



MODELLING DYNAMIC NORMATIVE UNDERSTANDING IN AGENT SOCIETIES

CHRISTOPHER K. FRANTZ*, MARTIN K. PURVIS, BASTIN TONY ROY SAVARIMUTHU† AND MARIUSZ NOWOSTAWSKI‡

Abstract. Agent-based Modelling appears as a promising analytical tool when it comes to a lasting question: in how far did different institutions affect the social and economic outcomes of societies? Taking an incremental step to address this question, we present a refined approach that combines existing institution representations (the *structure*) with a norm identification *process* to systematically ‘grow’ normative understanding from the bottom up without relying on any prior knowledge. The proposed mechanism provides agents with the ability a) to detect complex normative behaviour by developing and differentiating stereotypes of social actors, and b) to generalise behaviour beyond observed social entities, giving agents the ability to develop normative understanding as a potential precursor for predicting newcomers’ behaviours. We exemplify this approach using a simulated prototypical trader scenario that is evaluated with respect to behavioural diversity (different compositions of non-/cooperative agents) as well as structural diversity (different types of agents). Using the simulation results, we showcase the explanatory power of the derived normative understanding beyond the interpretation of quantitative results, and finally discuss the generalisability of the proposed approach.

Key words: Norms, Institutions, Institutional Grammar, Norm Generalisation, Norm Synthesis, Nested ADICO, nADICO, Dynamic Deontics, Normative Understanding, Maghribi Traders Coalition, Social Simulation, Agent-Based Modelling, Agent-Based Institutional Modelling

AMS subject classifications. 68T42, 68U20

1. Introduction. Institutional modelling has received increasing attention in the area of multi-agent systems, and multi-agent-based simulation, such as in [17, 2, 45]. One central driver is the continued interest in explaining socio-economic development based on the institutional environment that either fostered or restrained economic development, which is a key theme of the area of New Institutional Economics [48]. Agent-based modelling is particularly useful in this context, since it can model human interaction on multiple levels of social organisation (micro, meso, macro) and provide insight into the emergence of behavioural regularities.

This work contributes to this discussion by proposing mechanisms that allow us to ‘grow’ institutional understanding using the very individuals that are subject to the established institutions. In this context, we interpret institutions from a behavioural perspective as *manifestations of social behaviour* that transcend the individuals that constitute them and exist in the form of conventions (e.g. use of left or right road side), social norms (e.g. littering in the park), or formal laws and rules (e.g. contracts, traffic rules). The proposed mechanism is an extension of the approach previously introduced in [15] and attempts to make bottom-up emergence processes accessible using a human-readable structural representation (introduced in previous work [13, 14]) that captures the institutional spectrum. The concrete contribution includes a systematic formalisation of the individual norm generalisation steps, along with proposing a novel approach to multi-level norm generalisation based on social structures.

Doing so, this work attempts to integrate ‘structure’ and ‘process’ to provide a generalisable framework that allows for an accessible representation of institutional emergence, with particular focus on social norms. In the area of normative modelling, this work bears relationships to two research areas, namely institution representations as well as norm synthesis.

In the following Section 2 we lay out the motivation in more detail in the light of existing work from the area of normative modelling. In Section 3 we introduce the utilised structural and conceptual representations, before moving towards the description of the proposed norm generalisation process in Section 4. In Sections 5 and 6 we apply this mechanism to a prototypical scenario that models economic cooperation based on informal institutions to show the process’ inner workings as well as its explanatory value. We conclude our work with a discussion of the contribution and its potential applications in Section 7.

*College of Enterprise and Development, Otago Polytechnic, New Zealand (christopher.frantz@op.ac.nz)

†Department of Information Science, University of Otago, New Zealand ({martin.purvis, tony.savarimuthu}@otago.ac.nz)

‡Faculty of Computer Science and Media Technology, Gjøvik University College, Norway (mariusz.nowostawski@hig.no)

2. Background, Related Work & Motivation. Coordination in human societies is increasingly recognised as important, if not decisive, to set societies on a path towards prosperity [19, 1]. Realistic artificial models of human societies thus afford mechanisms that are able to replicate behaviour leading to the emergence of different forms of stabilising behavioural manifestations, or institutions. Even though this challenge is not novel in the context of computational modelling but goes back to the pioneering works on norm synthesis by Shoham and Tennenholtz [38, 40], recent works have made promising advances in the context of online norm synthesis, thus inferring normative behaviour at runtime, as opposed to identifying norms *ex post* (offline detection). Recent works in this area specifically include Morales et al. [28, 29, 27] and Riveret et al. [37] that put emphasis on learning-centred bottom-up inference of norms.

Morales et al. [27] propose a framework that identifies and selects norms based on their effectiveness and efficiency. It does so using a central monitor that observes behaviours and performs the balancing act of proposing the greatest possible generalisation while avoiding overregulation. Their model applies case-based reasoning to inform future action choices and uses the concept of a *Normative Network* to represent the interrelationships between systems of norms and to resolve norm conflicts. The efficiency of their system is described in the context of traffic scenario inspired by Shoham and Tennenholtz [39].

Riveret et al. [37] propose a transfiguration mechanism in which agents learn norms from the bottom up and use collective action to arrive at a consensus about the normative content, before enforcing this agreed-upon prescription collectively. The authors' intent is to retrace the establishment of self-governing systems by marrying bottom-up and top-down perspective on institutions, which they operationalise using stochastic games.

In contrast to the previous works the approach proposed in this work is not primarily informed from a technological perspective that yields towards the engineering of efficient normative outcomes. Instead of controlling norm emergence with top-down processes, we emphasise the sociological perspective, with the intent to remodel characteristics typical for actual human societies, including complex behaviour that produces undesirable or even seemingly inconsistent results. We further tolerate the establishment of inefficient norms (see e.g. [6]) that the regulatory artefacts of the mentioned works prevent by design.

We entertain a comparatively broad understanding of institutions [31, 19], interpreting institutions as 'manifestations of social behaviour'. Though being concise, this understanding captures the essence of institutions, which includes a) their grounding in interaction among individuals and/or groups, and b) their stability characteristics that allows institutions to potentially survive the circumstances (and thus purpose) that brought them about (see e.g. [35]). This understanding captures a broad range of institution types that reach from self-enforcing conventions, via decentrally enforced (informal) social norms, to centrally enforced (formal) rules. For this reason in this work we seek to operationalise a general representation for institutions, such as that found in Crawford and Ostrom's Grammar of Institutions [11], whose specification has previously been refined in [13]. For our purposes the effectiveness of the 'grammar' lies in its accessible human-readable interpretation, consideration of social structures (e.g. actors) as well as actions, along with its cross-disciplinary applicability (see e.g. [42, 41, 17]). In the context of social simulation, we can thus use it as an expressive interface between the experimenter and the observed artificial society.

The present work continues the refinement of this structural representation (building on [13] and [14]) and the associated systematic process (initially introduced in [15]) that describes how individuals can develop normative understanding based on generalised experience and observations.

The essential contribution of this work is firstly to provide a systematic formalisation of the individual norm generalisation steps, and secondly, to introduce a notion of multi-level generalisation of social structures, the latter of which provides the foundation to generalise behaviour beyond the observed social entities and potentially enabling the prediction of newcomers' behaviours. The introduced example further showcases how this generated normative understanding can support the interpretation of emerging behaviour, a challenge that is particularly prevalent in complex social scenarios.

This process is driven by three key themes, that is a) the inherent decentralised nature without central coordinating entity, b) the strong generalisability by making minimal assumptions about cognitive capabilities of agents as well as absence of a priori knowledge about actions or social structures, and c) the accessibility to the human experimenter.

To illustrate our present work, we invoke the metaphor of a long-distance trading scenario inspired by the Maghribi Traders Coalition [19] that operated around the 10th century along the North African coast. This scenario describes a collective that relied informal trade arrangements: traders delegated the transport and sale of goods to fellow traders in remote locations under the promise to reciprocate that service, an aspect that allowed them to expand the geographic range of their operations. The essential aspect for our scenario is the reliance on mutual compliance to assure a functioning economic system; individuals that were suspected of embezzling profits, i.e. cheating, faced exclusion from trade. At a later stage, we will complement that society with a counterpart that had to rely on formal institutions to govern compliance, the Genoese traders. They could not rely on a trust network to assure cooperation, but had to resort to contracts in order to sustain a functioning economy. We operationalise those scenarios to provide an accessible evaluation scenario with prototypical institutional configurations, but do not claim to recreate those societies accurately.

In our prototypical societies traders can at the same time adopt two roles: 1) *sender*, and 2) *operator*. Senders send goods to other operators who then facilitate the actual sale (and potentially cheat), and return the realised profits to the sender.

3. Norm Representation.

3.1. Nested ADICO (nADICO). The concept of Nested ADICO (nADICO) [13] builds on the essential purpose of the original institutional grammar [11] to represent conventions (or shared strategies), norms and rules in the form of *institutional statements*. It consists of five components (with the acronym ADICO), and is briefly explained in the following.

- **Attributes (A)** – describe the characteristics of individuals or groups of individuals that are subject to an institution;
- **Deontic (D)** – indicates whether the institutional statement is of prescriptive or permissive nature. The grammar relies on deontic primitives (e.g. may, must, must not) [46] to express the normative or discretionary nature of a statement. The
- **AI/m (I)** – describes an action or outcome associated with the institutional statement;
- **Conditions (C)** – describe the circumstances under which a statement applies, which can include spatial, temporal and procedural aspects. If unspecified, the applicability is unconstrained and thus implies “at all times and in all places” [11]. Finally, the
- **Or else (O)** – specifies the consequences or sanctions [36] that follow non-compliance with a statement specified by the previous four components. In the nADICO interpretation, consequences can both be attached to social norms and formal rules, whereas the original concept only permits the specification for consequences for rules. The ‘Or else’ itself can be a nested institutional statement.

This grammar allows for the expression of statements of varying nature and strength, representing different institution types, while allowing a straightforward transformation from natural language.

A convention, for example, can be adequately expressed using the components AIC, e.g.:¹ **Traders (A) trade fair (I) when being observed (C)**. Adding the *Deontic* component to the (then ADIC) statement extends it to a sanction-less norm statement: **Traders (A) must (D) trade fair (I) when being observed (C)**. Finally, adding a consequence (*Or else*), constitutes a norm or a rule [13]:²

Traders (A₀) must (D₀) trade fair (I₀) when being observed (C₀), or else observers

(A₁) must (D₁) report this (I₁) in any case (C₁).

Institutional statements can be *nested vertically* (as shown above), in which a *consequential statement* (Level 1 on the previous example) backs a statement it monitors (above: ADIC(ADIC)) (here the *monitored statement* on Level 0 in the example). This allows the representation of institutional regress, i.e. the backing of institutions by more deeply embedded institutions that describe monitors’ duties and potential consequences for their non-compliance. In contrast to the original institutional grammar, in nADICO both norms and rules can have specified sanctions. This conception follows the understanding that norms and rules are not so much differentiated by the absence of sanctions, but rather by sanction characteristics. In the case of rule sanctions, prescribed actions, monitoring entities, and enforcers are explicitly specified (i.e. formally appointed). For

¹In the examples we use parenthesised letters to define the relationship of natural language expressions with the respective grammar component, such as ‘Traders (A)’ referring to the *Attributes* component.

²We use indices on statement components to indicate nesting levels (0 being the top level, 1 being the first nested level, etc.).

norms, in contrast, sanctions are characterised by varying/unknown actions as well as fuzzy (and potentially conflated) conceptions of private monitors and enforcers [35].³

Ideally this enables the modelling of comprehensive chains of responsibility. Beyond this, institutional statements can be *horizontally nested*, i.e. combined by logical operators that describe their co-occurrence (e.g. ADIC *and* ADIC) or alternative occurrence (inclusive or: ADIC *or* ADIC; exclusive or: ADIC *xor* ADIC). In addition to that, statements can be negated using the *not* operator. *Horizontal Nesting*, as we refer to such co-occurrences, can be used both in the leading monitored statement and trailing consequential statement of a nested institutional statement. The syntax (in the Extended Backus-Naur Form [24]) is defined in Figure 3.1.

```

attributes = "A" ;
deontic    = "D" ;
aim       = "I" ;
conditions = "C" ;
aic       = attributes , aim , conditions ;
adic      = attributes , deontic , aim , conditions ;
and       = "and" ; (* Conjunction *)
or        = "or" ; (* Inclusive Disjunction *)
xor       = "xor" ; (* Exclusive Disjunction *)
not       = "not" ; (* Negation *)
ws        = " " ; (* Whitespace *)
LB        = "(" ;
RB        = ")" ;
nadico    = adic | (* Individual Norm Statement without Sanction *)
           [ not , LB , ( nadico | aic ) , RB ] | (* Negation *)
           [ nadico [ , LB , ( nadico | aic ) , RB ] ] | (* Vertical Nesting *)
           [ LB , ( nadico | aic ) , ws , ( and | or | xor ) ,
             ws , ( nadico | aic ) , RB ] ; (* Horizontal Nesting *)
statement = aic | nadico ; (* nADICO Statement (including Top-Level Conventions) *)

```

Fig. 3.1: nADICO Grammar Syntax Specification

In nADICO, vertical nesting represents notions of *structural institutional regress*, i.e. the activation of one or more consequential norms upon violation of monitored norm(s).⁴ Complementing this, horizontal nesting enables the representation of institutional complexity on a given nesting level. Combining both nesting approaches enables the comprehensive description of individual institutions.

The concept of nested institution levels has previously been discussed by Grossi et al. [20] with focus on the specification of interdependent norm types based on *substantive norms* that are decomposed into *check norms* (that monitor behaviour) and *reaction norms* (that prescribe sanctions associated with violations). López y López et al. [26] propose a similar approach to manage the interdependency between norms. Similar to Grossi et al. they separate norms into primary and secondary norms, but interlock those. The violation of normative goals of secondary norms activates the corresponding primary norms. The nADICO grammar affords the expression of such dependencies in a single comprehensive structural representation. The function of the norm is relative to its nesting level in an institutional statement, guarding a statement on a lower nesting level (i.e. being a consequential statement), while at the same time being guarded by statements on a higher nesting level (i.e. being a monitored statement).

To facilitate the derivation of normative understanding from behavioural feedback, we further utilise a continuous conception of deontics (as opposed to discrete deontic primitives) to reflect the evolution of normative understanding over time and integrate it with the institutional grammar.

³This is discussed in greater detail in [13], [14], and [16].

⁴The conventional understanding of institutional regress (as used by Aoki [3] and Hodgson [22]) emphasises the path-dependent formation of institutions over time, an aspect we label as *temporal institutional regress*, in contrast to the structural dependency highlighted by the institutional grammar.

3.2. Dynamic Deontics. In place of the institutional grammar’s use of conventional discrete deontic primitives to signify prohibitions, permissions, and obligations, we apply a continuous notion of deontics, previously introduced as Dynamic Deontics [14], which we briefly introduce at this stage. The essential intuitions of this approach are shown in Figure 3.2 and explored in the following.

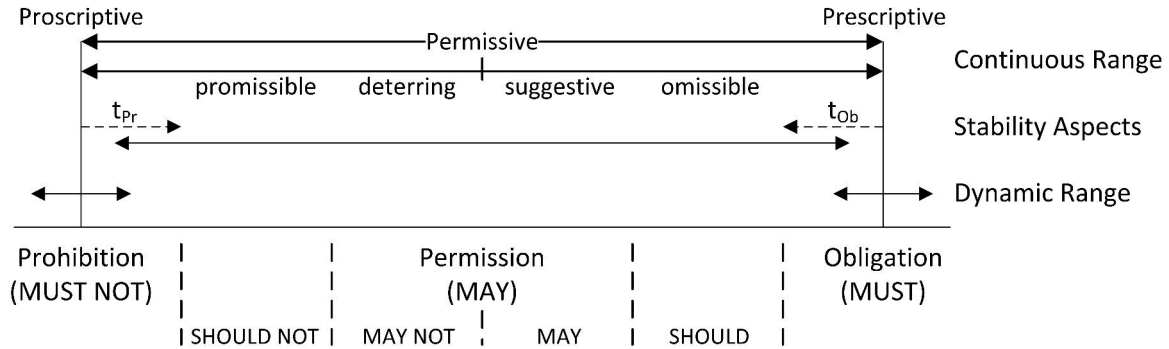


Fig. 3.2: Dynamic Deontic Range

Instead of relying on a simplified discretised representation of normative states, a continuous representation of deontics bridges the gaps between the extremal prohibitions and obligations, reflecting gradual evolving normative understanding over time. Similarly, a continuous conception can reflect the varying notions of oughtness associated with specific prescriptions [7]. Along the continuum between prohibition and obligation, what we conceive as deontic range, we can allocate deontic terms, or *deontics*. The original deontic primitives occupy the extremes and centre point (deontic centre) of the deontic range. To establish basic categories, the deontic range is composed from deontic compartments that carry labels such as *must* and *may*, and indicate relative levels of prescriptiveness,⁵ as shown in the ‘omissible’ (obligations that can be foregone) or ‘promissible’ (prohibitions that can be ignored). The range itself is dynamic and determined by the individual’s experience, a possible mapping being the direct association of the most positive and the most negative experiences with the respective ends of the deontic range.

To operationalise this approach, we rely on reinforcement learning, or more specifically, Q-Learning [47]. It offers the direct representation of extremal reinforcements (highest and lowest Q-values) as well as discounting of reinforcements over time, operationalising features such as the expansion and contraction of the deontic range.

In the context of this work, Dynamic Deontics are used to develop normative associations with specific actions (e.g. developing the understanding that one is *obliged to trade fair* and *prohibited from cheating*). Along with this, we maintain sliding windows of past extremal (i.e. positive and negative) reinforcements to smooth the dynamics of the deontic range, from which deontic compartments are dynamically derived. For a detailed account on the involved operational steps refer to [16].

4. From Experiences to Institutions. Although conventions and norms surround us, we are often barely conscious of their existence and origin. Generally, norms implicitly emerge on the basis of experiential [33] and social [4] learning in the contexts of existing institutions [49], whereas rules are explicitly specified [32], which allows their unambiguous communication and adoption.

To model the subconscious processes of evolving institutional understanding, we employ a data-driven approach that utilises the data structures we have described to facilitate agents’ understanding of the normative environment involving a minimal amount of explicit reasoning. In addition, the operationalisation proposed here aims to be generalisable, operates decentralised (i.e. without a central coordinating entity), and relies on a minimal amount of prior information. To achieve this, in Subsection 4.1 we operationalise an institution

⁵Note that the choice of intermediate deontic terms for this example is not systematically grounded but is chosen with analogy to other fields, such as the keywords for requirements specification in *Request for Comment* (RFC) documents [8] by the Internet Engineering Task Force.

representation that has the descriptive power to capture individual actions (action instances) as well as higher-level institutions. Based on those principles, we present a systematic process that describes the transition and derivation of institutional statements from individual observations (Subsection 4.2 onwards).

4.1. Action Representation. To operationalise actions, we use the syntax of the nADICO grammar’s AIC component that augments an action definition (*AIM*) with the subject (*Attributes*) and context/conditions (*Conditions*). Utilising the term **act** to identify an individual action instance (as opposed to an institutional statement of the form **aic**), an *action statement* is thus **act(attributes, aim, conditions)**, where **aim** represents the action definition.

Attributes. Operationalising the *Attributes* component, we assume that individuals carry observable properties reflecting the notion of social markers individuals display in real life [43], such as name, ethnic background, gender, etc. Such markers can be of individual nature (i.e. attached to the personal identity) and social character (i.e. one’s social identity), and be temporary and/or modifiable (e.g. clothing, situational role), or permanent (e.g. physical features). For the context of this work we represent attributes as two sets, with the first set **i** representing individual characteristics and the second one (**s**) highlighting group features [21], providing the generic structure **attributes(i, s)**. To facilitate the differentiated treatment of attributes as part of the generalisation process, we assume that individuals are able to categorise social markers by their natural kind, such as ‘role’, ‘occupation’, ‘origin’, as opposed to treating those as unstructured tags as done in previous work [15]. The essential advantage of categorising markers is the ability to account for asymmetric sets of markers, but also to permit refined characterisations by permitting combinations of markers (e.g. multiple occupations).

Individual and social markers are represented as key-value pairs $\langle k, v \rangle$, where the key **k** represents a unique label⁶ that describes a relevant category (e.g. **occupation**), and the value **v** is the set of associated markers (e.g. **trader**). Consequently, individual and social markers are sets of such pairs, as exemplified in the Figure 4.1.⁷




Fig. 4.1: Attributes Structure

Actions. To establish unambiguous symbolic references to an action and its properties, we require a structured action specification. We define actions using a label **a** and an associated set of properties **p**, the set of which depends on the nature of the action. Substituting the *AIM* component of the institutional grammar, we can thus say **aim(a, p)**. Taking an example from our scenario, the central properties of the action ‘send goods to trader’ are the *object* that is dispatched (‘goods’), as well as the destined *target* (‘trader’). This *action definition* can thus be represented as **aim(send, {object, target})**, and instantiated as **aim(send, {goods, trader})**. Note that the representation shown here simplifies the *Attributes* component by substituting the trader’s complete attribute specification with the token ‘trader’.

Conditions. In addition, we tailor the *Conditions* component to capture the context of action execution by allowing the specification of a potentially related preceding action (e.g. as a reaction to a previous action) as the first element in order to represent action sequences, or behavioural regularities that can be precursors for identified institutions.

Table 4.1 provides an overview of the operationalised component specifications.⁸

4.2. Generalisation. Individuals generally and unintentionally engage in processes of ‘implicit social cognition’ [18], one of which is the social generalisation process of ‘stereotyping’. Stereotyping attributes individual observations of behaviour or features to social structures generalised from observation instances. As such this

⁶Note that the label does not need to carry any semantic weight other than uniquely identifying a given category.

⁷For the sake of brevity the value sets in our example only contain individual values.

⁸To keep expressions concise, for given examples the complete attributes specification of agents (i.e. their social markers) is substituted with their name (e.g. **trader1**).

Table 4.1: Component Specifications

Component	Structure	Example/Instance
Attributes	<code>attributes(i, s)</code> , with <code>i/s</code> being sets of individual/social marker pairs	<code>attributes({name={trader1}}, {role={sender}}, occupation={trader}, ethnicity={maghribi})</code>
Action Definition	<code>aim(a, p)</code> , with <code>a</code> being a natural language action descriptor, and <code>p</code> being a set of action properties	<code>aim(send, {object, target})</code>
Conditions	<code>conditions(act, c)</code> , with <code>act</code> being a preceding action, and <code>c</code> being a set of further conditions	<code>act(trader2, aim(trade, {goods}), conditions(act(trader1, aim(send, {goods, trader2}), *)))</code>
Action Statement	<code>act(attributes, aim, conditions)</code>	<code>act(trader1, aim(send, {goods, trader2}), *)</code>

process bears similarities to inductive reasoning, but in contrast to the latter, stereotyping occurs subconsciously and does not consider explicit probabilistic reasoning. Notwithstanding the disputed underlying causes for its existence and attributed oversimplification in social contexts, this process can lead to uncertainty reduction and efficiency enhancement, effects that are characteristic for institutions [31, 50]. Since stereotyping builds on the assumption that actors repeat previously displayed behaviour, it enables individuals to develop predictors to anticipate others' behaviours and to call upon previously executed successful reactions. This specific aspect makes consideration of stereotyping important when modelling institutional evolution, since it provides the mechanism that links agents' minds with their environment, shaping *beliefs* from presumed *habits*, or as Peirce [34] put it: “[the] essence of belief is the establishment of habit.” In a similar vein, Castelfranchi links these conceptions by suggesting that minds themselves are social institutions [9].

Apart from the resemblance of human behaviour in agents, adopting this approach for artificial agents minimises the cognitive load for agents, limiting the need for complex agent architectures, and, along with this, potentially the demand for computational resources.

Following this understanding, the proposed mechanism operates entirely based on collected action instances as specified above. As highlighted in the individual transitions in Figure 4.2, the collected action instances are generalised into *descriptive norms* (or conventions), before aggregating associated experience to produce *valenced expressions*, which are used to generate *injunctive norms*.⁹ In the following we provide a detailed description of these individual steps.

Modelling subconscious generalisation processes shifts the perspective from the observation of *individual behaviour instances* to the inference of *social behavioural regularities*, narrowing the gap to what we perceive as institutions. We thus perform an aggregation of individual action statements to form generalised AIC statements, which we consider equivalent to observed conventions, or, in this case, *descriptive norms*. Operationally we achieve this by grouping the observed action statements based on their individual components.

To explore this generalisation process, let us use a running example, a trade action instance. `trader2` trades goods in the role of an `operator`, after having been sent goods by another fellow trader `trader1` (in the role of a `sender`):

⁹*Descriptive norms* describe how norm participants behave, whereas *injunctive norms* specify how norm participants should behave [10].

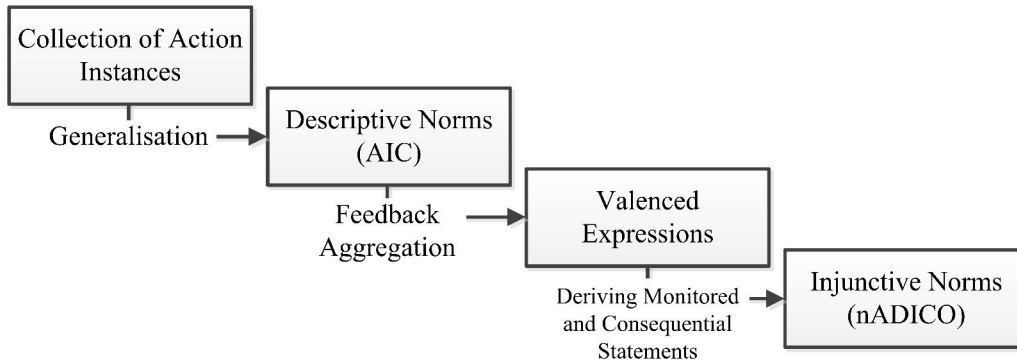


Fig. 4.2: Generalisation Process

```

act(attributes({name={trader2}}, {role={operator}, occupation={trader},
              ethnicity={maghribi}}),
    aim(trade, {goods}),
    conditions(act(attributes({name={trader1}}, {role={sender},
                            occupation={trader}, ethnicity={maghribi}}),
                  aim(send, {goods, attributes({name={trader2}},
                                                {role={operator}, occupation={trader},
                                                ethnicity={maghribi}})}),
                  conditions(*))))
  
```

Focal statement components with relevance for the generalisation of behavioural conventions or norms from individual observations include the *actor* (**attributes** component) and the *action* (**aim** component).

In an initial step, all observations are grouped by the action descriptor, ignoring additional action properties. In our running example this resolves to the action **trade**. In the resulting action statements we ignore the individual markers but concentrate on the aggregation based on social markers, i.e. markers that are shared among multiple individuals/groups.

The generalisation of grouped action statements can be attacked from two perspectives, one being the identification of *differences* among individual action instances in order to infer behavioural regularities that are specific to particular roles as described in [15].

An alternative is to approach generalisation from a bottom-up perspective by concentrating on *similar features* to identify behavioural regularities on different social aggregation levels. Let us exemplify this approach using a more realistic set of individual observations shown in Figure 4.3.

```

attributes(*, {role={sender},    occupation={trader},  origin={maghribi}})
attributes(*, {role={sender},    occupation={trader},  origin={maghribi}})
attributes(*, {role={operator},  occupation={trader},  origin={maghribi}})
attributes(*, {role={sender},    occupation={trader},  origin={genoese}})
attributes(*, {role={potter},    occupation={artisan}, origin={maghribi}})
  
```

Fig. 4.3: Observation Samples for Hierarchical Generalisation

To model more complex multi-level generalisation processes, we are inspired by agglomerative hierarchical clustering [25], i.e. the generation of clusters of decreasing granularity by incrementally increasing permissible distances of clustered elements. For the context of our work, we group individual observations based on their shared attribute categories/features, before incrementally reducing the number of shared categories to establish more coarse-grained generalisations that cluster wider ranges of observations.

The initial grouping of observations operates on all attributes features of given observations. We refer to this as **Level 0 Generalisation**, since it does not permit deviating attribute values in aggregated statements. Using the observation samples listed in Figure 4.3 we can group two observations on **Level 0**,¹⁰ with all other observations being unique with respect to their attribute sets as shown in Figure 4.4. For the sake of conciseness we only show combinations with at least two matching statements.

Level 2	{role=*, occupation={trader}, origin=*} (4), {role=*, occupation=*, origin={maghribi}} (4), {role={sender}, occupation=*, origin=*} (3)
Level 1	{role=*, occupation={trader}, origin={maghribi}} (3), {role={sender}, occupation={trader}, origin=*} (3), {role={sender}, occupation=*, origin={maghribi}} (2)
Level 0	{role={sender}, occupation={trader}, origin={maghribi}} (2)

Fig. 4.4: Hierarchical Generalisation

To model higher-level generalisations, we permit an increasing distance of observations. We operationalise the distance between observations as number of deviating category values, since our action representation operates based on nominal features. A **Level 5 Generalisation** would thus capture observations in which five features differ. A scenario consisting of n attribute categories thus permits a maximum generalisation level of $n - 1$, i.e. requiring at least one shared feature. In a scenario with six features the **Level 5 Generalisation** would represent the most coarse-grained generalisation.

Shifting to **Level 1**, we consequently permit a single deviating feature between attributes of statements considered for generalisation. From the sample set of five observations, **Level 1 Generalisation** results in three generalised attribute sets, one capturing all *Maghribi traders* and the second one describing *traders that operate as senders (of goods)*, as well as *Maghribi senders*.

Permitting **Level 2 Generalisation**, we cluster observations that showcase two deviating features. For our sample case, we extract three generalised attribute sets, *traders*, *Maghribis*, and *senders*.

Even though a uniform characterisation of attributes appears desirable, the established generalisations do not conflict, but reflect multiple possible perspectives of generalisation. The relevance of whose can potentially be determined by the number of observations captured, as well as the aggregated reinforcements of the associated actions that indicate the salience of behaviour (i.e. the extent to which all observations associated with this generalisation actually describe uniform behaviour), an aspect discussed in the upcoming Subsection 4.3.

Returning from the generalisation of attributes to the extraction of generalised observations, let us summarise the involved individual steps.

With S_{acts} being the collection of all individual observation action statements, S_{aims} being the individual aims, and the functions

- $attr(act)$ extracting a given action statement's (act) attributes, or all attribute components from a set of action statements, and
- $aim(act)$ extracting a given action statement's aim,

we can formalise the generalisation process as shown in Algorithm 1.

Initially all observations are grouped by the action they describe (**aim**), followed by the generalisation process that identifies intersections of attributes on different aggregation levels. The data structures keep track of the number of observations that back a generalised statement. This generalisation process operates across all attributes occurrences in observation statements, including attributes of preceding actions in the **conditions** component of an action statement.

This generalisation process marks the initial step of deriving normative understanding from individual observations in the form of descriptive norms, or conventions (AIC statements in nADICO (see Subsection 3.1).

¹⁰The number of captured observations per generalised attribute set is shown in parentheses.

Algorithm 1 Hierarchical Generalisation Algorithm

```

1: Initialise set Obs (observations grouped by aim);
2: Initialise set  $Gen_{(aim,level)}$  (generalised observations for aim aim and level level);
3: // Group action statements by aim
4: for aim in  $S_{aims}$  do
5:   for act in  $S_{acts}$  do
6:     if  $aim(act) == aim$  then
7:       Add act to set  $Obs_{aim}$  (Observations related to aim aim) contained in Obs;
8:     end if
9:   end for
10: end for
11: // Generalise grouped observations
12: for  $Obs_{aim}$  in Obs do
13:   Remove individual markers;
14:   // Determine maximum cluster distance based on number of attribute categories for observations
   grouped by action aim
15:    $maxDist = \max(attr(Obs_{aim})) - 1$ ;
16:   // Determine generalised statements for all feature distances
17:   for distance in  $maxDist$  do
18:      $level = maxDist - distance$ ;
19:     // Identify intersections of size level among attributes of grouped observations for action aim
     and save generalised statements
20:      $Gen_{(aim,level)} = \{x \in attr(Obs_{aim}), length(x) = level\}$ ;
21:   end for
22: end for
23: // At this stage  $Gen_{(aim,level)}$  contains generalised observations for all aims and all levels

```

From the perspective of the observing agent, this generalised understanding describes *how* agents in his social environment behave – independent of how they *should* behave.

Recalling the example provided at the outset of this subsection the generalisation process produces the following aic statement:

```

aic(attributes({role={operator}, occupation={trader}, ethnicity={maghribi}}),
  aim({action={trade}}, {object={goods}}),
  conditions(aic(attributes(*, {role={sender},
    occupation={trader}, ethnicity={maghribi}}),
    aim({action={send}, {object={goods}, target={
      attributes(
        {role={operator}, occupation={trader},
          ethnicity={maghribi}})})),
    conditions(*))))

```

From a normative perspective this translates into the understanding that ‘Maghribi traders in operator roles trade goods after having been sent goods by Maghribi traders acting as senders’.

In order to develop more complex institutional statements, beyond conventions or objectified descriptive norms, we need to assume that agents receive and associate feedback with individual action instances as part of their experiential learning process. This is not only important to develop a understanding of ‘oughtness’, as introduced in conjunction with the Dynamic Deontics conception (see Subsection 3.2), but also to support the evaluation of generalisation outcomes with respect to their contextual relevance as discussed in this subsection.

Any such feedback serves as input for a second aggregation process that converts the feedback into normative understanding. The conceptualisation and implementation of feedback naturally depends on the application scenario.

4.3. Feedback Aggregation. The central purpose of the feedback aggregation process is to build up information used for an agent’s overall understanding of a generalised convention. This is not to be confused with its attitude towards that convention, but instead represents the result of cyclic internalisation and socialisation processes that make agents experiential learners and shape their normative understanding, while, by means of

Table 4.2: Exemplified Action Instances with Valences

Simplified Action Example	Feedback
<code>act(attributes({name={trader1}}, {...}), aim(action={trade}), ...)</code>	30
<code>act(attributes({name={trader2}}, {...}), aim(action={trade}), ...)</code>	10
<code>act(attributes({name={trader3}}, {...}), aim(action={trade}), ...)</code>	-20
<code>act(attributes({name={trader4}}, {...}), aim(action={trade}), ...)</code>	20

conforming to normative behaviour, becoming enforcers in their own right. This aggregation operates on the action instances associated with a corresponding AIC statement as introduced in the previous Subsection 4.2. Table 4.2 shows (simplified) action instance observations for the action `trade`, along with hypothetical feedback values.

To derive an aggregated feedback, or *deontic value*, that is to be attached to the generalised institutional statement, we consider various *aggregation strategies*.

Since the purpose of this work is centred around the process of deriving institutional understanding from behaviour, as opposed to applying it to a specific scenario or types of agents, we concentrate on the rational strategy. The *rational aggregation strategy* is the purest analytical approach to deriving normative understanding, since it is based on the mean of all input values to determine a rationally expected feedback for future occurrences of a given action. Applying the rational deontic aggregation strategy, the exemplified observation sample in Table 4.2 would result in an aggregate feedback value (or *deontic value*) of 10. Recalling the deontic range conceptualisation discussed in Subsection 3.2, we can map this numeric value into an obligation, prohibition or permission that the individual associates with the evaluated action.

A (non-exhaustive) set of alternative possible strategies include the *opportunistic strategy*, in which individuals focus on extremal accounts of the most desirable or feared outcome (resolving to 30 for the sample in Table 4.2), as well as *optimistic* and *pessimistic* strategies that emphasise a bias towards best possible or worst possible outcomes (i.e. resolving to 30 and -20 respectively). The informed choice of those strategies can be used to express personal characteristics, such as a positivity or negativity bias [5] of individuals, but also to reflect cultural predispositions such as risk avoidance [23] when evaluating normative behaviour. The operationalisation of the mentioned alternative aggregation strategies is described in [16].

In preparation for the final step of this process of norm identification, the aggregated value – determined based on strategies discussed in this subsection – is associated with the generalised AIC statement as a precursor for the development of nADICO statements. For this purpose we rely on the concept of Dynamic Deontics to operationalise nADICO’s *Deontic* component.

4.4. Deriving nADICO Expressions. To derive higher-level nADICO expressions from AIC statements, we revisit the developed action sequences that not only reflect an individual’s actions but also multi-party actions. Actors can be generalised based on their social markers, such as roles (e.g. sender of goods, recipient, etc.). Consequently, agents do not only generate behavioural conventions that consider themselves or only of roles they play. In principle, individuals can form an institutional understanding about any role (or, more precisely, set of social markers) they observe, and use this knowledge predict unknown individuals’ behaviours based on their social markers. This satisfies an important purpose of institutions, the reduction of uncertainty [31], based on the ability to predict normative behaviour using the stereotypes attached to the observed individual. The data-driven mechanism to develop normative understanding introduced up to this stage gives agents the ability to generalise and abstractly observe different role perspectives, both of which are precursors for applying cognitive empathy [12], such as the ability to perform perspective taking.

However, to translate observed action instances from a mere descriptive (AIC statements) to a prescriptive perspective (ADIC/nADICO statements), we require the transformation of action sequences by *separating* those into *monitored statements* and *consequential statements* (see Subsection 3.1) as well as the assignment of deontic values.

In the following we describe the elementary steps involved in deriving nADICO statements, which include:

- Identification of monitored and consequential statements
- Assigning the deontic value to the monitored statement
- Inverting the deontic value and assigning it to the consequential statement

Identifying Monitored and Consequential Statements. In the course of their interaction, individuals collect action sequences, i.e. actions along with preceding actions, expressed as **act** statements with a subsequent **act** statement's **conditions** component (see Subsection 4.1). These action sequences represent action/reaction pairs that describe a 'because of', or 'on the grounds of' relationship, reflecting a direct causal dependency between both actions. Using our running example we would then arrive at the interpretation: "The Maghribi operator trades goods *because* he has been sent goods by the Maghribi sender.", which represents the descriptive norm perspective directly derived from observed behaviour. However, to represent an injunctive perspective that highlights an individual's perception of its duties [10], we require the transformation into '*Or else*' relationships for cases in which multiple actions occur in sequence. An example of such generalised action sequence is the initial statement '*Sender sending goods*', followed by '*Operator trading goods*'. To reflect the injunctive nature of the '*Or else*', we attach textual representations of perceived duty (deontics) to the individual statements and *invert* the derived consequential statement's deontic ('Senders *have to* send goods, or else Operators *will NOT* trade goods.'). As discussed in Subsection 3.2, note that the deontic terms associated with the deontic range (*may, should, must*) may not precisely reflect this understanding, but they capture the intuition of varying levels of duty.

To proceed along these lines, it is necessary to identify action sequences that involve one's social environment, i.e. multi-party interactions. Operationally, starting from the last element in the action sequence (i.e. the most recent action), the agent iterates backwards through the action sequence and identifies the first action statement whose attributes (i.e. generalised social markers) differ from the final action statement's attributes. The subsequence of all remaining action statements is then treated as *monitored statement* (i.e. the statement whose compliance is monitored), whereas the sequence of all parsed action statements is interpreted as *consequential statement*. Let us use the following example to explore these mechanics: '*Senders send goods*', followed by '*Operators cheating*', followed by '*Senders firing them*'. Working our way backwards from the last statement ('*Senders firing them*'), we find that the second last statement has a differing generalised actor (*Operators*), leaving us with the monitored statement '*Operators cheat after Senders sending goods*' and the consequential statement '*Senders firing them*'.

If an action sequence *does not* include multiple parties, i.e. all actions are performed by a single individual (or multiple individuals carrying identical social markers), the last elementary action is then interpreted as the previous action's consequence and treated as *consequential statement*, with the preceding action statements being the corresponding *monitored statement*.

Assigning Deontic Value to Monitored Statement. To give the statements normative value, we utilise the aggregated deontic value for the entire action sequence derived as part of the previously described feedback aggregation step (Subsection 4.3) and assign it to the identified monitored statement's *Deontic* component.

Assigning the Inverted Deontic Value to the Consequential Statement. Finally, to establish a deontic term's corresponding counterpart (e.g. *should* vs. *should not*), we rely on the individual's existing deontic range to invert the aggregated value. The inversion involves the proportional mapping of deontic values along the individual's deontic range to the inverted deontic compartment. The resulting deontic value is assigned to the consequential statement to express the presumed consequence for the violation of the inferred monitored statement. For example, assuming a deontic range with midpoint value of 0, a value of 5 – possibly mapping to *should* – is inverted to its opposite scale value and deontic term (*should not*) as schematically visualised in Figure 4.5. Note that the assumption of a symmetric deontic range is an exemplary simplification.¹¹

¹¹Alternative configurations are discussed in [16].




Fig. 4.5: Exemplified Deontic Inversion Process

Applying the deontic inversion to our running example, the corresponding nADICO statement reads:

```
adico(attributes({role={sender}, occupation={trader}}),
  deontic(5),
  aim(action={send}, {object={goods}, target={attributes(
    {role={operator}, occupation={trader}})})),
  conditions(*),
  orElse(adico(attributes({role={operator}, occupation={trader}}),
    deontic(-5),
    aim(action={trade}, {object={goods}}),
    conditions(*))))
```

Let us summarise the individual steps in algorithmic form.

Using $adic_k$ to describe an individual action statement at the k -th position within a sequence of statements $(adic_k)_{k=1}^{max}$ (with the first statement being the last performed action and max indicating the total number of action statements in the sequence), we can derive the corresponding nADICO statement $nadico$ and define the associated monitored statement (Function $nadic(nadico)$) and consequential statement (Function $o(nadico)$) as described in Algorithm 2.

In the context of the derivation of nADICO statements, it is important to highlight a special case. This derivation approach does not establish a consequential statement if no previous action has been observed, generating an injunctive norm without specified consequences, e.g. an obligation or prohibition without consequence. This corresponds to an ADIC statement in the nADICO syntax.

At this stage, the derived nADICO statements provide the experimenter with a comprehensive insight into individuals' experience-based normative understanding. The derived statements can be further generalised based on individual grammar components, such as

- *actions (Aims)* – by inferring an overall normative understanding of actions from aggregated nADICO statements – and
- *social structures (Attributes)* – by generalising normative behaviour for specific actor stereotypes (e.g. roles) on multiple generalisation levels (e.g. by occupation, by origin, etc.).

However, as in reality, this mechanism can only operate on action sequences known to an agent, i.e. actions an agent actually perceived. To afford the representation of deceitful behaviour and the associated information asymmetry, we assume that some actions are performed internally without being directly accessible to outside observers. Agent can thus perform more actions than can be perceived by observers. An example is the action of cheating, which may well be internalised by the cheating party, but will not be visible to the subject of cheating. Operationally, agents achieve this by internally maintaining action sequences that include the private actions they performed in addition to the publicly visible ones. The generalisation process itself is indifferent about the public or private nature of action statements in action sequences and operates on whichever statements have been perceived by the individual agent, potentially leading to an information asymmetry between the cheater (who knows he cheated) and the unknowing observer.

At this stage we will explore the introduced mechanism using the simulation scenario described at the outset of this article.

Algorithm 2 nADICO Statement Derivation Algorithm

```

1: // 1 is last (i.e. latest) statement in sequence, max is first (oldest)
2:  $k = 1$ ;
3: if  $max == 1$  then
4:   // Only one statement in sequence
5:   // Assign monitored statement
6:    $nadic(nadico) = adic_1$ ;
7:   // Assign aggregated feedback value (see Subsection 4.3) as monitored statement's deontic
8:    $d(nadic(nadico)) = d((adic_k)_{k=1}^{max})$ ;
9: else
10:  // More than one statement in sequence
11:  // Iterate backwards until attributes in statements differ
12:  while  $a(adick) == a(adick_{k+1})$  &  $k < max$  do
13:     $k++$ ;
14:  end while
15:  if  $k == max$  then
16:    // Make last statement consequential statement if no deviating attributes found
17:     $k = 1$ ;
18:  end if
19:  // Assign subsequence from last (i.e. most recent) action statement to action statement k as
    consequential statement
20:   $o(nadico) = (adic_l)_{l=1}^k$ ;
21:  // Assign all other action statements as monitored statement
22:   $nadic(nadico) = (adic_l)_{l=k+1}^{max}$ ;
23:  // Assign aggregated feedback value (see Subsection 4.3) as monitored statement's deontic
24:   $d(nadic(nadico)) = d((adic_k)_{k=1}^{max})$ ;
25:  // Assign the inverted deontic to consequential statement
26:   $d(o(nadico)) = inv(d(nadic(nadico)))$ ;
27: end if

```

5. Simulating Normative Understanding amongst Maghribi Traders. The Maghribi trader society highlighted in the beginning of this article provides an example of a medieval society that flourished based on informal cooperation – without the use of existing centrally enforceable institutional instruments (such as contracts). The mechanism described in this work thus bears characteristics that are compatible with this prototypical scenario, that is the decentralised nature, the informal normative perspective, and the purist approach towards inference of norms from experience and observation. We use this scenario to highlight the features of this approach, specifically the ability to inspect agents' minds individually, but also to develop an understanding of the collective normative landscape. Instead of exploring the test scenario comprehensively, we explore distributions of non-/cooperative behaviour under which general economic cooperation can be sustained, but more importantly, how it is reflected in the emerging normative understanding.

Traders, potentially acting as **sender** and **operator** establish a maximum number (*maxRelationships*) of mutual trade relationships to other traders based on random requests, to which they then send goods. The receiver trades those goods at a profit that is determined by a fixed factor *profitFactor*.¹² If initialised as cheater, the trader cheats by embezzling a fraction of the profit (as described below) and returns the remainder to the original sender. Sending agents memorise a parameterised number of last returned profits (as action feedback), which they can query for specific individuals or across all their partners in order to decide on the correctness of the returned profit. In cases of presumed cheating, traders can fire the suspect and memorise it as a cheater, or reward it with *wageFactor* of the generated profit. The interactions represent actions of the structure introduced in Subsection 4.1, with preceding actions stored in the *Conditions* component, successively building up action sequences that represent the comprehensive transactions between individual agents.

To reflect the ongoing trade relationship, rewards are represented as the accumulated profit from previous interactions with the rewarding party (i.e. the profits a respective trader had generated for the service provider over time). If the operator's cheating was not detected, the cheater adds the embezzled fraction as part of his

¹²We use a fixed factor since it simplifies the detection of cheating. The focus of this work does not lie in the evaluation of an agent's ability to detect cheating, but to understand the impact of cheating on the emerging normative understanding produced by the generalisation process.

reward from a trade interaction. In the case of firing, the fired trader memorises the inverted trade value and associates it with the action sequence.

For each transaction agents generate corresponding action sequences. To represent deceitful actions, agents maintain private actions which they include in their internalised action sequences in addition to the overtly observable actions, as alluded to in Subsection 4.4. If an operator cheats, he internalises the action ‘cheating’ (in addition to ‘trading’) but merely indicate that he traded the goods. He can then monitor this action sequence and internalise if his cheating (in combination with trading) was successful. Another case is the suspected cheating (from the perspective of the original sender of the goods). If suspecting cheating by an observed agent, the original sender can add the private action ‘cheating’ to his internalised action sequence of the observed individual without sharing this suspicion.

Based on their experience agents derive normative statements as outlined in the previous sections. Traders can query the derived statements to guide their decisions whether to continue sending goods by aggregating deontic values for given actions (e.g. sending goods) across one’s overall experience. Given our characterisation of norms as continua, agents stop committing to actions such as ‘sending goods’ as soon as those actions appear undesirable from a normative perspective. Using the labels of the prototypical deontic range (see Subsection 3.2), this would be the case if such actions are associated with the deontic compartments *may not*, *should not*, or *must not*. However, instead of having agents rigidly act according to those compartments, we introduce an element of randomness. To achieve this all deontic values are normalised across the deontic range (i.e. lowest deontic value to 0 (*must not*), deontic centre as 0.5 (neutral), and highest deontic value as 1 (*must*)). Those normalised values are then used as input probabilities for the random number generator to determine whether the behaviour is performed (and economic transactions sustained). Defecting from performing such actions naturally limits market interactions if they are deemed normatively undesirable. In consequence, this affects the overall economic performance, with the best possible case being complete market participation, and the worst case being the breakdown of any market activity. The notion of normalised deontic values is likewise used to determine the fraction of embezzled profits – the embezzled fraction increases with the internalised normative desirability of cheating.

Algorithm 3 outlines the agents’ execution cycles described in the previous paragraphs; Algorithm 4 specifies agents’ reactions to possible types of incoming requests (which are marked as different ‘cases’).

Algorithm 3 Agent Execution Cycle

```

1: if < maxRelationships relationships to other traders then
2:   Pick random another trader this agent does not have relationship with and send relationship
   request;
3:   if request is accepted then
4:     Add accepting trader to set of relationships;
5:   end if
6: end if
7: if agent has relationships to other traders then
8:   Pick random fellow trader from set of relationships;
9:   Determine random Boolean to decide whether to send goods (using normalised deontic
   value for action ‘send goods’ as probability);
10:  if sending goods then
11:    Send goods to selected trader and await return of profit;
12:    if profit < 0 and mean value of memorised past transactions from trader is < 0 then
13:      Extend received action statement with action cheat;
14:      Memorise trader as cheater;
15:      Fire trader;
16:    else
17:      Reward trader by paying a reward of wageFactor * profit;
18:    end if
19:    Memorise action statement in association with profit;
20:  end if
21: end if
22: Derive nADICO statements from memory;

```

Algorithm 4 Agent Reactions

```

1: Initialisation: Initialise agent as cheater with probability  $p_{cheater}$ ;
2: Case 1: Receipt of Relationship Request
3: if requester is not memorised as cheater and  $< maxRelationships$  relationships then
4:   Accept request and add requester to relationships;
5: else
6:   Reject request;
7: end if
8: Case 2: Receipt of Trade Request
9: Perform market transaction (goods' value *  $profitFactor$ );
10: if initialised as cheater then
11:   Determine fraction  $f$  of profit to embezzle as normalised deontic value for action 'cheating'
      ( $0 < f \leq 1$ );
12:   Create private copy of action statement and add action cheat;
13:   Memorise extended action sequence;
14: end if
15: Return (remaining) profit to sender;
16: Case 3: Receipt of Firing Notification
17: Remove sender from own relationships;
18: Mark sender as cheater;

```

6. Evaluation. Turning to the evaluation, we develop a baseline scenario (Subsection 6.1) in order to understand how the generalisation mechanism can be operationalised. Following this, we modify the parameter set of the base model by introducing cheating (Subsection 6.2) and analyse how this affects the overall simulation outcome. As a third step, we change the composition of the artificial society to see the impact of changing social structures on the emerging normative understanding (Subsection 6.3), before exploring the generalisation of norms beyond the observed actor characteristics as a precursor for predictive capabilities (Subsection 6.4).

6.1. Establishing Normative Understanding in the Base Scenario. We initialise the simulation with the base parameter set outlined in Table 6.1, one of which we vary in the course of the evaluation. Since our intent is to measure economic cooperation, we use the number of performed transactions per round as a target variable to indicate overall economic performance. With the exception of the $wageFactor$ that is derived from documented profit distributions in commenda relationships [44], the chosen parameter set has not been empirically explored. Instead it is initialised to produce return for both sender and operator in the case of cooperative interaction to show how such outcomes are reflected in the generated normative understanding, without assuming an accurate representation of historical happenstance.

Table 6.1: Simulation Parameters

Parameter	Default Value	Variation in Section
Number of Agents	50	
$maxRelationships$	20	
$profitFactor$ (Factor of Goods' Value)	1.3	
$wageFactor$ (Factor of Total Profit)	0.33	
$p_{cheater}$ (Fraction of Cheaters)	0	Subsection 6.2
Deontic Aggregation Strategy	Rational	
Number of Profit Memory Entries	100	
Number of Cheater Memory Entries	10	

Running a base scenario without any cheating we can establish the maximum level of cooperation, that is 100 transactions per simulation round. In this scenario each agent is involved twice per round, once for sending goods to initiate trade by another agent, and once for returning goods to the original sender. For all simulation runs we evaluate the emerging normative understanding after 1000 rounds, a period after which behaviour stabilised in all simulation scenarios.

In the base scenario the absence of cheaters offers limited surprise. However, it is a good starting point to identify how our generalisation mechanism translates individual agents' observations into normative understanding at large. Inspecting an individual's traces in Figure 6.1, we can identify the situational deontic range boundaries as well as the generated nADICO statements.

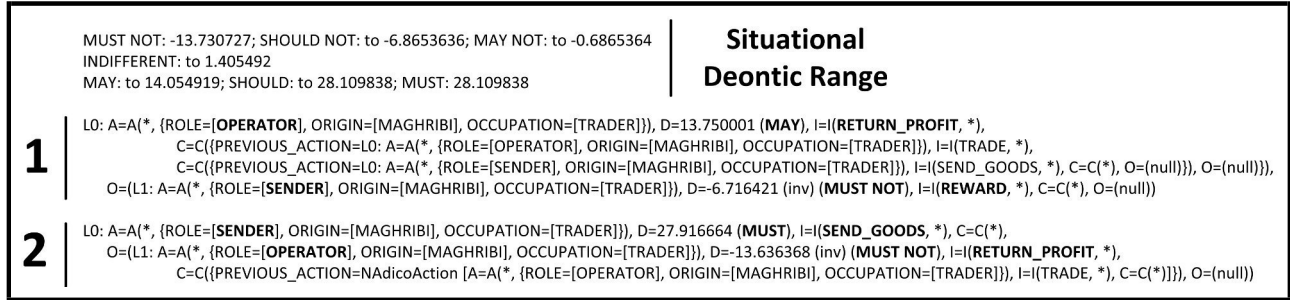


Fig. 6.1: Situational nADICO Statements for Individual Agent in Baseline Scenario

For this particular agent, the deontic range spreads from around -13.73 to +28.11, with the deontic centre lying at 0. The subranges from centre to proscriptive end, and from centre to prescriptive end are divided into equally-sized deontic compartments within the respective subranges.

Based on its experience this particular agent has inferred two nADICO statements, both of which are identified by indices. The essential features, such as actor attributes, deontic, and actions (aim) are highlighted to ease the interpretation. The statements consist of monitored statements on nesting level 0 (L0) that includes all preceding actions, as well as consequential statements, nested on Level 1 (L1), that attach an experientially derived presumed sanction for non-compliance. Translating the first statement into natural language, we can read 'Operators may want to return profit (Line 1 of Statement 1), after having been sent goods by senders (Line 3) and traded those (Line 2), or else senders must not reward operators (Line 4).' As noted before, the preceding actions leading to the monitored action are stored in the *Conditions* component of a nADICO statement on a given nesting level and may carry preceding actions on their own, as found in this example.

The second statement concerns the action 'sending goods' and reads: 'Senders must send goods (Line 1 of Statement 2), or else operators must not return profit (Line 2) after having traded those (Line 3).'

Both statements manage to capture the complexity of the interactions and role associations, while being accessible for human interpretation. However, even though both statements describe inferred obligations their deontic strength varies, with strong inclinations for sending goods (Statement 2) and initiating trade interactions themselves, while the return of profit (Statement 1) is of lesser priority. This is associated with the varying levels of reinforcement both actions experience based on the internalised feedback.

This aspect is better reflected when observing the normative understanding on the macro-level. For this purpose we rely on a Kiviat-inspired chart that shows distributions across adjacent ordinally scaled values, such as deontic terms associated with the individual deontic compartments. Figure 6.2 shows the distribution of normative understanding for identified nADICO statements, with focus on the monitored statements' *Attributes* component (in our case 'Operator' or 'Sender' and further social markers), followed by the action sequence (to be read in reverse order similar to the individual statements above, with the last action being the guarded action in the monitored statement; the consequential statement's action is omitted).

Observing the aggregated overview, we can see that the vast majority of agents (82%) feels obliged to send goods, with a minor fraction of 18% perceiving it as omissible (generally obligatory, but exceptionally foregone). Returning profit finds more moderate reinforcement, with 70% only feeling mildly inclined to do so, mostly because of the lesser feedback compared to the aggregated feedback from sending goods (which includes the incurred profit). However, in all cases the normative understanding suggests economic cooperation. This example highlights the relationship between reinforcement and the salience of a given norm for our generalisation mechanism, thus offering a possible approach to model the prioritisation of conflicting or alternative norms.

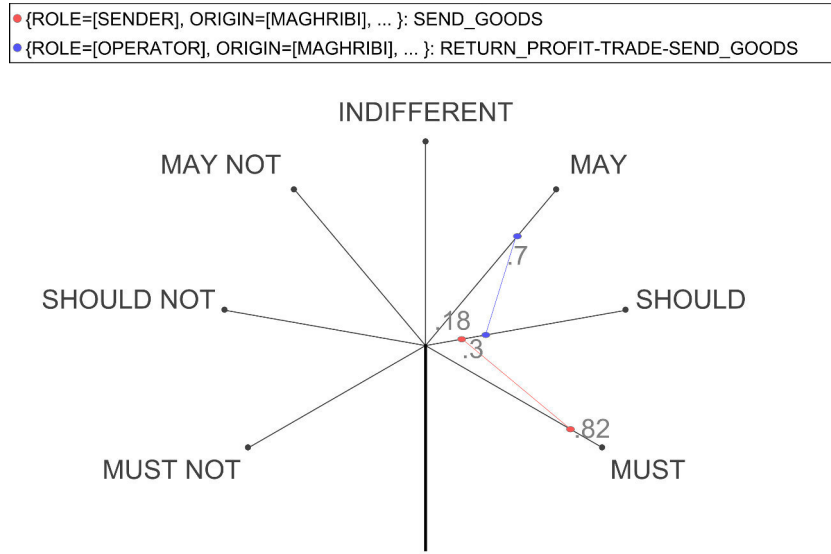


Fig. 6.2: Distribution of Normative Understanding for Baseline Scenario

With this baseline in mind, we can introduce different levels of cheating into the social composition and observe how this is reflected in the level of cooperation, as well as individuals’ normative understandings.

6.2. Modifying Behaviour: Introducing Cheating. To analyse the effect of cheating on the artificial society, we incrementally increase $p_{cheater}$, i.e. the fraction of all agents that involve in cheating.

Introducing a fraction of 25% as cheaters, the society can still sustain nearly full cooperation (as shown in Table 6.2). This implies that the modelled society can sustain a moderate amount of cheating. However, to gain insight into what exactly facilitates this, we can inspect the normative understanding shown in Figure 6.3, which, in addition to the previous figure, includes an additional third action sequence, reflecting experience of cheating (to be read as ‘returning profit after having been sent goods, traded, as well as cheated’).

Observing the traces in the radar chart, we can see that the majority of agents has a stronger inclination to trade fair (the blue trace) than to cheat. This dominant commitment to honour the operator’s obligation to return goods over cheating ensures the continued sending of goods (red trace), even though to a slightly lesser extent than in the baseline case (see Figure 6.2 in Subsection 6.1).

Only for higher levels of cheating we notice an incremental breakdown of market operations, with mean transactions counts for selected cheater fractions shown in Table 6.2. All statistical values are based on 30 simulation runs per configuration.

Table 6.2: Transaction Counts for Different Cheater Fractions

Cheater Fraction	Mean	Std. dev.
0	98.873	5.164
0.25	99.805	0.565
0.5	84.378	5.338
0.75	35.428	5.544
1	11.442	4.485

More interesting than the breakdown of market cooperation are the means by which we can retrace this

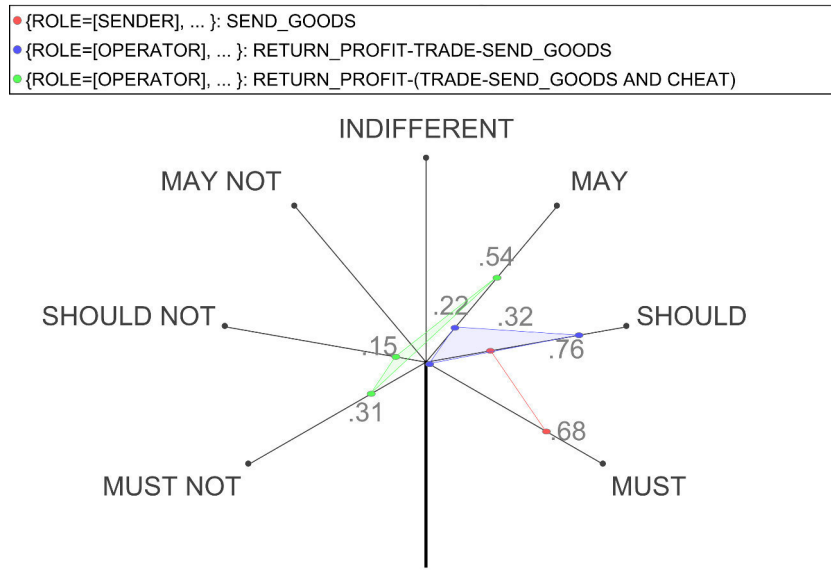


Fig. 6.3: Distribution of Normative Understanding for a Cheater Fraction of 25%

shift in normative understanding. While cheating levels up to 0.5 only affect the prioritisation of actions on the prescriptive part of the deontic range (see Figure 6.4(a)), higher levels afford a shift of the action ‘send goods’ into the proscriptive range, thus promoting the undesirability of that action (red trace) as shown for a cheater fraction of 75% in Figure 6.4(b). This shift is based on the declining revenue for initiating trades in the light of increasing levels of cheating (note the desirability of cheating (green trace) over honest trading (blue trace)) and thus the incremental breakdown of cooperation.

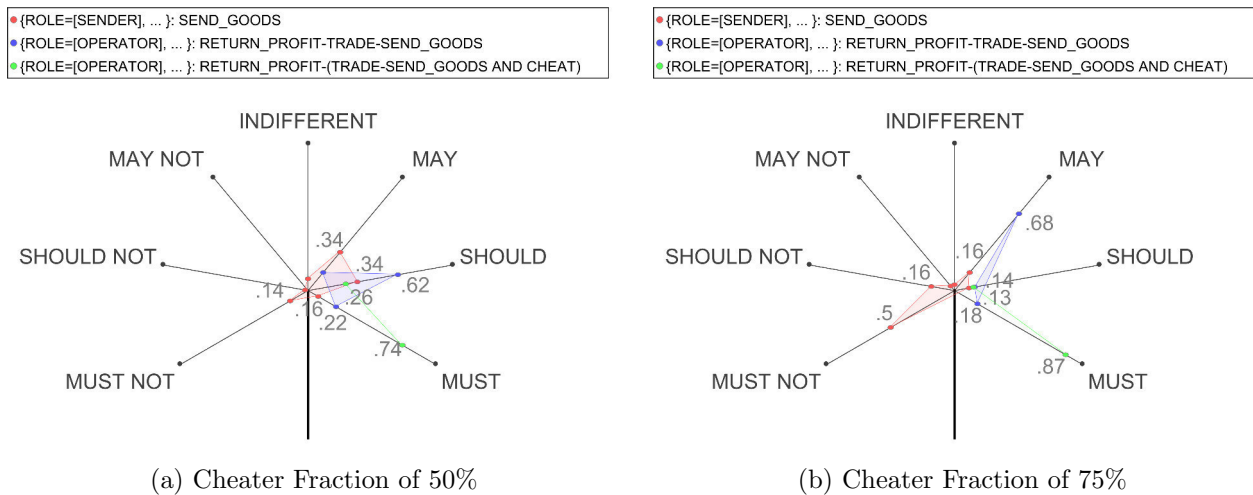


Fig. 6.4: Distribution of Normative Understanding for Different Cheater Fractions

To get a better understanding of an individual agent’s motivation to adjust its behaviour, we can observe the generated statements for an individual trader in a scenario with a cheater fraction of 75% as shown in Figure 6.5.

In this example, the agent has generated three nADICO statements, with the first one highlighting the *desirability of cheating*, the second one indicating the *undesirability of sending goods*, and the third one indi-

```

LO: A=A(*, {ROLE=[OPERATOR], ORIGIN=[MAGHRIBI], OCCUPATION=[TRADER]}), D=30.12369 (MUST), I=I(RETURN_PROFIT, *), C=C({PREVIOUS_ACTION=
  (LO: A=A(*, {ROLE=[OPERATOR], ORIGIN=[MAGHRIBI], OCCUPATION=[TRADER]}), I=I(TRADE, *), C=C({PREVIOUS_ACTION=
    LO: A=A(*, {ROLE=[SENDER], ORIGIN=[MAGHRIBI], OCCUPATION=[TRADER]}), I=I(SEND_GOODS, *), C=C(*), O=(null))), O=(null)) AND
    (LO: A=A(*, {ROLE=[OPERATOR], ORIGIN=[MAGHRIBI], OCCUPATION=[TRADER]}), I=I(CHEAT, *), C=C(*), O=(null))),
  O=(L1: A=A(*, {ROLE=[SENDER], ORIGIN=[MAGHRIBI], OCCUPATION=[TRADER]}), D=-0.4248632 (inv) (MUST NOT), I=I(FIRE, *), C=C(*), O=(null))

LO: A=A(*, {ROLE=[SENDER], ORIGIN=[MAGHRIBI], OCCUPATION=[TRADER]}), D=-0.64500105 (MUST NOT), I=I(SEND_GOODS, *), C=C(*),
O=(L1: A=A(*, {ROLE=[OPERATOR], ORIGIN=[MAGHRIBI], OCCUPATION=[TRADER]}), D=45.731926 (inv) (MUST), I=I(RETURN_PROFIT, *),
  C=C({PREVIOUS_ACTION=NADicoAction [A=A(*, {ROLE=[OPERATOR], ORIGIN=[MAGHRIBI], OCCUPATION=[TRADER]}), I=I(TRADE, *), C=C(*)]}), O=(null))

LO: A=A(*, {ROLE=[OPERATOR], ORIGIN=[MAGHRIBI], OCCUPATION=[TRADER]}), D=16.916664 (SHOULD), I=I(RETURN_PROFIT, *), C=C({PREVIOUS_ACTION=
  LO: A=A(*, {ROLE=[OPERATOR], ORIGIN=[MAGHRIBI], OCCUPATION=[TRADER]}), I=I(TRADE, *), C=C({PREVIOUS_ACTION=
    LO: A=A(*, {ROLE=[SENDER], ORIGIN=[MAGHRIBI], OCCUPATION=[TRADER]}), I=I(SEND_GOODS, *), C=C(*), O=(null))), O=(null))),
  O=(L1: A=A(*, {ROLE=[SENDER], ORIGIN=[MAGHRIBI], OCCUPATION=[TRADER]}), D=-0.23859188 (inv) (SHOULD NOT), I=I(REWARD, *), C=C(*), O=(null))

```

Fig. 6.5: Situational nADICO Statements for Individual Agent in Scenario with a Cheater Level of 75%

cating a *moderate inclination to trade fair*. We highlighted the essential *Attributes* components and actions for monitored and consequential statements. To emphasise the composite nature of the first statement (specifically in contrast to the third one), we additionally emphasised the horizontal nesting of the composite action ‘trade and cheat’ that precedes the returning of goods.

While the second and third statement bear limited novel insights, the first one clarifies the motivation as to why cheating is desirable. Interpreted literally, the agent presumes that ‘operators must trade, cheat and return the profit, or otherwise they are not fired’. This essentially reflects the inferred expectation of being fired as a result of cheating. However, even though the agent identifies this norm, it does not have any prior conception of what ‘cheating’ or ‘returning profits’ entails, but infers this based on long-term reinforcement. In this case the accumulated net pay-off of cheating and potentially being fired afterwards exceeds the pay-off received from trading fair (third statement), which makes cheating a profitable option. This interpretation shows that agents are able to identify the normative behaviour, while their actions choices remain utilitarian.

This collective shift in normative understanding with increasing penetration of cheating highlights the conditions under which cooperation can be sustained in this agent society, and beyond the quantitative evaluation of the dependent variable ‘transaction count’, gives us insight into the normative understanding on micro- and macro-level.

6.3. Modifying Structure: Introducing Outsiders. To explore the impact of introducing a wider set of different social actors (as opposed to merely changing behavioural composition as done up to this stage), we introduce a secondary type of trader into the social pool. While the Maghribis’ economic success built on the successful use of informal institutional instruments, European societies predominantly relied on formal institutional instruments such as contracts [19]. They did so in the light of changing social composition of their societies (especially in City States such as Genoa), lacking long-term commitment of individuals, and thus an increased incentive to cheat – since only long-term employment would have provided the necessary incentive to act compliantly.¹³

We thus conceive a second prototypical group of traders that is more inclined to cheat¹⁴, which we will refer to as ‘Genoese’. Note that this composition is synthetic. Historically, Maghribis and Genoese hardly engaged in trade interactions with each other [19].

With an initial fraction of 25% as Genoese (and otherwise non-cheating Maghribis), after 1000 rounds we arrive at a transaction level of around 49, i.e. effectively half of the possible transactions. The incremental breakdown for changing social compositions (in fraction levels from 0.0 to 1.0 in steps of 0.25 across 30 simulation runs each) is shown in Table 6.3.

To maintain an accessible representation of the normative understanding, we separate the aggregated statements for Maghribis and Genoese, with the Maghribi representation shown in Figure 6.6(a) and Genoese one in Figure 6.6(b).

¹³For an in-depth elaboration refer to [19] and [16].

¹⁴For Genoese we induce higher cheating levels by increasing the fraction of embezzled profits by 20%.

Table 6.3: Transaction Counts for Different Fractions of Genoese

Genoese Fraction	Mean	Std. dev.
0	98.873	5.164
0.25	48.997	9.48
0.5	35.52	6.966
0.75	32.003	6.634
1	8.11	3.501

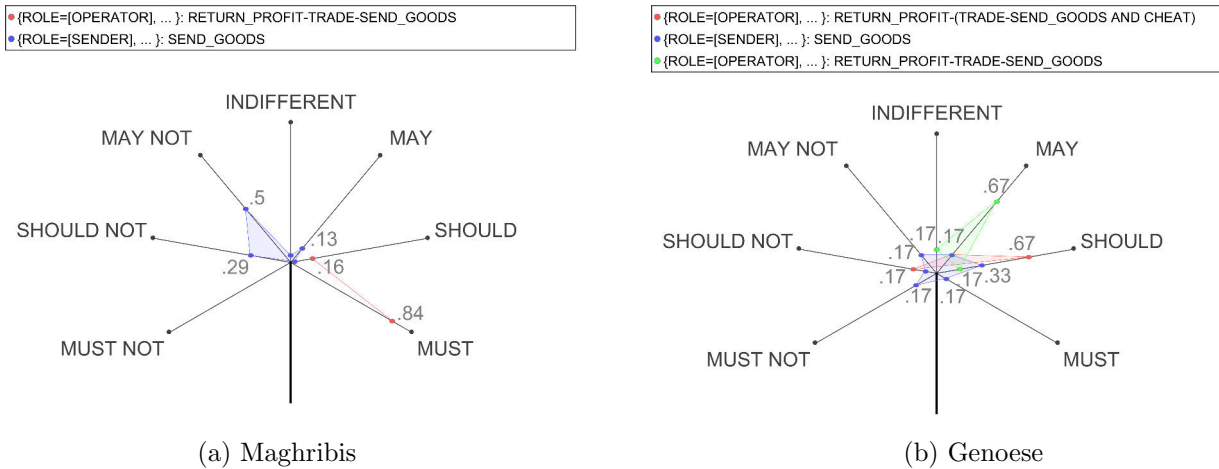


Fig. 6.6: Distribution of Normative Understanding for Maghribis and Genoese

Observing the non-cheating Maghribis, we can see how committing to trade requests and returning the generated profits is considered an obligation, since it yields in promising returns. Due to the increasing cheating levels, sending of goods becomes increasingly undesirable as of the considerable risk of being cheated by the introduced Genoese.

This complements with the impression of the Genoese traders. Genoese initiate the bulk of trade transactions, but have a clear preference for cheating, with 67% treating it as a *should*, with the same 67% of interpreting only mild inclinations for fair trading.

In essence, integrating two different trader society’s operating based on different institutional principles shows interesting interaction patterns, which we can derive from the emerging normative understanding. The Maghribis concentrate their honesty on a role as operators, i.e. pursuing Genoese’ trade dealings, while this honest trading behaviour sustains the Genoese’ commitment to inject a moderate level of goods into the market, thus preventing the economic collapse. Maintaining this base level of fair trading by Genoese senders and Maghribi operators, Genoese operators can retain their preference for cheating.

These interdependencies become clearer when inspecting individuals’ normative understandings – this time from a Genoese perspective shown in Figure 6.7.

The produced set of statements is double in size than for the uni-cultural case. Recall that Level 0 generalisation (as introduced in Subsection 4.2) operates on all social markers (nADICO’s *Attributes*), that now not only differentiates senders and operators as roles, but further includes the marker ‘origin’ (Maghribi vs. Genoese), for which normative behaviour can be characterised.

Specifically statements 2 and 3 are of particular interest. Both describe the agent’s normative understanding of cheating in the context of trading goods and returning profits, but they do so for different attribute sets. Statement 2 highlights that a Genoese operator *should* return profit to a Genoese sender when embezzling profits, whereas when dealing with a Maghribi sender, cheating is strongly undesirable. For this case, the

```

LO: A=A(*, {ROLE=[SENDER], ORIGIN=[GENOESE], OCCUPATION=[TRADER]}), D=96.73126 (MUST), I=I(SEND_GOODS, *), C=C(*),
O=(L1: A=A(*, {ROLE=[OPERATOR], ORIGIN=[MAGHRIBI], OCCUPATION=[TRADER]}), D=-85.12692 (inv) (MUST NOT), I=I(RETURN_PROFIT, *),
C=C({PREVIOUS_ACTION=NAdicoAction [A=A(*, {ROLE=[OPERATOR], ORIGIN=[MAGHRIBI], OCCUPATION=[TRADER]}), I=I(TRADE, *), C=C(*))}), O=(null))

LO: A=A(*, {ROLE=[OPERATOR], ORIGIN=[GENOESE], OCCUPATION=[TRADER]}), D=58.2375 (SHOULD), I=I(RETURN_PROFIT, *), C=C({PREVIOUS_ACTION=
(L0: A=A(*, {ROLE=[OPERATOR], ORIGIN=[GENOESE], OCCUPATION=[TRADER]}), I=I(TRADE, *), C=C({PREVIOUS_ACTION=L0: A=A(*, {ROLE=[SENDER],
ORIGIN=[GENOESE], OCCUPATION=[TRADER]}), I=I(SEND_GOODS, *), C=C(*), O=(null))), O=(null)) AND
(L0: A=A(*, {ROLE=[OPERATOR], ORIGIN=[GENOESE], OCCUPATION=[TRADER]}), I=I(CHEAT, *), C=C(*), O=(null))),
O=(L1: A=A(*, {ROLE=[SENDER], ORIGIN=[GENOESE], OCCUPATION=[TRADER]}), D=-51.251053 (inv) (SHOULD NOT), I=I(FIRE, *), C=C(*), O=(null))

LO: A=A(*, {ROLE=[OPERATOR], ORIGIN=[GENOESE], OCCUPATION=[TRADER]}), D=-72.62001 (MUST NOT), I=I(RETURN_PROFIT, *), C=C({PREVIOUS_ACTION=
(L0: A=A(*, {ROLE=[OPERATOR], ORIGIN=[GENOESE], OCCUPATION=[TRADER]}), I=I(TRADE, *), C=C({PREVIOUS_ACTION=L0: A=A(*, {ROLE=[SENDER],
ORIGIN=[MAGHRIBI], OCCUPATION=[TRADER]}), I=I(SEND_GOODS, *), C=C(*), O=(null))), O=(null)) AND
(L0: A=A(*, {ROLE=[OPERATOR], ORIGIN=[GENOESE], OCCUPATION=[TRADER]}), I=I(CHEAT, *), C=C(*), O=(null))),
O=(L1: A=A(*, {ROLE=[SENDER], ORIGIN=[MAGHRIBI], OCCUPATION=[TRADER]}), D=-82.51943 (inv) (MUST), I=I(FIRE, *), C=C(*), O=(null))

LO: A=A(*, {ROLE=[SENDER], ORIGIN=[GENOESE], OCCUPATION=[TRADER]}), D=-27.59 (MAY NOT), I=I(SEND_GOODS, *), C=C(*),
O=(L1: A=A(*, {ROLE=[OPERATOR], ORIGIN=[GENOESE], OCCUPATION=[TRADER]}), D=31.351017 (inv) (MAY), I=I(RETURN_PROFIT, *),
C=C({PREVIOUS_ACTION=NAdicoAction [A=A(*, {ROLE=[OPERATOR], ORIGIN=[GENOESE], OCCUPATION=[TRADER]}), I=I(TRADE, *), C=C(*))}), O=(null))

LO: A=A(*, {ROLE=[OPERATOR], ORIGIN=[GENOESE], OCCUPATION=[TRADER]}), D=4.9500003 (MAY), I=I(RETURN_PROFIT, *), C=C({PREVIOUS_ACTION=L0: A=A(*,
{ROLE=[OPERATOR], ORIGIN=[GENOESE], OCCUPATION=[TRADER]}), I=I(TRADE, *), C=C({PREVIOUS_ACTION=L0: A=A(*, {ROLE=[SENDER],
ORIGIN=[MAGHRIBI], OCCUPATION=[TRADER]}), I=I(SEND_GOODS, *), C=C(*), O=(null))), O=(null))),
O=(L1: A=A(*, {ROLE=[SENDER], ORIGIN=[MAGHRIBI], OCCUPATION=[TRADER]}), D=-4.356175 (inv) (MAY NOT), I=I(REWARD, *), C=C(*), O=(null))

LO: A=A(*, {ROLE=[OPERATOR], ORIGIN=[GENOESE], OCCUPATION=[TRADER]}), D=4.9500003 (MAY), I=I(RETURN_PROFIT, *), C=C({PREVIOUS_ACTION=L0: A=A(*,
{ROLE=[OPERATOR], ORIGIN=[GENOESE], OCCUPATION=[TRADER]}), I=I(TRADE, *), C=C({PREVIOUS_ACTION=L0: A=A(*, {ROLE=[SENDER],
ORIGIN=[GENOESE], OCCUPATION=[TRADER]}), I=I(SEND_GOODS, *), C=C(*), O=(null))), O=(null))),
O=(L1: A=A(*, {ROLE=[SENDER], ORIGIN=[GENOESE], OCCUPATION=[TRADER]}), D=-4.356175 (inv) (MAY NOT), I=I(REWARD, *), C=C(*), O=(null))

```

Fig. 6.7: Situational nADICO Statements for Individual Genoese Agent in Scenario with 25% Genoese

Genoese operator favours fair trading instead (Statement 5).

To understand how this differentiation comes about, we recall the principles of evaluating cheating: agents evaluate the situational profit for negative values, but also keep track of historical cheating both on individual level (i.e. past returned profits of individual agents) and global level (i.e. average across all past trade interactions), which makes the evaluation relative to the context agents interact in. For example, interacting within a pool of cheaters the embezzling of profit is less likely to be identified, in contrast to a predominantly compliant environment. As mentioned above, Maghribis stop initiating trade interactions (by sending goods) over the course of the simulation as of declining pay-offs (since cheaters eventually learn that cheating is profitable). In consequence, since they only act as operators they are hardly subject to cheating. However, if cheating occurs they operate based on past observations (from their earlier experience of fair trading), which leads to improved identification of cheating, resulting in negative feedback for the cheater. The Genoese, in contrast, continue to operate as senders in this mixed society and ‘got used’ to a moderate amount of cheating, which makes it less likely to identify and fire cheating individuals. Strategically, for Genoese it makes more sense to cheat their own kind, than embezzling profits from non-cheating Maghribis.

6.4. Generalising Identified Behaviour. This introduced complexity provides grounds for more differentiated social behaviour, with agents learning about actions as well as their social environment. However, as seen in the previous example, this complexity reflects all known categories of social structures, but does not give agents the ability to predict or stereotype behaviour on a more general level. Currently, agents can predict behaviour for ‘Genoese traders’ or ‘Maghribi traders’, but fail to characterise behaviour for traders more generally.

In Subsection 4.2 we proposed a conceptual solution by affording multi-level generalisation that systematically aggregates normative understanding by subsets of social marker combinations. In Figure 6.8 we can observe an extract (2 out of 6 statements) of the same individual’s understanding (as from Figure 6.7) on the highest generalisation level, in our case Level 2 – the generalisation of two attributes (out of three ones in our case: role, origin, occupation).

The traces concentrate on generalisations from a Genoese perspective, and result in the aggregation of statements based on the marker ‘Trader’ or ‘Genoese’. Generalisation based on the category ‘Genoese’ produces the

```

L0: A=A(*, {OCCUPATION=[TRADER]}), D=27.663889 (MAY), I=I(TRADE, *), C=C({PREVIOUS_ACTION=L0: A=A(*,
    {OCCUPATION=[TRADER]}), I=I(SEND_GOODS, *), C=C(*), O=(null)}),
O=(L1: A=A(*, {OCCUPATION=[TRADER]}), D=-24.345198 (inv) (MAY NOT), I=I(RETURN_PROFIT, *), C=C(*), O=(null))

L0: A=A(*, {ORIGIN=[GENOESE]}), D=-16.859093 (MAY NOT), I=I(TRADE, *), C=C({PREVIOUS_ACTION=L0: A=A(*,
    {ORIGIN=[GENOESE]}), I=I(SEND_GOODS, *), C=C(*), O=(null)}),
O=(L1: A=A(*, {ORIGIN=[GENOESE]}), D=19.157293 (inv) (MAY), I=I(RETURN_PROFIT, *), C=C(*), O=(null))

L0: A=A(*, {OCCUPATION=[TRADER]}), D=-14.461121 (MAY NOT), I=I(RETURN_PROFIT, *), C=C({PREVIOUS_ACTION=
    (L0: A=A(*, {OCCUPATION=[TRADER]}), I=I(TRADE, *), C=C({PREVIOUS_ACTION=L0: A=A(*,
        {OCCUPATION=[TRADER]}), I=I(SEND_GOODS, *), C=C(*), O=(null)}), O=(null)) AND
    (L0: A=A(*, {OCCUPATION=[TRADER]}), I=I(CHEAT, *), C=C(*), O=(null))}),
O=(L1: A=A(*, {OCCUPATION=[TRADER]}), D=16.432432 (inv) (MAY), I=I(FIRE, *), C=C(*), O=(null))

L0: A=A(*, {ORIGIN=[GENOESE]}), D=58.2375 (SHOULD), I=I(RETURN_PROFIT, *), C=C({PREVIOUS_ACTION=
    (L0: A=A(*, {ORIGIN=[GENOESE]}), I=I(TRADE, *), C=C({PREVIOUS_ACTION=L0: A=A(*, {ORIGIN=[GENOESE]}),
        I=I(SEND_GOODS, *), C=C(*), O=(null)}), O=(null)) AND
    (L0: A=A(*, {ORIGIN=[GENOESE]}), I=I(CHEAT, *), C=C(*), O=(null))}),
O=(L1: A=A(*, {ORIGIN=[GENOESE]}), D=-51.251053 (inv) (SHOULD NOT), I=I(FIRE, *), C=C(*), O=(null))

L0: A=A(*, {OCCUPATION=[TRADER]}), D=4.9500003 (MAY), I=I(RETURN_PROFIT, *), C=C({PREVIOUS_ACTION=L0: A=A(*,
    {OCCUPATION=[TRADER]}), I=I(TRADE, *), C=C({PREVIOUS_ACTION=L0: A=A(*, {OCCUPATION=[TRADER]}),
        I=I(SEND_GOODS, *), C=C(*), O=(null)}), O=(null)}),
O=(L1: A=A(*, {OCCUPATION=[TRADER]}), D=-4.356175 (inv) (MAY NOT), I=I(REWARD, *), C=C(*), O=(null))

L0: A=A(*, {ORIGIN=[GENOESE]}), D=4.9500003 (MAY), I=I(RETURN_PROFIT, *), C=C({PREVIOUS_ACTION=L0: A=A(*,
    {ORIGIN=[GENOESE]}), I=I(TRADE, *), C=C({PREVIOUS_ACTION=L0: A=A(*, {ORIGIN=[GENOESE]}),
        I=I(SEND_GOODS, *), C=C(*), O=(null)}), O=(null)}),
O=(L1: A=A(*, {ORIGIN=[GENOESE]}), D=-4.356175 (inv) (MAY NOT), I=I(REWARD, *), C=C(*), O=(null))

```

Fig. 6.8: Extract of Situational nADICO Statements on Generalisation Level 2 for Individual Genoese Agent in Scenario with 25% Genoese

same results as for the Level 0 generalisation (see Figure 6.7), whereas an aggregation by ‘Trader’ integrates the previously separated statements for ‘Maghribi’ and ‘Genoese’. Focusing on the cheating behaviour (Statements 1 and 2 in Figure 6.8), we find that Statement 2 corresponds to Statement 2 in Figure 6.7, while Statement 1 in Figure 6.8 is an aggregation of Statements 2 and 3 in Figure 6.7. For this example, we can now develop a more abstract picture of normative behaviour: while cheating is a desirable strategy for dealings amongst Genoese, for traders in general (which includes trades with Genoese and Maghribis) the likeliness of being fired leads to a mild disinclination (*may not*) from cheating. In principle, a Genoese trader could now ascribe this behavioural stereotype to newly entering non-Maghribi and non-Genoese traders (e.g. Venetian traders), and avoid cheating as of the anticipated negative consequence.

With this derived abstract understanding, we can reflect stereotyping capabilities that extend beyond the ‘known’ of the individual. This approach thus facilitates the *reduction of uncertainty* when dealing with newcomers, as well as providing a shortcut for decision-making, which offers the potential for an *efficiency increase* – both of which are functions characteristic for institutions [32, 50].

7. Summary, Discussion & Outlook. In this work we extend a comprehensive approach to model norm emergence and identification in artificial agent societies previously introduced in [15]. It builds on an expressive institutional grammar that captures different types of institutions, including conventions, norms, and rules (Subsection 3.1). The essential benefits of this structure are a) its generic and unified nature (one representation for different institution types), while b) providing what we believe is an acceptable compromise of accessible interpretation by human observers, and c) capturing institutions with great detail, facilitated by

notions of nested institutional statements.

This structural representation is augmented with a continuous conception of normative duty, capturing prohibitions, permissions, and obligations (Subsection 3.2). This not only facilitates a nuanced expression of norms and a better reflection of their evolution over time, but provides the foundation for prioritising conflicting norms along a continuous and dynamically changing deontic range. This representation is individualised, thus reflecting individuality on the micro-level, while maintaining generalisability for macro-level inspection of the normative landscape. This analysis of the normative landscape can be performed on arbitrary social aggregation levels, i.e. for all agents (macro-level), specific sub-groups (meso-level), such as cheaters or non-cheaters, as well as for an individual-centric micro-perspective.

Apart from a refinement of the aforementioned structures and their specification, the essential contribution of this work lies in the systematic formalisation of the norm identification process and its extension with a multi-level generalisation of behaviour that operates on an individual's experience of its social environment. The mechanism attempts to reconstruct subconscious norm internalisation processes by relying on concepts such as (experiential or social) learning and multi-level generalisation (stereotyping).

The integrated use of both 'structure' (nADICO & Dynamic Deontics) and 'process' (Generalisation) has been demonstrated in a multi-agent trader scenario in which agents generate experience-based normative understanding that feeds back into their action decisions. As part of the evaluation we explored the impact of cheating behaviour and structural diversity on the emerging normative understanding.

Our evaluation scenario is inspired by a realistic medieval scenario from the area of comparative economics, in which we explored the effect of cheating in societies that operate on different institutional principles, with one society being the Maghribi collective that relied on informal coordination mechanisms, and a second one, the Genoese, that inherently relied on formal coordination mechanisms to foster long-distance trade interactions. Even though the results themselves bear limited surprise with respect to the introduction of cheating – eventually economic cooperation collapses – the emerging normative understanding does. Using the proposed generalisation mechanisms we can shed light on the inner workings of individuals as well as the motivations that underlie their action selection, such as understanding why specific levels of cheating can be tolerated while sustaining economic cooperation. Another interesting finding relates to the hypothetical mixing of both societies, with Genoese goods senders developing a normative preference to interact with compliant Maghribis, while preferring to cheat their own kind when acting as operators. These conceptions cannot be validated, since no interaction between both societies has been documented [19]. However, it sheds an interesting insight into what could have happened if both collectives had interacted, and perhaps even sketching a possible reason why they did not (or ceased to do so).¹⁵

Our approach aims to provide a pragmatic account of realistic norm identification processes as evident in human beings as well as non-human species. This is in contrast to human-inspired norm engineering approaches that rely on additional assumptions, such as central coordinating entities to guarantee the identification of efficient norms (as is the case with Morales et al. [28, 29, 30]), or explicit collective action processes to facilitate consensus (as used by Riveret et al. [37]). In contrast to their work, our generalisation of social structures does not rely on prior knowledge, such as ontologies or the like, but is inferred as part of the norm identification process – the individuals themselves derive structure from behaviour. As with human societies, our conception of social norms is inherently virtual. No central entity maintains or governs a social norm. Individuals share norms via continuous internalisation and socialisation processes, but the normative understanding only exists on the micro-level. A system-wide identification of norms only occurs for evaluative purposes, a feature that is facilitated by the expressive but uniform institution representations.

The proposed process bears strong generalisability characteristics, since it operates on the subconscious level and makes minimal assumptions about agents' cognitive capabilities; it merely requires a representation of individual and social markers, as well as actions and associated feedback specifications – equivalents of which are found in most social simulation scenarios.

We further believe that this approach bears strong potential for customisation beyond scenario-specific marker and action specifications, including different configurations of the deontic range (see Subsection 3.2), or

¹⁵Recall that the Maghribis stopped sending goods in the course of the simulation because of the Genoese cheating behaviour, but continued to facilitate trading for others.

the personality- and/or culture-dependent choice of feedback aggregation strategies (see Subsection 4.3).

Similarly, the means by which agents can attain normative understanding are not constrained to reinforcement learning, but could involve notions of imitation, social learning, or even explicit communication. The latter further opens the prospect of injecting institutional (i.e. normative and/or legal) prescriptions, thus complementing the bottom-up perspective as emphasised in this work with a top-down perspective. The generalisation mechanism itself can furthermore be combined with other components to form a more realistic cognitive agent architecture, or be ‘plugged into’ existing architectures and complement explicit reasoning capabilities (such as the case with BDI-style architectures) with a subconscious norm generalisation process.

We believe that the proposed approach is an important contribution towards a more realistic representation of institutions in general, and norms in particular, moving agents an incremental step closer towards a comprehensive representation of human behaviour in silico.

REFERENCES

- [1] D. Acemoglu and J. Robinson. *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*. Crown Business, New York (NY), 2012.
- [2] G. Andrighetto, D. Villatoro, and R. Conte. Norm Internalization in Artificial Societies. *AI Communications*, 23(4):325–339, Dec. 2010.
- [3] M. Aoki. *Toward a Comparative Institutional Analysis*. MIT Press, Cambridge (MA), 2001.
- [4] A. Bandura, D. Ross, and S. A. Ross. Transmission of Aggression through Imitation of Aggressive Models. *Journal of Abnormal and Social Psychology*, 63(3):575–582, 1961.
- [5] R. F. Baumeister, E. Bratslavsky, and C. Finkenauer. Bad is Stronger than Good. *Review of General Psychology*, 5(4):323–370, 2001.
- [6] C. Bicchieri. *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge University Press, New York (NY), 2006.
- [7] G. Björnsson and R. Shanklin. Must, Ought and the Structure of Standards. In F. Cariani, D. Grossi, J. Meheus, and X. Parent, editors, *Deontic Logic and Normative Systems*, volume 8554 of *Lecture Notes in Computer Science*, pages 33–48. Springer International Publishing, 2014.
- [8] S. Bradner. Key Words for Use in RFCs to Indicate Requirement Levels. <https://www.ietf.org/rfc/rfc2119.txt>, 1997. Accessed on: 1st December 2015.
- [9] C. Castelfranchi. Minds as Social Institutions. *Phenomenology and the Cognitive Sciences*, 13(1):121–143, 2014.
- [10] R. B. Cialdini, R. R. Reno, and C. A. Kallgren. A Focus Theory of Normative Conduct: Recycling the Concept of Norms to Reduce Littering in Public Places. *Journal of Personality and Social Psychology*, 58(6):1015–1026, 1990.
- [11] S. E. Crawford and E. Ostrom. A Grammar of Institutions. In *Understanding Institutional Diversity*, chapter 5, pages 137–174. Princeton University Press, Princeton (NJ), 2005.
- [12] J. Decety and J. Grèzes. The Power of Simulation: Imagining One’s Own and Other’s Behavior. *Brain Research*, 1079(1):4–14, 2006.
- [13] C. Frantz, M. K. Purvis, M. Nowostawski, and B. T. R. Savarimuthu. nADICO: A Nested Grammar of Institutions. In G. Boella, E. Elkind, B. T. R. Savarimuthu, F. Dignum, and M. K. Purvis, editors, *PRIMA 2013: Principles and Practice of Multi-Agent Systems*, volume 8291 of *Lecture Notes in Artificial Intelligence*, pages 429–436, Berlin, 2013. Springer.
- [14] C. Frantz, M. K. Purvis, M. Nowostawski, and B. T. R. Savarimuthu. Modelling Institutions using Dynamic Deontics. In T. Balke, F. Dignum, M. B. van Riemsdijk, and A. K. Chopra, editors, *Coordination, Organizations, Institutions and Norms in Agent Systems IX*, volume 8386 of *Lecture Notes in Artificial Intelligence*, pages 211–233, Berlin, 2014. Springer.
- [15] C. Frantz, M. K. Purvis, B. T. R. Savarimuthu, and M. Nowostawski. Modelling Dynamic Normative Understanding in Agent Societies. In H. K. Dam, J. Pitt, Y. Xu, G. Governatori, and T. Ito, editors, *Principles and Practice of Multi-Agent Systems - 17th International Conference, PRIMA 2014*, volume 8861 of *Lecture Notes in Artificial Intelligence*, pages 294–310, Berlin, 2014. Springer.
- [16] C. K. Frantz. *Agent-Based Institutional Modelling: Novel Techniques for Deriving Structure from Behaviour*. PhD Thesis, University of Otago, Dunedin, New Zealand, 2015. Available under: <http://hdl.handle.net/10523/5906>.
- [17] A. Ghorbani, P. Bots, V. Dignum, and G. Dijkema. MAIA: a Framework for Developing Agent-Based Social Simulations. *Journal of Artificial Societies and Social Simulation*, 16(2):9, 2013.
- [18] A. G. Greenwald and M. R. Banaji. Implicit Social Cognition: Attitudes, Self-esteem, and Stereotypes. *Psychological Review*, 102:4–27, 1995.
- [19] A. Greif. *Institutions and the Path to the Modern Economy: Lessons from Medieval Trade*. Cambridge University Press, New York (NY), 2006.
- [20] D. Grossi, H. Aldewereld, and F. Dignum. Ubi Lex, Ibi Poena: Designing Norm Enforcement in E-Institutions. In P. Noriega, J. Vázquez-Salceda, G. Boella, O. Boissier, V. Dignum, N. Fornara, and E. Matson, editors, *Coordination, Organizations, Institutions, and Norms in Multi-Agent Systems II*, volume 4386 of *Lecture Notes in Computer Science*, pages 107–120, Berlin, 2006. Springer.

- [21] S. A. Haslam, N. Ellemers, S. D. Reicher, K. J. Reynolds, and M. T. Schmitt. The Social Identity Perspective Today: An Overview of its Defining Ideas. In T. Postmes and N. R. Branscombe, editors, *Rediscovering social identity*, pages 341–356. Psychology Press, 2010.
- [22] G. M. Hodgson. The Evolution of Institutions: An Agenda for Future Theoretical Research. *Constitutional Political Economy*, 13(2):111–127, 2002.
- [23] G. Hofstede, G. J. Hofstede, and M. Minkov. *Cultures and Organizations: Software of the Mind*. McGraw-Hill, New York (NY), 3rd edition, 2010.
- [24] ISO. ISO/IEC 14977 – Information Technology - Syntactic Metalanguage - Extended BNF, December 1996. Accessed on: 1st December 2015.
- [25] S. C. Johnson. Hierarchical Clustering Schemes. *Psychometrika*, 2:241–254, 1967.
- [26] F. López y López, M. Luck, and M. d’Inverno. Normative Agent Reasoning in Dynamic Societies. In N. R. Jennings, C. Sierra, L. Sonenberg, and M. Tambe, editors, *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multi-Agent Systems*, AAMAS ’04, pages 535–542, New York (NY), 2004. ACM Press.
- [27] J. Morales, M. López-Sánchez, J. A. Rodríguez-Aguilar, W. Vasconcelos, and M. Wooldridge. Online Automated Synthesis of Compact Normative Systems. *ACM Trans. Auton. Adapt. Syst.*, 10(1):2:1–2:33, Mar. 2015.
- [28] J. Morales, M. López-Sánchez, J. A. Rodríguez-Aguilar, M. Wooldridge, and W. Vasconcelos. Automated Synthesis of Normative Systems. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-Agent Systems*, AAMAS ’13, pages 483–490, 2013.
- [29] J. Morales, M. López-Sánchez, J. A. Rodríguez-Aguilar, M. Wooldridge, and W. Vasconcelos. Minimality and Simplicity in the On-line Automated Synthesis of Normative Systems. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-Agent Systems*, AAMAS ’14, pages 109–116, 2014.
- [30] J. Morales, I. Mendizabal, D. Sanchez-Pinsach, M. López-Sánchez, and J. A. Rodríguez-Aguilar. Using IRON to Build Frictionless On-line Communities. *AI Commun.*, 28(1):55–71, 2015.
- [31] D. C. North. *Institutions, Institutional Change, and Economic Performance*. Cambridge University Press, New York (NY), 1990.
- [32] D. C. North. Institutions. *Journal of Economic Perspectives*, 5(1):97–112, 1991.
- [33] T. Parsons. *The Social System*. Routledge, New York (NY), 1951.
- [34] C. S. Peirce. How to Make Our Ideas Clear. *Popular Science Monthly*, 12:286–302, January 1878.
- [35] E. A. Posner. Law, Economics, and Inefficient Norms. *University of Pennsylvania Law Review*, 144:1697–1744, 1996.
- [36] R. A. Posner and E. B. Rasmusen. Creating and Enforcing Norms, with Special Reference to Sanctions. *International Review of Law and Economics*, 19(3):369–382, 1999.
- [37] R. Riveret, A. Artikis, D. Busquets, and J. Pitt. Self-governance by Transfiguration: From Learning to Prescriptions. In F. Cariani, D. Grossi, J. Meheus, and X. Parent, editors, *Deontic Logic and Normative Systems*, volume 8554 of *Lecture Notes in Computer Science*, pages 177–191. Springer, Springer, 2014.
- [38] Y. Shoham and M. Tennenholtz. On the Synthesis of Useful Social Laws for Artificial Agent Societies. In *Proceedings of the 10th National Conference on Artificial Intelligence (AAAI ’92)*, pages 276–281, San Jose (CA), July 1992.
- [39] Y. Shoham and M. Tennenholtz. On Social Laws for Artificial Agent Societies: Off-line Design. *Artificial Intelligence*, 73(12):231 – 252, 1995. Computational Research on Interaction and Agency, Part 2.
- [40] Y. Shoham and M. Tennenholtz. On the Emergence of Social Conventions: Modeling, Analysis, and Simulations. *Artificial Intelligence*, 94(1–2):139–166, July 1997.
- [41] S. Siddiki, C. M. Weible, X. Basurto, and J. Calanni. Dissecting Policy Designs: An Application of the Institutional Grammar Tool. *The Policy Studies Journal*, 39(1):79–103, 2011.
- [42] A. Smajgl, L. R. Izquierdo, and M. Huigen. Modeling Endogenous Rule Changes in an Institutional Context: The ADICO Sequence. *Advances in Complex Systems*, 2(11):199–215, 2008.
- [43] H. Tajfel. Social Identity and Intergroup Behaviour. *Social Science Information*, 13(2):65–93, 1974.
- [44] Q. van Doosselaere. *Commercial Agreements and Social Dynamics in Medieval Genoa*. Cambridge University Press, Cambridge (UK), 2009.
- [45] D. Villatoro, G. Andrighetto, J. Sabater-Mir, and R. Conte. Dynamic Sanctioning for Robust and Cost-efficient Norm Compliance. In T. Walsh, editor, *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, volume 1 of *IJCAI’11*, pages 414–419. AAAI Press, 2011.
- [46] G. H. von Wright. *An Essay in Modal Logic*. North Holland, Amsterdam, 1951.
- [47] C. Watkins. *Learning from Delayed Rewards*. PhD thesis, Cambridge University, Cambridge (UK), 1989.
- [48] O. E. Williamson. *Markets and Hierarchies, Analysis and Antitrust Implications: A Study in the Economics of Internal Organization*. Free Press, New York (NY), 1975.
- [49] O. E. Williamson. *The Economic Institutions of Capitalism*. Free Press, New York (NY), 1975.
- [50] O. E. Williamson. Transaction Cost Economics: How it works; Where it is headed. *De Economist*, 146(1):23–58, 1998.

Edited by: Hoa Dam

Received: April 16, 2015

Accepted: January 13, 2016