making Phase. Now rules for making a decision about the next action to be performed are considered. Mainly, this process consists of generating plans for reaching the desired state given that the agent knows the existence of actions that could achieve it. For example, agent a knows that a contract among two agents is formalised when both agents sign this contract and the contract is registered. Thus, the agent generates an intention to execute a plan for signing and registering the contract. Then, this plan is executed.

Perception Phase: Belief Update Rules. As a consequence of the execution of the plan, agent a formalises a marriage contract and updates its beliefs accordingly—i.e. $\text{belief}(\text{formalised}(\text{marriage}, a, b), 1)$ is inserted into B .

Deliberation Phase: Norm Reasoning Rules. Since the agent believes that the marriage constitutive norm is relevant (i.e. agent a believes that agent a and b are a couple) and it believes that a marriage contract has been formalised, then the ‘Believing institutional facts being true’ rule is applied. As a consequence, a new belief that represents that the agents a and b are married is created as follows:

$$\frac{\exists \text{norm}(\langle BF, IF, C \rangle) \in N \wedge \exists \sigma : \text{belief}(\sigma(BF), \rho_{\sigma(BF)}) \in B \wedge \text{belief}(\sigma(C), \rho_{\sigma(C)}) \in B}{\langle B, D, I, N \rangle \rightarrow \langle B', D, I, N \rangle}$$

where $BF = \text{formalised}(\text{marriage}, X, Y)$; $IF = \text{married}(X, Y)$; $C = \text{couple}(X, Y)$;
 $\sigma = \{X/a, Y/b\}$; $\rho_{\text{formalised}(\text{marriage}, a, b)} = 1$; $\rho_{\text{couple}(a, b)} = 1$;
 $B' = f_{\text{Realism}}(B, \text{belief}(\text{married}(a, b), f_{\text{Expansion}}(1, 1)))$

Again, $f_{\text{Expansion}}(1, 1) = 1$ and the belief set is updated—i.e. $\text{belief}(\text{married}(a, b), 1)$ is inserted into B . Thanks to this belief, the abstract desire of being married can be retracted, since it has been achieved. This process is carried out by the *desire update rules* on Algorithm 1 that are responsible for creating new desires and dropping those ones that have been achieved.

Moreover, the new belief about the marriage institutional fact allows agent a to determine how the desire of gaining a mortgage can be translated in terms of states of affairs that can be achieved by it. Specifically, the mortgage constitutive norm is a second-order norm: i.e. a constitutive norm whose context is expressed in terms of institutional facts. In this case the ‘Desiring to bring about brute facts’ rule is executed and a new desire to formalise a deed contract is created as follows:

$$\frac{\exists \text{norm}(\langle BF, IF, C \rangle) \in N \wedge \exists \sigma : \text{desire}(\sigma(IF), \rho_{\sigma(IF)}) \in D \wedge \text{belief}(\sigma(C), \rho_{\sigma(C)}) \in B}{\langle B, D, I, N \rangle \rightarrow \langle B, D', I, N \rangle}$$

where $BF = \text{formalised}(\text{deed}, X, Y)$; $IF = \text{mortgage}(X, Y)$; $C = \text{married}(X, Y)$;
 $\sigma = \{X/a, Y/b\}$; $\rho_{\text{mortgage}(a, b)} = 1$; $\rho_{\text{married}(a, b)} = 1$;
 $D' = f_{\text{Orientation}}(D, \text{desire}(\text{formalised}(\text{deed}, a, b), f_{\text{Expansion}}(1, 1)))$

As a result, a new desire is generated inside the desire set—i.e. a desire such as $\text{desire}(\text{formalised}(\text{deed}, a, b), 1)$ is inserted into D . This implies that the agent has been able to create a more specific desire about formalising a deed that will allow it to achieve a more abstract desire about gaining a mortgage.

The marriage example is a metaphor for how agents are able to perform those actions (e.g. following the specific protocols and sequences of actions) that entail the modification of the institutional state (e.g. the performance of trading operations).

7 Experimental results

This section evaluates the performance of our proposal when agents reason about constitutive norms. We have carried out several experiments to evaluate if our approach allows agents to keep track of the institutional state precisely.²³ Specifically, we demonstrate that: (i) our proposal allows agents to keep track of the institutional state with more precision than agents that simply observe the institutional state; (ii) even if agents are able to perceive the institutional facts, it is better to complement this information with the information that is derived from reasoning about constitutive norms; and (iii) agents that use their perceptions to reason about constitutive norms and making inferences about the institutional state are not more sensitive to the precision of these perceptions than agents that simply observe the institutional state. Moreover, these experiments allow us to provide a characterization of the different agent types.

7.1 Experiment description

We have considered a scenario in which there is a set of agents that belong to the same institution (e.g. an e-market). In this institution a set of constitutive norms has been registered in a repository. All agents have access to the norm repository. Constitutive norms regulate the creation of institutional facts (e.g. constitutive norms may define the specific actions that *count-as* trading operations). Agents may perform different actions during their execution causing changes in the environment (e.g. sold items may be delivered) and in the institutional state (e.g. a purchase contract is executed). Agents are able to observe their environment. Moreover, agents are able to observe a part of the institutional state (e.g. there may be a facilitator agent in the e-market that informs about the execution of the purchase contracts signed by a specific agent). The rest of institutional facts are not visible by agents and their truth-value can only be determined by means of constitutive norms (e.g. an agent that observes that the actions that *count-as* executing a purchase contract have been performed by another agent can determine that the purchase contract has been executed).

We have built a simulator of this scenario with the parameters that we sum up in Table 1. Algorithm 3 contains the pseudo code of this simulator. According to this algorithm, in our simulator there is a set of agents that belong to an institution in which there are 10 constitutive norms.²⁴ Agents have access to a repository that contains the information about these constitutive norms (this corresponds to the instruction named *Norm Acquisition* on Algorithm 3). Once the agents have read the repository of constitutive norms, they observe the environment to keep track of the institutional state. Specifically, in each step of the experiment they observe the changes in the environment and in the institutional state. With this information they update their beliefs and the norms. Then, they carry out the deliberation phase and the decision-making phase. The experiment is executed during 400 steps. Moreover, each experiment has been repeated 1000 times to support the findings.

Each agent acts as a binomial classifier that determines which of the institutional facts hold and which ones not. In each step, we inspect the belief set of agents and compare the estimation made by agents against the institutional state. Specifically, in each iteration we update: the number of true positives (*TP*), which is the number of times that an agent considers that an institutional fact is true and it is actually true; the number of true negatives (*TN*), which is the number of times that an agent

²³For simplicity, we will only focus on detecting the institutional changes; i.e. this experiment only takes into account how agents extend their belief base. As a consequence, the results described in this section only take into account the generation of beliefs from constitutive norms. However, similar results were obtained when we also considered the generation of desires.

²⁴Note that we assume that the set of constitutive norms is not changed during an experiment.

TABLE 1. Parameters used in the norm recognition experiment

Parameter	Value
# of agents	3
# of constitutive norms	10
# of steps	400
# of repetitions	1000
# of institutional facts	10
Visibility of institutional facts	0.25
# of atomic formulas	20
Agent accuracy	[0,1]
Opacity Coefficient	0.5

Algorithm 3 Pseudocode of algorithm executed by our simulator

```

function SIMULATOR
   $N \leftarrow \text{randomNormCreation}()$  ▷ Norm Definition
   $IF, BF \leftarrow \text{randomEnvironmentCreation}()$  ▷ Environment Definition
   $\mathcal{A} \leftarrow \text{createAgents}()$  ▷ Agent Creation
  for  $a \in \mathcal{A}$  do
     $a.N \leftarrow \text{NormAcquisitionRules}(N)$  ▷ Norm Acquisition
  end for
  while the simulator has not been executed all the steps do
     $Actions \leftarrow \emptyset$ 
     $VIF \leftarrow \text{VisibleInstitutionalFacts}(IF, \text{visibility})$  ▷ Visibility of Institutional Facts
    for  $a \in \mathcal{A}$  do
       $p \leftarrow \text{PerceiveEnvironmentRules}(VIF, BF, a.\text{accuracy}, \text{opacity})$  ▷ Perception
       $a.B \leftarrow \text{BeliefUpdateRules}(a.B, p)$ 
       $a.D \leftarrow \text{DesireUpdateRules}(a.D, a.B, a.I)$  ▷ Deliberation
      if  $a$  is a normative agent then
         $a.B, a.D \leftarrow \text{NormReasoningRules}(a.B, a.D, a.N)$ 
      end if
       $R \leftarrow \text{UpdateResults}(R, a.B, IF)$  ▷ Inspection
       $I \leftarrow \text{IntentionGenerationRules}(a.B, a.D, a.I, \text{PlanLibrary})$  ▷ Decision Making
       $Actions \leftarrow Actions \cup \text{ActionSelectionRules}(a.I)$ 
    end for
     $IF, BF \leftarrow \text{evolveEnvironement}(IF, BF, N, Actions)$  ▷ Environment Evolution
  end while
  return  $\text{CalculateMCC}(R)$ 
end function

```

considers that an institutional fact is not true and it is actually false; the number of false positives (FP), which is the number of times that an agent considers that an institutional fact is true and it is not true; and the number of false negatives (FN), which is the number of times that an agent considers that an institutional fact is false and it is actually true.

Matthews Correlation Coefficient. The Matthews Correlation Coefficient (MCC) [4] is used as a measure of the quality of binary classifications. It takes into account true and false positives and negatives and is generally regarded as a balanced measure that can be used even if the classes are of very different sizes. The MCC can be calculated using the formula:

$$MCC = \frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}}$$

The MCC is in essence a correlation coefficient between the observed and predicted binary classifications; it returns a value between -1 and $+1$. A coefficient of $+1$ represents a perfect prediction, 0 an average random prediction and -1 an inverse prediction. While there is no perfect way of describing the confusion matrix of true and false positives and negatives by a single number, the Matthews correlation coefficient is generally regarded as being one of the best measures.

7.1.1 Norm definition

In each experiment, 10 constitutive norms are randomly created (this corresponds to the instruction named *Norm Definition* on Algorithm 3). These norms regulate the creation of 10 institutional facts. As previously mentioned, the institutional state is not fully visible and the truth-value of some of the institutional facts cannot be observed. To simulate this, we consider that the truth-value of any institutional fact is observable with a probability of 0.25 (this corresponds to the instruction named *Visibility of Institutional Facts* on Algorithm 3). When it is not observable, agents have to use constitutive norms to infer the truth-value. Specifically, agents have to observe their environment to determine whether the brute fact (*BF*) and the context (*C*) of a constitutive norm hold or not. We assume that for each constitutive norm $\langle BF, IF, C \rangle$ the elements *BF* and *C* are defined as an atomic formula that represents a state of affairs. Specifically, we consider a set of 20 different atomic formulas. Thus, for each constitutive norm the brute fact and the context of each are randomly assigned to a proposition of this set.

7.1.2 Environment definition

As previously mentioned, in this experiment we assume that the physical environment is described in terms of 20 different atomic formulas. Each one of these formulas represents a state of affairs (i.e. a brute fact) of the physical environment. At the beginning of each experiment, these states of affairs are randomly defined as holding or not; i.e. the truth-value of the 20 atomic formulas is randomly defined as true or false (this corresponds to the instruction named *Environment Definition* on Algorithm 3). We also assume that the institutional environment is described by 10 institutional facts. At the beginning of each experiment, the institutional facts are defined as holding or not. In this case, the initial truth-value of the institutional facts is defined considering the constitutive norms. Therefore, only those institutional facts that can be inferred from the initial state of the environment and the constitutive norms are true. The rest of institutional facts are initially false.

As the agents execute actions during the experiment, both the physical and the institutional environment evolve (this corresponds to the instruction named *Environment Evolution* on Algorithm 3). Therefore, the truth-values of the 20 atomic formulas that represent the physical environment change

according to the actions performed by agents. Similarly, the truth-value of the 10 institutional facts change as a consequence of the changes in the physical environment.²⁵

In our simulator we define the *opacity* of the environment as a feature that determines the difficulty to observe the environment. Specifically, the *opacity* coefficient is a real value that determines the magnitude or the error made by all agents when they observe the environment. The more opaque the environment is, the more difficult is for agents to observe it. Thus, the error made by agents when they observe the environment would increase as the opacity increases. In the experiments we fixed the *opacity* coefficient to 0.5.

7.1.3 Agent implementation

As previously mentioned, at the beginning of each experiment agents read the norm repository and, as a consequence, all of them know the same 10 constitutive norms defined in each experiment.²⁶ The initial belief and desire bases of each agent are randomly defined at the beginning of the experiment (this corresponds to the instruction named *Agent Creation* on Algorithm 3). Specifically, agents initialize the certainty of the brute and institutional facts as random real values within the $[0, 1]$ interval. Recall that these initial beliefs will be updated during the agent execution by the perception and deliberation phases (see Algorithm 3). Agents also pursue a set of goals that represent desired (versus undesired) situations that the agent wants to achieve (versus avoid). For simplicity, we assume that agents pursue different desires that are randomly generated during the experiment.²⁷

According to Algorithm 3, in each step of the experiment agents execute three different phases (see Figure 1).

1. *Perception*. Agents are able to observe their environment. Specifically, they are able to observe the truth-value of the atomic formulas that describe the physical and the institutional state. However, the agents have limited capabilities for observing their environment. The accuracy of each agent to observe its environment is randomly assigned to a real value within the $[0, 1]$ interval. The higher the accuracy, the more exact the observations are. Hence, the observations of agents are affected by a random normally distributed noise.²⁸ We consider a normally distributed noise with mean 0.0 and a varying standard deviation depending on the agent accuracy and the difficulty to observe the environment (i.e. the *opacity* coefficient). Specifically, we consider the distribution:

$$\mathcal{N} \sim (0, (1 - accuracy) * opacity)$$

where the value of *opacity* coefficient is shared by all agents.

Thus, in each step agents observe their environment and apply belief update rules for translating their observations into beliefs. In this experiment we have defined the following belief

²⁵Recall that, the truth-values of institutional facts are always consistent with the state of the environment and the constitutive norms.

²⁶Note that our goal is to evaluate the rules that we propose for reasoning about constitutive norms. Thus, we want that the only difference among the agents is the way in which they consider constitutive norms, not the specific norms that they know.

²⁷Note that we assume that agents always have at least one desire that can be pursued.

²⁸Notice that we denote the falsity value as 0 and the true value as 1.

update rules:

$$\frac{\text{observation}(\alpha, \rho) \wedge \nexists \rho' : \text{belief}(\alpha, \rho') \in B}{\langle B, D, I, N \rangle \rightarrow \langle B', D, I, N \rangle}$$

where $B' = B \cup \{\text{belief}(\alpha, \rho)\}$

$$\frac{\text{observation}(\alpha, \rho) \wedge \exists \rho' : \text{belief}(\alpha, \rho') \in B}{\langle B, D, I, N \rangle \rightarrow \langle B', D, I, N \rangle}$$

where $B' = B \setminus \{\text{belief}(\alpha, \rho')\} \cup \{\text{belief}(\alpha, \rho)\}$

where α is an atomic formula representing a state of the environment or the institution, and ρ is the certainty of this observation.

2. *Deliberation.* Agents revise their desires in the deliberation phase. Specifically, in this phase agents remove their desires from the desire set as soon as they achieve them.

Moreover, in the deliberation phase agents apply the norm reasoning rules for keeping track of the institutional state. In this experiment, we create agents that have different capabilities for reasoning about norms:

Non-Normative. These agents are BDI agents that do not use our methods for reasoning about constitutive norms and they are not able to take constitutive norms into account.

Norm-Realistic These agents are Normative BDI agents that use our methods for reasoning about constitutive norms and they are able to take constitutive norms into account. Therefore, they can determine the truth-value of institutional facts according to their observations and the inferences that they made by means of constitutive norms. These agents are *norm-realistic* and give more priority to the information that they infer, when they are able to both observe and infer the truth-value of an institutional fact.

Observation-Realistic These agents are quite similar to *norm-realistic* ones. However, they give more priority to the information that they observe, when they are able to both observe and infer the truth-value of an institutional fact.

In each experiment we create one agent of each type, thus a total of three agents.

After the deliberation phase, we inspect the agent beliefs to determine which of the institutional facts are believed to be true or false according to the information that each agent has (this corresponds to the instruction named *Inspection* on Algorithm 3). As aforementioned, the BDI agents used in this article use graded logics to represent mental propositions. Thus, they have graded beliefs that represent their knowledge about the institutional state. In general, when a belief about a proposition has a low certainty, then the agent considers that the proposition is false. In this experiment, we assume that institutional facts are false when their certainty is lower than an internalization threshold ($\delta_{\text{internalization}}$).

Since the different types of agents carry out a different reasoning process, we have to determine the most suitable values for the internalization threshold ($\delta_{\text{internalization}}$) for each agent type. To this aim, we have performed a set of previous experiments varying the value of the threshold. In each of the previous experiments, a set of three agents (one agent of each type) are informed about 20 constitutive norms. In each iteration, agents perceive their environment and determine (i.e. conjecture) which institutional facts hold and which ones not. The conjecture made by agents is compared against the institutional state and the number of *TP*, *TN*, *FP* and *FN* is updated accordingly. The experiment is executed during 400 steps. Once the experiment ends we calculate the MCC achieved by each agent type. For each value of the threshold we have performed 1000 experiments. Figure 4 shows the MCC obtained on average in the

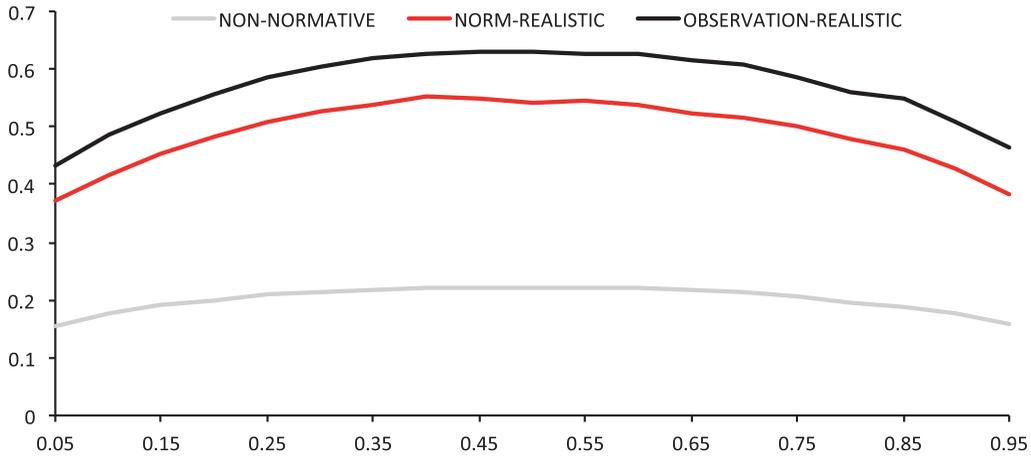


FIG. 4. MCC (Y axis) with respect to the internalization threshold (X axis).

100 experiments by each agent type with respect to the value of $\delta_{internalization}$. As illustrated by Figure 4, the best results of non-normative agents are obtained when $\delta_{internalization}$ is 0.5, the best results of norm-realistic agents are obtained when $\delta_{internalization}$ is 0.4, and the best results of observation-realistic agents are obtained when $\delta_{internalization}$ is 0.45. In the rest of the experiment types, we have fixed the internalization threshold to the values that allow each agent type to obtain its best results.

3. *Decision Making*. In this phase, agents generate intentions to execute one plan that achieves one or more desires. For simplicity, agents select randomly one action to be executed in each step of the experiment.²⁹

7.2 Results

We have performed 3 different experiment types to illustrate the performance of our methods for reasoning about constitutive norms with respect to: (i) number of steps that the experiment is executed; (ii) the visibility of the institutional state; and (iii) the opacity of the environment. The results of these experiment types are described below.

7.2.1 Number of steps

This experiment is aimed at determining the number of steps that the simulator must be executed to obtain stable results (i.e. results in which the MCC does not vary much). Moreover, this experiment demonstrates that in a typical situation³⁰ our proposal allows agents to keep track of the institutional state with more precision.

Figure 5 illustrates the MCC achieved per each agent type on average with respect to the number of steps that the experiment is executed. The error bars represent the 95% confidence interval for the

²⁹Note that this experiment is aimed at evaluating the performance of our proposal for reasoning about constitutive norms when agents keep track of the institutional state. Thus, the functionality in charge of updating the agent desires and carrying out the decision-making phase has been simplified in the simulator (i.e. agents select the next action to be executed randomly).

³⁰A typical situation is when most of the institutional facts are not visible (i.e. when $visibility=0.25$) and the variance of the noise that affects agent observations is half of the agent accuracy (i.e. when $opacity=0.5$)

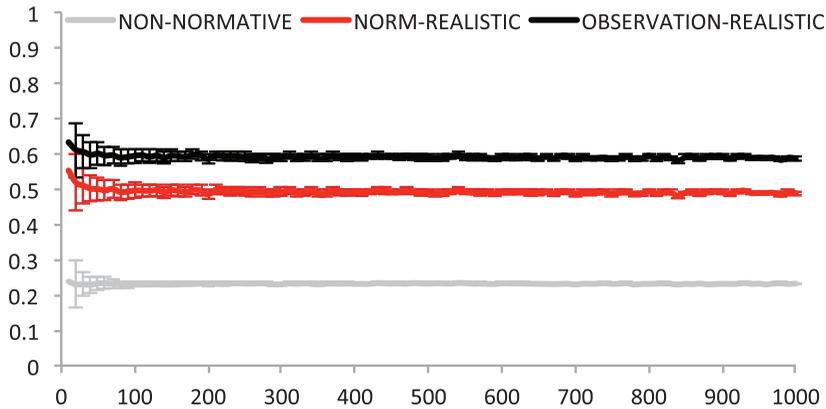


FIG. 5. MCC (Y-axis) achieved per agent type with respect to the number of steps that the experiment is executed (X-axis). The error bars (Y-axis) represent the 95% confidence interval for the MCC obtained in the experiments.

MCC obtained in the experiments. Non-normative agents keep track of the institutional state with less precision. On average, the MCC obtained by non-normative agents is 0.23. This is explained by the fact that they can only determine the truth-value of institutional facts when they are able to observe it. On the contrary, agents that use our proposal for reasoning about constitutive norms are able to keep track of the institutional state with more precision. Specifically, the MCC obtained by norm-realistic and observation-realistic agents is 0.49 and 0.59, respectively. The improvement on the MCC achieved by norm-realistic and observation-realistic is 112.38%³¹ and 154.91%, respectively. In light of these results, we can conclude that agents that use our methods for reasoning about constitutive norms attain a significant improvement in this experiment. From 400 steps, the results achieved, in terms of the average MCC and the 95% confidence intervals for the MCC, remain quite stable. For this reason, we have fixed the number of steps to 400 in the following experiment types.³²

7.2.2 Visibility of institutional state

With this experiment we aim to demonstrate that agents that consider the information that is inferred from reasoning about constitutive norms keep track of the institutional state with more precision, even if agents can perceive the institutional facts. Specifically, this experiment illustrates the performance of the different agent types when the visibility of the institutional state changes. Moreover, we have determined which type of normative agent is more suitable according to the visibility of the institutional state.

As explained before, the visibility of the institutional state is simulated by the probability in which the truth-value of any institutional fact is observed. In this experiment we analyse the MCC achieved by each agent when the probability of observing institutional facts varies within the $[0, 1]$ interval.

Figure 6 illustrates the MCC achieved per each agent type on average with respect to the probability that the truth-value of any institutional fact is observable. In this picture we can observe

³¹The improvement on the MCC is calculated as $(0.49 - 0.23) / 0.23 \times 100 = 112.38\%$

³²Note that in the following experiment types the 95% confidence intervals for the MCC are quite small and we have preferred not to draw them.

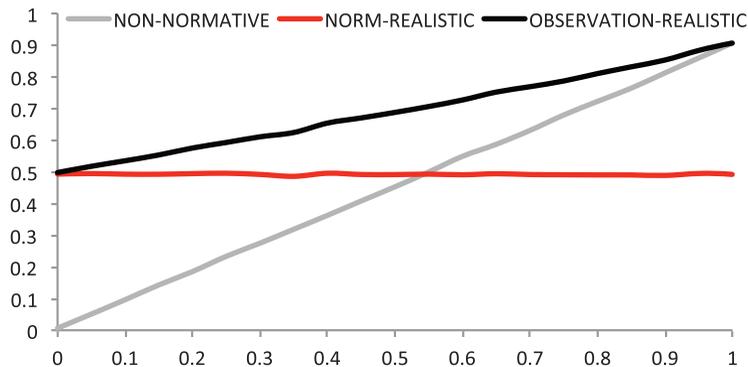


FIG. 6. MCC (Y-axis) achieved per agent type with respect to the visibility of institutional facts (X-axis).

that the MCC achieved by norm-realistic agents is not affected by the visibility of the institutional facts. This is due to the fact that norm-realistic agents consider that the information that is inferred has more priority than the information that is already believed. In our simulator there is one constitutive norm per each institutional fact, then norm-realistic agents are able to infer from constitutive norms the truth-value of all institutional facts. As a result the performance of norm-realistic agents is not affected by the visibility of institutional facts. In contrast, the MCC achieved by observation-realistic and non-normative agents increases linearly with the visibility. Since non-normative agents only can determine the truth-value of institutional facts on the basis of their observations, and thus they are the most affected by the visibility of institutional facts (i.e. the slope of the line that represents the MCC achieved by non-normative agents is higher).

When the probability is very low (i.e. $visibility \ll 0.1$), then observation-realistic and norm-realistic agents achieve the best results. In this case, agents are not able to observe the truth-value of almost any institutional fact and they only have the information that they can infer from constitutive norms. Therefore, observation-realistic and norm-realistic behave similarly. As the probability increases, then more institutional facts can be observed and the performance of both observation-realistic and non-normative agents gets better. Specifically, observation-realistic agents achieve the best results, when the probability is within the interval $[0.05, 0.95]$. When the truth-value of all the institutional facts can be observed (i.e. when the probability is 1), then non-normative and observation-realistic agents achieve the best results. In this case, observation-realistic agents behave as non-normative agents. In this situation, our methods for reasoning about constitutive norms are not necessary since agents are able to observe the institutional state completely and there is no need for reasoning about constitutive norms.

In light of these results, we can conclude that observation-realistic agents are the most suitable ones. They behave as norm-realistic when their institutional state is not visible and they rely on the inferences that they made from constitutive norms. However, when the institutional state is visible, observation-realistic agents behave as non-normative agents and rely on their observations. When the institutional state is partially visible, observation-realistic agents combine both information that is inferred from norms and information that is observed and they achieve the best results. Thus, we demonstrate that agents that make use of our proposal for reasoning about constitutive norms obtain the best results whatever the visibility of institutional facts is.

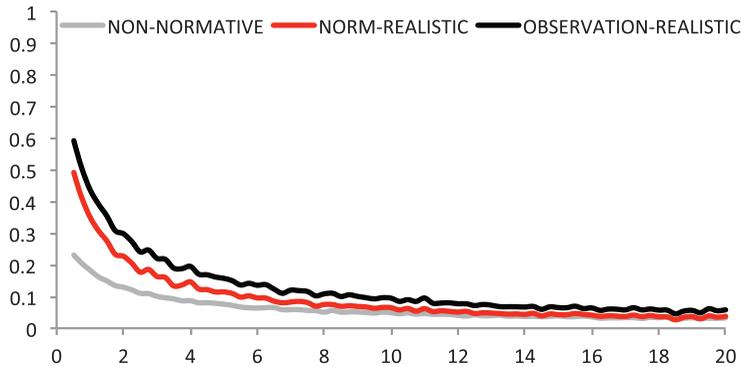


FIG. 7. MCC (Y-axis) achieved per agent type with respect to the opacity of the environment (X-axis).

7.2.3 Opacity of the environment

In our proposal, agents make use of their perceptions to: (i) update their beliefs; and (ii) reason about constitutive norms and infer beliefs and desires about the institutional state. Given that in our proposal agents make a further use of their perceptions, it may be argued that our proposal makes agents more sensitive to the precision of these perceptions. In this experiment, we aim to demonstrate that agents that consider the information that is inferred from constitutive norms are similarly affected by the precision of perceptions. Specifically, this experiment illustrates the performance of the different agent types as the *opacity* of the environment increases and the environment becomes more and more difficult to discern.

As explained before, the opacity of the environment determines the difficulty to observe the environment. It has been modelled by the *opacity* coefficient that is shared by all agents and determines the magnitude of the error made by all agents when they observe the environment.³³

Figure 7 illustrates the MCC achieved per each agent type on average with respect to opacity of the environment. In all agent types the MCC decreases as the opacity of the environment increases. In light of these results, we can conclude that the opacity of the environment affects all agent types in a similar way. Thus, observation-realistic agents obtain the best results regardless of the opacity of the environment. This is explained by the fact that the opacity determines the precision of observations about institutional and brute facts. The observations about institutional facts are used by all agents to determine the truth-value of institutional facts. The observations about brute facts are used by norm-realistic and observation-realistic agents when they reason about constitutive norms to determine the truth-value of institutional facts. Thus, all agent types use observations to keep track of the institutional state somehow or other. As a consequence, all agent types are affected by the opacity of the environment.

7.3 Discussion

As shown in the results provided in this section, the methods for reasoning about constitutive norms proposed in this article allows agents to keep track of the institutional state with more precision.

³³Note that the observation of the environment includes the observation of the brute facts and those institutional facts that are visible. The more opaque the environment is, the less precise the observations about brute and institutional facts are. Because of this, both non-normative and normative agents are affected in the same way by the opacity of the environment.

Besides that, we have demonstrated that observation-realistic agents are the ones that keep track of the institutional state with the highest precision. This is explained by the fact that observation-realistic agents are able to adapt their behaviour; i.e. they behave as norm-realistic agents when the institutional state is not visible, they behave as non-normative when the institutional state is visible, and they combine the information that they observe and infer from constitutive norms when the institutional state is partially visible. Finally, we demonstrate that our agents are the ones that keep track of the institutional state with the highest precision regardless of the opacity of the environment.

8 Conclusions

As mentioned before, the role of constitutive norms on agent reasoning processes is an open problem that has been identified by previous proposals, but that has not been faced yet. To address this problem, in this article we propose an information model, knowledge representation and an inference mechanism for reasoning about constitutive norms in BDI agents. Our proposal considers that agents have a physical presence in some real-world environment. Agents have a partial and uncertain knowledge of it. The main aim of our proposal is to endow agents with capabilities for reasoning about constitutive norms and being able to know the institutional repercussions of their actions. We have provided a classification of the different agent types that can be developed with our proposal. Moreover, we have also shown how our proposal works in a case study. Finally, we have carried out an experimental evaluation of the performance of our methods for reasoning about constitutive norms and the results are shown in this article.

The set of constitutive norms considered by an agent might be in conflict, since agents belong to different institutions simultaneously. Thus, the development of mechanisms for resolving conflicts and inconsistencies among constitutive norms is an interesting issue that will be addressed in future work. Another interesting line of future work is the interpretation of constitutive norms. In this article, we assume that norms are unambiguously interpreted by agents. However, in environments that are populated by heterogeneous agents, agents might give different interpretations to constitutive norms and, as a consequence, inconsistencies and conflicts may arise (e.g. agents may consider that a constitutive norm is relevant in different contexts).

References

- [1] H. Aldewereld, S. Álvarez-Napagao, F. Dignum, and J. Vázquez-Salceda. Making norms concrete. In *Proceedings of International Conference on Autonomous Agents and Multiagent Systems*, van der Hoek, Kaminka, Lespérance, Luck and Sen, eds. pp. 807–814. International Foundation for Autonomous Agents and Multiagent Systems, 2010.
- [2] H. Aldewereld. Autonomy vs. conformity: an institutional perspective on norms and protocols. *The Knowledge Engineering Review*, **24**, 410–411, 2009.
- [3] G. Andrighetto, D. Villatoro, and R. Conte. Norm internalization in artificial societies. *AI Communications*, **23**, 325–339, 2010.
- [4] P. Baldi, S. Brunak, Y. Chauvin, C. Andersen, and H. Nielsen. Assessing the accuracy of prediction algorithms for classification: an overview. *Bioinformatics*, **16**, 412, 2000.
- [5] I. Bloch. Information combination operators for data fusion: A comparative review with classification. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, **26**, 52–67, 1996.

- [6] G. Boella, J. Broersen, and L. van der Torre. Reasoning about constitutive norms, counts-as conditionals, institutions, deadlines and violations. In *Intelligent Agents and Multi-Agent Systems*, T. D. Bui, T. V. Ho and Q. T. Ha, eds, pp. 86–97. Springer, 2008.
- [7] G. Boella and L. Van Der Torre. Regulative and constitutive norms in normative multiagent systems. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, D. Dubois, C. A. Welty and M.-A. Williams, eds, pp. 255–265. AAAI Press, 2004.
- [8] G. Boella and L. van der Torre. Constitutive norms in the design of normative multiagent systems. In *Proceedings of the International Conference on Computational Logic in Multi-Agent Systems*, F. Toni and P. Torroni, eds, pp. 303–319. Springer, 2005.
- [9] P. Bourdieu. *Practical Reason: On the Theory of Action*. Stanford University Press, 1998.
- [10] J. Broersen, M. Dastani, J. Hulstijn, Z. Huang, and L. van der Torre. The BOID architecture: conflicts between beliefs, obligations, intentions and desires. In *Proceedings of International Conference on Autonomous Agents and Multiagent Systems*, pp. 9–16, 2001.
- [11] A. Casali, L. Godo, and C. Sierra. A graded BDI agent model to represent and reason about preferences. *Artificial Intelligence*, **175**, 1468–1478, 2011.
- [12] R. Conte, C. Castelfranchi, and F. Dignum. Autonomous norm acceptance. In *Intelligent Agents V: Agents Theories, Architectures, and Languages*, J. Müller, M. P. Singh and A. S. Rao, eds, pp. 99–112. Springer, 1999.
- [13] N. Criado, E. Argente, and V. Botti. A BDI architecture for normative decision making. In *Proceedings of International Conference on Autonomous Agents and Multiagent Systems*, W. van der Hoek, G. A. Kaminka, Y. Lespérance, M. Luck and S. Sen, eds, pp. 1383–1384. IFAAMAS, 2010.
- [14] N. Criado, E. Argente, and V. Botti. Open issues for normative multi-agent systems. *AI Communications*, **24**, 233–264, 2011.
- [15] N. Criado, V. Julian, V. Botti, and E. Argente. A norm-based organization management system. In *Coordination, Organizations, Institutions, and Norms in Agent Systems V*, vol. 6069, J. Padget, A. Artikis, W. Vasconcelos, K. Stathis, V. T. da Silva, E. Matson and A. Polleres, eds, pp. 19–35. Springer, 2010.
- [16] N. Criado. *Using Norms To Control Open Multi-Agent Systems*. PhD Thesis, Universitat Politècnica de València, 2012.
- [17] M. Fitting. *First-order Logic and Automated Theorem Proving*. Springer, 1996.
- [18] G. Governatori, J. Gelati, A. Rotolo, and G. Sartor. Actions, institutions, powers: Preliminary notes. In *Proceedings of the International Workshop on Regulated Agent-Based Social Systems: Theories and Applications*, vol. 318, G. Lindemann, D. Moldt, M. Paolucci and B. Yu, eds, pp. 131–147. University of Hamburg, 2002.
Place published Hamburg, Germany Publisher University of Hamburg
- [19] G. Governatori and A. Rotolo. Bio logical agents: norms, beliefs, intentions in defeasible logic. *Autonomous Agents and Multi-Agent Systems*, **17**, 36–69, 2008.
- [20] D. Grossi, H. Aldewereld, J. Vázquez-Salceda, and F. Dignum. Ontological aspects of the implementation of norms in agent-based electronic institutions. *Computational & Mathematical Organization Theory*, **12**, 251–275, 2006.
- [21] D. Grossi and F. Dignum. From abstract to concrete norms in agent institutions. In *Formal Approaches to Agent-Based Systems*, M. G. Hinchey, J. L. Rash, W. F. Truszkowski and C. A., Rouff, eds, pp. 12–29. Springer, 2005.
- [22] D. Grossi, J. Meyer, and F. Dignum. Counts-as: classification or constitution? An answer using modal logic. *Deontic Logic and Artificial Normative Systems*, pp. 115–130, 2006.

- [23] J. Hübner, O. Boissier, R. Kitio, and A. Ricci. Instrumenting multi-agent organisations with organisational artifacts and agents. *Autonomous Agents and Multi-Agent Systems*, **20**, 369–400, 2010.
- [24] A. Jones and M. Sergot. A formal characterisation of institutionalised power. *Logic Journal of IGPL*, **4**, 427, 1996.
- [25] S. Joseph. *Coherence-Based Computational Agency*. Phd Thesis, Universitat Autònoma de Barcelona, Barcelona, 2010.
- [26] M. Kollingbaum and T. Norman. NoA-a normative agent architecture. In *Proceedings of the International Joint Conferences on Artificial Intelligence*, vol. 18, pp. 1465–1466. Morgan Kaufmann Publishers Inc., 2003.
- [27] A. Nonnengart and C. Weidenbach. Computing small clause normal forms. In *Handbook of automated reasoning*, Vol. 1, A. Robinson and A. Voronkov, eds, pp. 335–367. Elsevier Science and MIT Press, 2001.
- [28] G. Plotkin. A structural approach to operational semantics. 1981.
- [29] J. Rawls. Two concepts of rules. *The Philosophical Review*, **64**, 3–32, 1955.
- [30] R. Rubino, A. Omicini, and E. Denti. Computational institutions for modelling norm-regulated mas: an approach based on coordination artifacts. In *Coordination, Organizations, Institutions, and Norms in Multi-Agent Systems*. Vol. 3913 of *Lecture Notes in Computer Science*, pp. 127–141. 2006.
- [31] J. R. Searle. *The Construction of Social Reality*. Free Press, 1997.
- [32] J. Searle. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press, 1970.
- [33] J. Searle. *Making the Social World: The Structure of Human Civilization*. Oxford University Press, 2010.
- [34] V. Silva. From the specification to the implementation of norms: an automatic approach to generate rules from norms to govern the behavior of agents. *Autonomous Agents and Multi-Agent Systems*, **17**, 113–155, 2008.
- [35] J. Vázquez-Salceda. The role of norms and electronic institutions in multi-agent systems applied to complex domains. The HARMONIA framework. *AI Communications*, **16**, 209–212, 2003.

Received 20 September 2012