

Daniel Villatoro, Giulia Andrighetto, Rosaria Conte and Jordi Sabater-Mir (2015)

Self-Policing Through Norm Internalization: A Cognitive Solution to the Tragedy of the Digital Commons in Social Networks

Journal of Artificial Societies and Social Simulation 18 (2) 2
<http://jasss.soc.surrey.ac.uk/18/2/2.html>

Received: 01-May-2013 Accepted: 01-Jan-2015 Published: 31-Mar-2015

 Abstract

In the seminal work 'An Evolutionary Approach to Norms', Axelrod identified internalization as one of the key mechanisms that supports the spreading and stabilization of norms. But how does this process work? This paper advocates a rich cognitive model of different types, degrees and factors of norm internalization. Rather than a none-or-all phenomenon, we claim that norm internalization is a dynamic process, whose deepest step occurs when norms are complied with thoughtlessly. In order to implement a theoretical model of internalization and check its effectiveness in sustaining social norms and promoting cooperation, a simulated web-service distributed market has been designed, where both services and agents' tasks are dynamically assigned. Internalizers are compared with agents whose behaviour is driven only by self-interested motivations. Simulation findings show that in dynamic unpredictable scenarios, internalizers prove more adaptive and achieve higher level of cooperation than agents whose decision-making is based only on utility calculation.

Keywords:

Self-Organisation, Norms, Emergent Behavior, Cognitive Modelling, Artificial Social Systems

 Introduction

- 1.1 Norm internalization has long been studied in the social-behavioral sciences, social psychology and moral philosophy (Piaget 1965; Kohlberg 1981). By 'norm internalization', we refer to the process that transforms the motivations of agents for complying with social norms from those of external reward or punishment to that of following norms as an end in themselves (Andrighetto et al. 2010b). Consistent with Ullmann-Margalit (1977), we consider norms as prescribed or proscribed ways of behaving that are widespread in a group or population. We claim that norms influence people by pushing them to modify their mental representations and form new ones (Conte et al. 2013; Andrighetto et al. 2010b; Conte & Castelfranchi 2006). Humans may decide to comply with norms for many reasons, such as external or internal sanctions, which may be positive or negative, the legitimacy of the normative source, or because "norms should be respected" (Kantian morality).
- 1.2 As suggested by Axelrod (1986), norm internalization is a key process that supports norms that are only partially established, and it entails several advantages for maintaining social order. One is *greater robustness*. Compliance with a norm is generally more robust when a norm has been internalized by agents than when their compliance is solely motivated by external incentives. Consider that if everybody in a population internalizes a norm, there is no incentive to defect and the norm remains stable even without external enforcement (Gintis 2004a). Moreover, once internalized, individuals in many circumstances want, and influence, others to observe the same norms that they do. An effect of this is that norm internalization is decisive, if not indispensable, for achieving *distributed social control*.
- 1.3 While the importance of norm internalization as a promoter of social order is now widely recognized, some fundamental questions, most of which revolve around the issue of the proximate causes, remain to be definitively answered. Why and how do people internalize social inputs like commands, values, norms and tastes, and transform them into endogenous states? Which types of mental properties and ingredients must individuals possess to exhibit the multiple forms of compliance, from fully a deliberative choice to less considered or even automatic ones? What are the specific implications of different forms of norm compliance for society and its governance? These questions have, so far, received no conclusive answer. In particular, no explicit, controllable, and reproducible model of the process of norm internalization within the mind is currently available.
- 1.4 The present work aims to start filling this gap, by (1) sketching the building blocks of a cognitive model of norm internalization, (2) turning it into a simulation model, (3) checking, through agent-based simulation, the potential of our cognitive model in achieving norm internalization and (4) observing whether and to what extent this process of norm internalization affects the performance of a group of agents. The central purpose of this paper is to present an operational model of the internal processes and mechanisms of norm internalization. We identify sufficient conditions for agents to move from a deliberative cost-benefit compliance decision to one in which they pursue norm compliance as an end in itself. Far from depicting norm internalization as mindless conformity, automatic execution of a norm, we characterize it as a *multi-step* process, occurring at various levels of mental depth and giving rise to more or less robust compliance, and describe it as a *flexible* phenomenon enabling agents to de-internalize it and retrieve full control over those norms which they previously internalized.
- 1.5 We begin by identifying the set of mental components that characterize and constitute different levels and degrees of internalization: from the more deliberate to the more automatic. We suggest that, while sinking in the mind to different levels of depth, norms generate and activate different types of representations and mechanisms that can gradually lose track of their normative origin and become fully endogenous or even integrated with action plans. Thus, norm internalization is not an all-or-none phenomenon, but a *multi-step* process which consists of degrees and levels characterized by different mental ingredients (see Section 3). Additionally, norm internalization will be characterized as a *flexible* phenomenon, allowing norms to be de-internalized, automatic compliance blocked, and deliberation restored in certain circumstances (Hassin et al. 2009; Kennedy & Bugajska 2007). Internalized norms are not inexorably bound to remain as such: under extreme conditions, agents may retrieve awareness of their exogenous source and of their external enforcement. Even though automated responses are often highly efficient, they are rarely needed all the time. A completely automatized response may become counterproductive, or even dangerous in some cases. When certain conditions activate inconsistent prescriptions a given routine is blocked, allowing agents to easily adapt even to dynamic and unpredictable situations.
- 1.6 Operationalizing a multi-step and flexible model of norm internalization requires a versatile agent architecture. The modular normative architecture EMIL-I-A (EMIL Internalizer Agent) is a good candidate for this undertaking (Andrighetto et al. 2013; Conte et al. 2013; Andrighetto et al. 2010b). It is an agent architecture with normative mental modules, allowing it to internalize norms and to block automatic normative behavior when necessary. A key feature of this normative architecture that allows the process of norm internalization to take place is agents' ability to track the salience of a norm; they are able to measure and record the perceived prominence of a norm within a group. We investigate the process of norm salience estimation by agents and contributing factors. A proof-of-concept simulation, aimed to test the introduced concepts of norm internalization, is then described. Agent based simulation helps us to understand the internal mechanisms and processes of norm internalization and allows us to test the effects of this 'deep' form of norm compliance in promoting and sustaining social order. In particular, the performance of EMIL-I-A agents are compared with that of agents unable to internalize norms and whose decision-making is based on self-interest in a simulation scenario that recreates a "Tragedy of the Digital Commons" (Adar & Huberman 2000). Results show that in a social dilemma-like situation agents able to internalize norms in a flexible way achieve a higher level of cooperation and adapt better to new and unpredictable situations than agents whose decision-making is based on self-interest only.
- 1.7 The paper is organized as follows. After a brief overview on the work on norm internalization in Section 2, we will present our theory, focusing on different types and levels of norm internalization (Section 3). In subsection 4.1, the EMIL-I-A architecture will be described. In subsection 5.3, a simulation model aimed to test EMIL-I-A in a social dilemma scenario will be presented, and the results will be discussed in a final section.

 Related work

- 2.1 In recent years, an interest with norm internalization has reemerged in the (computational) social sciences and in the evolutionary game theoretic study of pro-social behavior (Bicchieri 2006; Epstein 2006; Gintis 2004b). Joshua Epstein considers norm internalization as a process that leads to automatic, or thoughtless, conformity (Epstein 2006). In his view, internalization consists of learning not to think about norms. As Epstein argues: 'When I had my coffee this morning and I went upstairs to get dressed, I never considered being a nudist for the day' (Epstein 2006). Once entrenched, people, observes Epstein, conform to norms without thinking about them. The more often they have taken an action in the past, the more likely it is that they will do it again in the future.
- 2.2 Although interesting, we challenge the idea that internalizing a norm makes one just a mindless 'norm-executor' (Conte 2008). Individuals have many other goals, some internalized, that constantly compete with internalized norms, which can imply that sometimes even internalized norms are violated. It could happen that one leaves the house partially undressed because of an emergency, like the outbreak of a fire. Also, a highly entrenched and ubiquitous norm like that of truth-telling is violated in some circumstances; for example, when the damage inflicted by the truth to the recipient is more severe than that caused by a lie. It is even possible to regain control over normative actions that have become automatic. For example, a driver stopping at the red light might see a policeman asking her to move on. In such a case, the car driver needs to be able to retrieve control of her action, block the automatism, and decide which normative input should be given priority.
- 2.3 The norm internalization process has several advantages for generating social order, such as increasing compliance and distributing norm enforcement and spreading. Individuals who internalize norms

are not only much better at complying with, but also at defending them than externally enforced individuals are (Gintis 2004b). Once an individual has internalized the norm, she expects it from herself and also from others. However, as argued by Gintis (2004b), to be advantageous, norm internalization must also be *flexible* and allow people to adapt in highly unpredictable and rapidly changing environments. The faster individuals' goals change, the likelier it is that they adapt to a dynamic environment. Hence, the higher their chances of survival and reproduction. Norm internalization leads agents to endogenize certain goals. For example other-regarding motivations, which had less time to evolve through natural selection, might result from norm internalization.

- 2.4 We are sympathetic with the explanation provided by Gintis of the adaptive role of norm internalization and the necessity of the flexibility of this process. This work is aimed to model the internalization process not as mindless conformity, but as flexible conformity by providing an analysis of the cognitive mechanisms that allow norm internalization and goal altering to take place.



A multi-step and flexible model of norm internalization

- 3.1 As claimed in Section 1, norm-internalization is the process by means of which agents comply with norms as an end in themselves and not because of external reward or punishment. We consider it as a *multi-step* process that occurs at different levels, from the fully deliberative to the fully automatic. In this section, we sketch a preliminary model of different levels and degrees of norm internalization (Kennedy & Trafton 2007). We suggest that the first level of norm internalization consists of the mental process that takes the normative belief (a belief that a given behavior, in a given scenario, for a given set of agents, is either forbidden, obligatory, or permitted) and provides an *internalized normative goal*¹¹ as an output. In other words, the normative goal is no longer relativized to the enforcement normative belief (i.e., a belief that normative compliance and violation are supported or enforced by positive or negative (informal) sanctions) but is now endogenized, i.e., it has become an end in itself, needing no external enforcement to be complied with. When an internalized normative goal is created, enforcement, if any, is self-administered through feelings of guilt, self-deprecation, loss of self-esteem, or other negative self-evaluations in case of violations, and pride, enhanced self-esteem, security, or other favorable self-evaluations in case of conformity (Reykowsky 1982).
- 3.2 As normative goals sink deeper into the mind, norms gradually lose track of their normative origin and become fully endogenized. At an intermediate level, the norm becomes an *internalized goal* of the agent. The normative belief may still persist, but the agent pursues the corresponding goal irrespective of it. Consider specific dietary regimes. Initially, an individual may decide to follow a vegetarian diet for ethical reasons. After time, she may come to strongly dislike the taste of meat and continue abstaining from eating meat, no longer for ethical reasons but for satisfying a personal goal. Kingsley (1949, p. 55) refers to this level of norm internalization as the process that occurs when:
- a norm [...] is a part of the person, not regarded objectively or understood or felt as a rule, but simply as a part of himself [...]
- 3.3 At the last and deepest level of internalization, *no decision-making* takes place, and compliant behavior is fully automatic. A specific perceived event triggers a conditioned action. For instance, stopping when traffic lights turn red, which consists of a sequence of movements necessary to apply the breaks - a behavioral response so deeply internalized that one can hardly make it explicit. At this level of internalization, as in the previous one, the normative beliefs may still be present in the agent's mind, but they are not the reason why she applies the norm. This last type of internalization corresponds to what Epstein calls thoughtless conformity. Other authors Tobias (2009) might see it as an example of habituation.
- 3.4 As this brief taxonomy shows, there is more than one form of norm internalization, and each of them is characterized by a specific mental configuration. Moreover, all of them are reversible. Perceiving that a norm has lost its salience can cause the de-internalization of the normative goal, and normative conflict can make an individual re-gain control over an automatic action and refrain from applying a given routine.
- 3.5 Here, we present a normative architecture, EMIL-I-A, that accounts for the first level of norm internalization. It explicates the mental process that takes a norm as an input and provides an internalized normative goal as output. At this point, the internalized normative goal is endogenous and does not need external enforcement to ensure compliance. We will also show how, in certain circumstances, the norm can be de-internalized and the norm-adoption decision restored. The remaining levels of norm internalization will be the object of future inquiry.



Two factors favoring norm internalization

- 4.1 In this section, we discuss two factors that can favor the process of norm internalization: *norm salience* and *calculation cost saving*. Other factors, such as cognitive consistency, self-enhancing effect, self-determination, emotions, etc., have been identified and discussed in Andrighetto et al. (2010b). Here, we will focus only on norm salience and calculation cost saving. These are the conditions that have been implemented in the EMIL-I-A agents.
- 4.2 First, we suggest that *highly salient norms* are ideal candidates for internalization. A norm can be perceived by an individual as more or less prominent and active within a group, and we refer to this perceived value as *norm salience* (Cialdini et al. 1990; Andrighetto et al. 2013; Houser & Xiao 2010; Bicchieri 2006). Psychological evidence shows that the more salient a norm is perceived to be, the greater its impact on the goal to comply with it (Cialdini et al. 1990).
- 4.3 The importance of a norm is one of the factors that affects the perception of its salience. The ability to monitor variations in a norm's salience allows individuals to better predict the actions, expectations, and willingness of others to react to violations, and in turn adapt their own behavior to others' conduct. Individuals update the salience of their normative beliefs and goals according to the information that they gather from their social, and non-social, environment using a variety of cues, of which others' actions is one, to infer how salient a given norm is. For example, the amount of compliance, the efforts and costs sustained in educating the population to form a certain norm, the visibility and explicitness of the norm, the credibility and legitimacy of the normative source are all signs by which people infer how important and active a social norm is in a specific context (Faillo et al. 2012; Cialdini et al. 1990).
- 4.4 As claimed in previous work, the way and the degree to which a norm is enforced also plays a crucial role in the dynamic of its salience (Andrighetto et al. 2013; Conte et al. 2013; Andrighetto & Castellfranchi 2013). If properly designed, the enforcement mechanisms can have both a coercive and a norm-signaling function. To make this distinction more vivid, we distinguish between two enforcement mechanisms, *punishment* and *sanction*, which have different capacities to convey normative information. Punishment works only by imposing a cost on the 'wrongdoer', reducing his or her material payoffs. In addition to inflicting a cost, sanction also communicates that the sanctioned behavior is not approved of because it violated a norm (Galbiati & D'Antoni 2007; Houser & Xiao 2010; Masclot et al. 2003; Giardini et al. 2010; Sunstein 1996). Sanctions convey a great deal of norm-relevant information that has the effect of making norms explicit and increasing their salience. To deter wrongdoers from future violations, *punishment* relies only on its coercive component, while *sanction* combines the coercive component with the norm-signaling one, also exploiting the motivational power of norms.
- 4.5 Salience may increase to such a degree that a norm becomes internalized. Conversely, if salience falls under a certain threshold, the norm ceases to be internalized, as it happens when people perceive violators, indicating that the norm in question is losing importance or is no longer operative.
- 4.6 Second, humans are *parsimonious calculators*. Under certain conditions, norms are internalized in order to save calculation processing and execution time (Bicchieri 1990). After weighing the costs and benefits of complying or not with a norm a certain number of times, each time reaching the same decision, individuals stop calculating and take norm compliance, or norm breaking, as their automatic best choice. By doing so, they save time and avoid errors that lead to punishment. Evolutionary analyses suggest that natural selection may have favored the internalization of norms because this mechanism saves on information processing costs and the associated errors (Gintis 2004b; Chudek & Henrich 2011).
- 4.7 As we will discuss in the next section, EMIL-I-A agents are designed so that when both the norm salience and calculation cost saving conditions are satisfied, the norm is internalized and the normative act is performed automatically.

EMIL-I-A: The internalizer architecture

- 4.8 In this section, we present EMIL-I-A (see Figure 1), an extension of the EMIL-A architecture, that includes the capacity of norm internalization. This work, to our knowledge, is the first to provide a normative architecture that is able to orchestrate the processes of internalization and de-internalization of norms in a dynamic and flexible way (see also Criado et al. 2010). EMIL-A is an agent architecture applied to the simulation of norm emergence, innovation and spread. It is a cognitive architecture endowed with modules that allow the recognition of norms, detecting their salience, updating it and deciding whether to comply or not. The main extensions of EMIL-I-A compared to EMIL-A are the introduction of a) the norm salience module, b) the urgent events management module, and c) the internalization process.

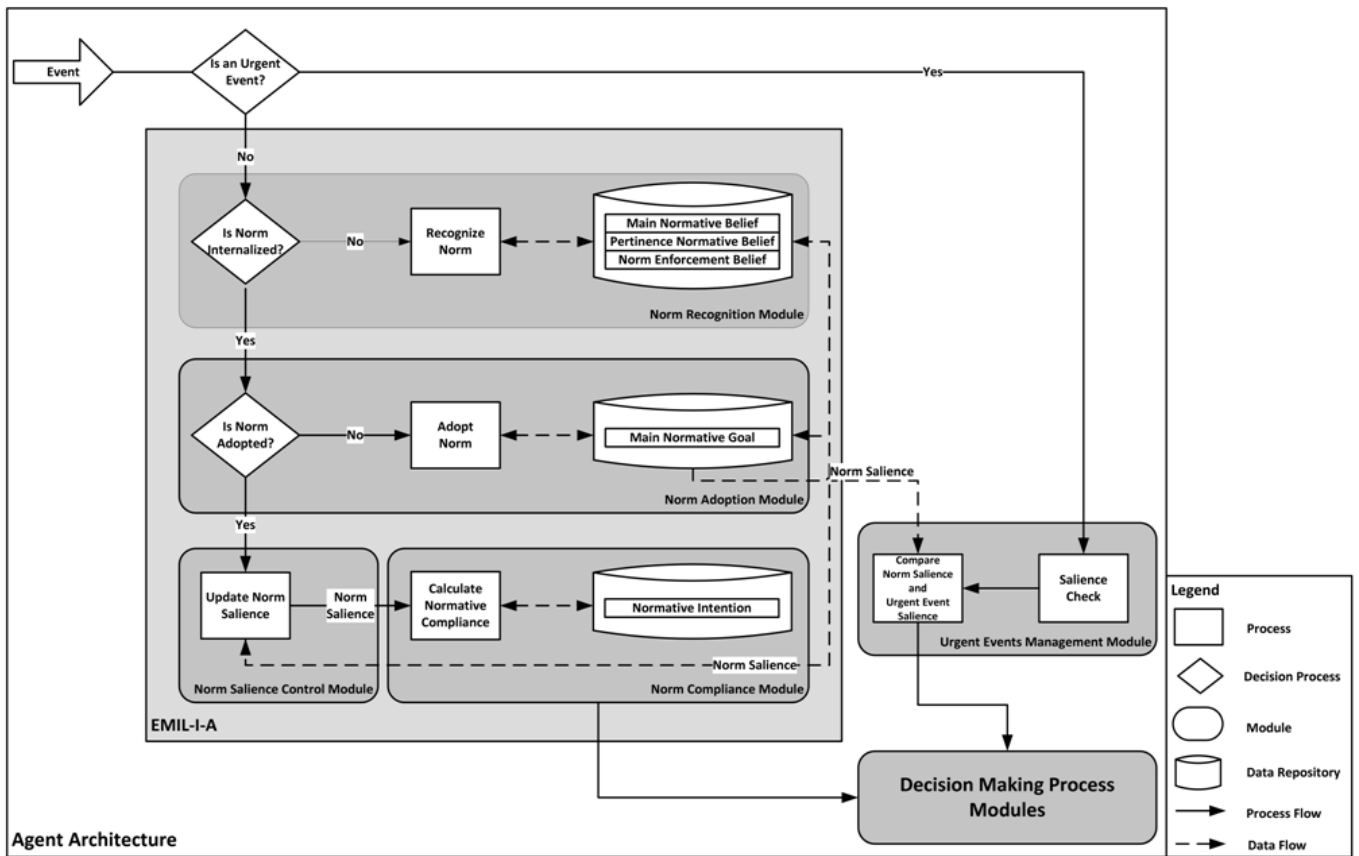


Figure 1. The main components of the EMIL-I-A architecture and the Urgent Events Management Module

Norm recognition and norm adoption modules

- 4.9 The Norm Recognition Module allows EMIL-I-A agents to interpret an observed behavior or communicated social request as normative and to form the corresponding normative beliefs and goals (see Conte et al. 2013; Andrighetto et al. 2010b; Campenni et al. 2009) for a detailed description of how the norm recognition module works). After exposure to the normative behaviors of others and to their explicit or implicit normative requests, agents potentially acquire those normative beliefs. Normative beliefs consist of three components: 1) the *main normative belief*, stating that there is a norm, *N*, prohibiting, prescribing, permitting action *a*, 2) the *normative belief of pertinence*, indicating that the agent who holds the belief belongs to the set of agents to which the norm applies, and finally 3) the *norm enforcement belief*, stating that norm compliance and violation are supported by positive or negative (informal) sanctions.^[2] In order to motivate agents to comply with it, the norm, once recognized, has to generate the corresponding normative goal in the agents' mind. Normative beliefs, because of the Norm Adoption Module, can generate normative goals as a means to avoid punishment and obtain rewards. Before internalization, the agent must calculate the costs and benefits of complying with or violating the norm in order to decide whether to observe it or not. But, as mentioned before, once internalized, this cost-benefit analysis is interrupted and the norm addressee will start to comply with it independently of external sanctions and rewards.

Saliency control module

- 4.10 In previous work (Andrighetto et al. 2010a), norm recognition heuristics have been applied to scenarios ruled by a single norm. When more than one norm is in force (especially when there are conflicting norms), agents need a way to discriminate between them and rank them. The Saliency Control Module provides EMIL-I-A agents with the ability to detect the relative importance of each norm and to update its saliency. The module is dynamically fed by the social information acquired by observing and interacting with the other agents.^[3]
- 4.11 Each of the normative cues gathered by the agents (see Table 1) is aggregated with different weights. Those behavioral or communicative acts, such as acts of norm obedience (C, O), norm invocation (E) and norm defense (e.g., sanctions S or punishments P) (C, O, P, S, E), that are interpreted as compliant with or defending a norm, make its saliency increase (see Table 1). Violations (NC) and unpunished violations (NPD), on the contrary, make norm saliency decrease, signaling that the social group is losing interest in that norm and does not invest in its enforcement. The values and the ranking order reported in Table 1 have been extracted from Cialdini et al. (1990).^[4]

Table 1: Social Cues and Weights for the Norm Saliency Aggregation Function. *C* represents the self-normative compliance or defection. *O* represents the normalized number (with respect to the total amount of neighbors) of observed norm compliance in the local environment. *NPD* represents the normalized number (with respect to the total amount of neighbors) of non-punished defectors in the local environment. *P* represents the normalized number (with respect to the total amount of punishments that can potentially be observed) of punishments in the local environment. *S* represents the normalized number (with respect to the total amount of sanctions that can potentially be observed) of sanctions in the local environment. *E* represents the normalized number (with respect to the total amount of explicit norm invocations that can potentially be observed) of explicit norm invocations in the local environment.

Social Cue	Weight
Self Norm Compliance/Violation (C)	$w_C = (+/-) 0.99$
Observed Norm Compliance (O)	$w_O = (+) 0.33$
Non Punished Defectors (NPD)	$w_{NPD} = (-) 0.66$
Punishment Observed/Applied/Received (P)	$w_P = (+) 0.33$
Sanction Observed/Applied/Received (S)	$w_S = (+) 0.99$
Norm Invocation Observed/Received (E)	$w_E = (+) 0.99$

- 4.12 Norms' saliency is updated according to the formula below and the social weights described in Table 1:

$$Sal_n^t = Sal_n^{t-1} + \frac{1}{\alpha \cdot \phi} (w_C \cdot O + w_{NPD} \cdot NPD + w_P \cdot P + w_S \cdot S + w_E \cdot E)$$

where Sal_n^t represents the saliency of the norm *n* at time *t*, α the number of neighbors of the agent, ϕ the normalization value that ensures that the saliency value does not go below 0 or exceed 1^[5] $w_{[C, O, NPD, P, S, E]}$ the weights specified in Table 1, and finally *O*, *NPD*, *P*, *S*, *E* indicate the registered occurrences of each cue. The resulting saliency measure $Sal_n^t \in [0, 1]$ (0 representing minimum saliency and 1 maximum saliency) depends on the social information that each agent gathers from the environment. Since every agent has access only to limited information, the value of the norm saliency may differ from agent to agent.

- 4.13 The norm saliency measure allows agents to decide which mechanism to use to enforce norms. EMIL-I-A agents can use two different enforcing mechanisms: *punishment* or *sanction* (see Section 4). An agent punishes by inflicting a cost on the norm breaker, reducing the payoffs of the punished agent. In addition to inflicting a cost, a sanctioner also communicates information about the norm that has been violated. The initial probability of inflicting a punishment act when another agent is observed acting against norms is set to 0.5 and is subsequently negatively affected by the number of defectors.

The probability of sending a normative message (associated or not with punishment) is a direct function of the perceived salience of the norm. When EMIL-I-A recognizes the existence of a norm, then it uses sanction to enforce it, instead of punishment. The probability of inflicting a sanction increases as a function of the perceived salience of the norm. Punishment and sanction have different impacts on norm recognition and salience (see Table 1): because of an explicit signaling component, sanction is more likely to be interpreted as an act enforcing and defending a norm than punishment.

- 4.14 Moreover, when agents have the belief that a norm is highly salient, one of the two necessary conditions for the internalization of the correspondent normative goal (see Sec. 4.1.4), they are more likely to internalize the correspondent normative goal.
- 4.15 Once a norm has been internalized, the salience mechanism remains active and continues to monitor the changes in its salience value. For example, if the salience of the internalized norm decreases after a certain threshold, the normative goal will be de-internalized. Even though the computational costs for norm compliance are not reduced to zero (because the salience mechanism keeps working), they are substantially decreased.
- 4.16 The Norm Control Salience Module is a key feature that improves agents' performance in several ways: it allows EMIL-I-A agents to *dynamically monitor* if the normative scene is changing and *adapt* to it. For example, if norm enforcement suddenly decreases in an unstable social environment, agents with highly salient norms are less inclined to violate them than agents with less salient norms. Vice versa, if a specific norm decays, EMIL-I-A agents are able to detect this change, stop complying with it and adapt to the new state of affairs. Considering the normative decision making of agents as a rule based system, the norm salience module enables them to dynamically modify the preference order of these rules, scoring the more salient norms higher in the list.
- 4.17 Finally, when facing a normative conflict, norm salience allows the agents to decide which action to perform, providing them with a criterion for comparing the norms applicable to the context. For example, the agent will decide whether to wait or move on the basis of the respective salience of the two norms "stop at red traffic light" and "clear the road when an ambulance comes along". Compared to rule-based engines, this module allows the temporal insertion of a higher priority rule when triggered by an external signal (the ambulance alarm in the previous example), leaving the other rules unattended while the signal is active.

Urgent events management module

- 4.18 In some situations, like emergencies, agents must be able to recognize the situation promptly and subsequently break a norm. By this means, agents can decide to violate a norm, for example, deciding to stop at green traffic lights when an ambulance with its sirens activated comes along at high speed from the opposite direction. The urgent events management module allows EMIL-I-A agents to manage these extreme and unpredictable situations.
- 4.19 In the present model, in order to violate norms when required to do so by contingent circumstances, agents need to explicitly record and compare the situations faced (in the ambulance example, the light and sound of the sirens cues the socially appropriate, and desirable, action in the mind of the driver, despite the norm *cross at green traffic light*). All the *Urgent Events* are assigned a salience value ($Sal_t^{urgent_event}$) that agents can compare with the salience of the norm regulating the situation, and decide accordingly (see Figure 1).^[6] In situations classified as *urgent*, the decision-making module compares the degree of the urgent event with the salience of the norms that usually regulate that situation. If the salience of the urgent event is higher than that of the norms, the agent performs the action associated to the urgent event. Otherwise, it will use the norms to decide what to do. The salience mechanism allows this management module to be compared with an adaptive rule-based system: the order of the rules dynamically changes according to the salience of each norm, and to urgent events when these are to be taken into consideration.

Internalization process

- 4.20 The internalization module is responsible for selecting which norm should be internalized; converting it into an internalized normative goal that is fired without decision-making. An agent can internalize only one norm regulating a certain situation (e.g., people greet each other by waving from a respectful distance, by a friendly handshake, or finally by kissing; depending on the environment, only one of these three actions can be internalized). Moreover, because of the information provided by the *Salience Control Module*, this mechanism is also in charge of the de-internalization process.
- 4.21 EMIL-I-A agents are designed as *parsimonious calculators*: under certain conditions, they internalize norms in order to save calculation and execution time. Upholding a norm that has led one to succeed reasonably well in the past is a way of economizing on the calculation costs that one would sustain in a new situation. To record repetition in their decision-making, agents track (*Eval*) the number of times that they have performed this decision-making calculation and arrived at the same decision. If the decision-making calculation returns a different decision, the counter will be reset, representing a break up of the repetition.
- 4.22 The norm n is necessarily transformed into an *internalized* normative goal by EMIL-I-A,^[7] when both the following conditions are satisfied: (1) the salience of the candidate norm is at its maximum value ($Sal_t^n = 1$), and (2) the decision-making calculation has returned the same decision a certain number of times (in the present model we fixed the calculation repetition tolerance to 10, therefore $Eval^n \geq 10$). On the contrary, a norm n_x is *de-internalized* when the following conditions apply: (1) the salience of the internalized norm decreases to its minimum value ($Sal_t^n = 0$), and (2) the salience of another norm (ruling exactly the same type of interaction) exceeds the internalized one ($Sal_t^y > Sal_t^n$). These two conditions ensure that in highly norm-competitive scenarios, where two norms have similar salience, both norms are not continuously internalized and de-internalize, but only when one of them is fully overridden by the other. In any case, these conditions should be tested empirically with laboratory experiments.



Self-Regulated Distributed Web Service Provisioning





- 5.1 Our simulation scenario is a web-service market populated with agents whose task is to find out which services are offered by other agents (distributed) and control the behavior of the rest of their peers (self-regulated). The simulated scenario captures fundamental features of the real-world phenomenon that it models: it is dynamic (new services with different capabilities can be created during the simulation), unpredictable and populated by heterogeneous agents. To test the performance of the proposed model of norm internalization in a multi-norm scenario, we implemented one example where agents are regulated by two possible norms and an urgent event may regulate the situation of sharing a common resource. This variant is a novel contribution with respect to previous works in the normative self-adapting community (such as Blanc et al. 2005; de Pinninck et al. 2007). The presented simulation scenario aims to represent a realistic setting while reducing the number of variables that one needs to define. The result is an elementary P2P simulation model that can be specified for different interaction protocols, underlying social networks, population size, exchanges and growing rates, etc. The exploration of these parameters and their relationship with real-world phenomena is out of the scope of this work and remains open for future research. At this stage we maintain the simplest version of the simulation model that allows us to test the implemented internalizing agent architecture and compare it with other architectures.

Motivation

- 5.2 The presented simulation model aims to reproduce a self-regulated web-service market. Inspired by the "Tragedy of the Digital Commons", we present a provider-consumer scenario, populated by providers who may suffer a tragedy of the commons caused by consumers' over-exploitation. The services we refer to present the following features: (1) they are available for use and provide utility to agents, (2) they are a common good, and (3) they can only handle a limited number of requests before their service starts losing quality.
- 5.3 While looking for services, consumers might query neighbors in their social network about who can provide them with the service they are looking for. Social networks are a source from which people can obtain information (given their topological structure reducing the distances amongst members while preserving simplicity) - for example provided by word-of-mouth - and represent an alternative source, with respect to traditional methods. Conventional approaches in multi-agent systems, such as *registries* (centralized repositories that store information used for matching agents offering services with agents needing services) or *matchmakers* (auxiliary agents that have the information necessary to couple agents offering services with those needing them), partially address the problem of finding service providers or information owners (Decker et al. 1997). In highly dynamic environments, however, some information that is valuable to agents cannot be stored in a centralized repository, as it might not be accessible or not defined functionally for design purposes. Such information, the updated details about the quality of the service or the availability of the service for instance, may be accessible only through social networks.
- 5.4 Here, we propose a hybrid approach: similarly to the *white pages* in UDDI (Curbera et al. 2002), our agents will query a central server to obtain pointers to service providers, and then, all other important information about the service will be provided to them by the service providers. Our system functions similarly to the one implemented by Napster (2006). More specifically, the following dynamics occur in our case study. Agents learn tasks, or look for and find services, that are necessary for them to satisfy their needs. Based on the number of agents using it, providers supply a certain quality of service to the agent, the quality of which can increase or decrease during the simulation. Agents obtain the services that they need by finding other agents who are willing to share their services. For the sake of simplicity, the allocation of new services and needs is done automatically by the system, with different allocation probability distributions depending on the experiment and specified accordingly in each experiment description.
- 5.5 By finding a service that fulfills their need, and which is of an adequate quality, agents receive a reward. If the quality of the service is not sufficiently high, agents requesting a *service*, *requesters*, will continue their search.
- 5.6 The service is a resource offered by an agent called the *service provider*. If a *service-provider* decides to share its service, this is automatically shared, becoming accessible at that same moment for the requester. When sharing, a *service-provider* remains blocked for a number of time-steps (representing the *transaction* time), thus reducing its possibility to look for and find its own needed service, and consequently perform any interaction. This *transaction* time is inspired by P2P networks; in these type of networks, when a service is shared, the bandwidth of both agents, one uploading and the other downloading, is reduced.^[8]
- 5.7 In the proposed scenario, purely self-interested agents will always free-ride and use the services provided by others but will never reciprocate by providing their own services in return. Widespread self-interested behavior leads to a "tragedy of the commons" situation in which a depletion of the shared resources occurs, as predicted in Adar & Huberman (2000).
- 5.8 Endowing agents with the EMIL-I-A architecture allows them to behave dynamically, changing the rate of service allocation and requests across time, and reflecting the performance of real systems. Our aim is to test how agents with EMIL-I-A architecture perform when facing this "tragedy of the digital commons" situation and to compare their performance against agents endowed with different

architectures.

Norms in the Web Service Scenario

- 5.9 As previously specified, service providers can share any of the held services with a requester. However, only service providers know, at run-time, the quality of the service they can offer, and the quality needed by requesters.
- 5.10 We define N as the set of norms in our Web Service Scenario that regulate the action of sharing a resource. A norm is specified as follows^[9]:

- 5.11 Similarly, an urgent event is specified as:

- 5.12 Finally, a service is specified as:

- 5.13 In the algorithms, we access the different elements of the specification structures using a C like syntax. For instance, if we have the norm n , we can access its antecedents using the form $n.ancecedents$.
- 5.14 There are two different norms in our scenario, a single urgent event that regulates the sharing of a generic service s , and three possible actions of agents ((Share_requested_service, Not_allowed_to_share, Do_not_share)). The first two actions are cooperative behaviors while the third is classed as a defection. The difference between Not_allowed_to_share and Do_not_share is that in the former, there is a socially accepted norm that prevents the agent from sharing the service (in the benefit of the society that imposes that norm) while in the second it is the agent's self interest that determines that it is better not to share the service.

- 5.15 The first norm (N1), "always share", is a norm of unconditional cooperation. The second (N2) is instructs agents to "share only if the capacity of the service provider has not reached the limit to guarantee a high quality service". Finally, the "Urgent event" reflects an exceptional generic necessity that requires the service to be shared.
- 5.16 Depending on the environmental conditions (the service's capacity distributions, in this scenario) and on the agents' needs (the services' expected quality, in this scenario), one or the other norm will become more salient, and thus govern the behavior of the system.
- 5.17 As described in Section 4, in order to enforce the two norms, N1 and N2, and to deter agents from violating them, EMIL-I-As can use two different enforcing mechanisms: *punishment* or *sanction*. The initial probability of inflicting a punishment act when another agent is observed acting against either N1 or N2 is set to 0.5 and then it is negatively affected by the number of defectors. Once agents recognize the existence of the norms, EMIL-I-A agents can use sanctions; the probability of inflicting a sanction increases as a function of the perceived salience of N1 and N2.

Simulation Model

- 5.18 As shown in Fig. 2, we employ a *Hybrid Decentralized* model (as defined in Androutsellis-Theotokis & Spinellis 2004) in which each agent can share a service with the rest of the network. All agents are connected to a central directory server that maintains (a) a table in which all the users' information is recorded (typically the IP addresses and the connection bandwidth, which in our system are represented, respectively, as IDs and the capacity of the service), and (b) a table listing the services from each agent, along with metadata descriptions of the services (such as the type, capacity and so on). Each agent that is looking for a service sends a query to the central server. The server searches for matches in its index, returning a list of users that could supply the required service. The user can then open direct connections with one or more of the peers that can satisfy the needed service. The final decision, taken by the peers, is to share the service or not. Kazaa, Gnutella2, eMule, and bitTorrent are real systems that use this specific type of resource-exchange P2P environment.
- 5.19 The central server maintains the following information: *Agents*, the set of agents in the system, *SM*, the social network that connects the agents, and *T*, the types of services that can be provided and looked for in the system. The complete knowledge of the social network is a strong assumption that simplifies the object of our research, and could be modified in future versions of this work. Other authors have applied similar approaches to P2P networks where the central agent does not possess complete knowledge of the network (de Pinninck et al. 2007).

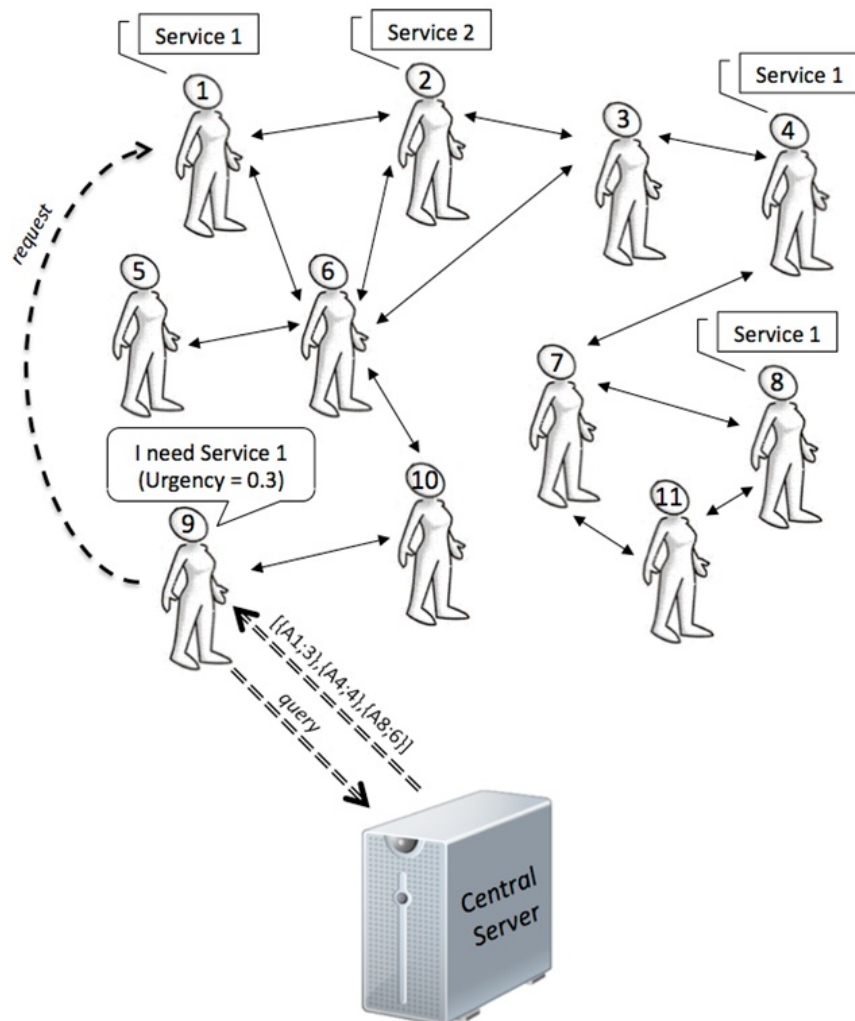


Figure 2. Social Network of the Web-Service Provisioning Scenario.

5.20 Agents are located in a social network that restricts their interactions and communications. The social network is described in the following way:

$$SN = \langle Agents, Rel \rangle,$$

where $Agents$ is the set of agents populating the network and Rel is a neighborhood function. The neighborhood function is symmetrical, therefore is $\forall a, b \in Agents$ if $Rel(a, b)$ then $Rel(b, a)$.

5.21 During the simulation, agents find services (or learn tasks) that are offered as services. Following the specification in section 5.2, a service is defined as:

$$Service = \langle provID, type, clients, capacity \rangle$$

where $provID$ represents the identification number of the service provider (e.g. IP address), $type \in T$ is the type of service offered by the service provider (e.g. storage, calculation time, data analysis), $clients$ indicates the number of clients currently using that service and $capacity$ is the number of clients that a service provider can satisfy while offering a high quality service.

5.22 Moreover, agents find out new needs that can be fulfilled by obtaining services from other agents:

$$Needs = \langle a, type, quality, deadline \rangle$$

where $a \in Agents$ is the agent ID for which this need applies, $type \in T$ represents the type of service an agent needs with a minimum Quality_of_Service level ($quality$), before the $deadline$ that is fixed to 100 time-steps after the generation of the request. The Quality_of_Service of a provider $provID$ is calculated in the following way:

$$QoS(provID) = \max\left(1, \frac{capacity}{clients}\right)$$

5.23 For the sake of simplicity, the allocation of new services and needs to the agents is done automatically by the system, with different allocation probability distributions depending on the experiment. Moreover, the service's capacity and expected quality follow two different probability distributions, both defined by the system and dependent on the environmental conditions that will be simulated. The expected quality of the service assigned to agents is fixed and cannot be changed by the agents. Moreover, the offered quality of the service changes dynamically, depending on the service capacity and the number of clients, forming a linear relationship (the more clients there are the worse the QoS).

```

Data: Server the central server.
Data: serviceAssignmentProbability which specifies the probability that an agent will be assigned a newly
created service.
Data: needAssignmentProbability which specifies the probability that an agent will be assigned the need of a
newly created service.
Data: S contains the list of all created services.
Data: deadline which specifies the deadline by when a service needs to be found, and fixed to 100 time-steps
after its creation.
1 for time-steps do
  /* Environmental Phase: Resource & Needs Assignment */
  2 for i  $\in$  Agents do
    3 coin = random([0,1]); /* generates a random number between 0 and 1 */
    4 if coin  $\leq$  serviceAssignmentProbability then
      5 s = createNewService(); /* generates a new service */
      6 S.add(s);
      7 i.assignService(s); /* assigns a service to agent i */
      8 for j  $\in$  Agents do
        9 coin = random([0,1]);
        10 if coin  $\leq$  needAssignmentProbability then
          11 j.assignNeed(s.type, deadline); /* assigns a need to agent j */
          12 end
        13 end
      14 end
    15 end
  /* Search & Exchange Phase */
  16 for i  $\in$  Agents do
    /* From all needs of agent i, select the most urgent one */
    17 n = i.mostUrgentNeed();
    /* Query the server about the service and get the list of service providers */
    18 serviceProviders = Server.serviceRequest(n);
    /* Ask for service to the nearest of the available service providers */
    19 contact = serviceProviders.nearestTo(i);
    20 response = contact.askForService(n.type, n.urgent event);
    21 if response == Share_requested_service then
      22 i.Block(contact.distance);
      23 contact.Block(contact.distance);
      24 i.removeNeed(n);
    25 end
    /* Evaluate response for norm violations, using Alg. ?? */
    26 normativeResponse = i.decideNormativeReaction(response, contact, n.type, n.urgent event);
    27 applyNormativeResponse(contact, normativeResponse);
    28 end
  /* Strategy Update Phase */
  29 forall the i  $\in$  agents do
    30 nc = i.ObserveNormativeCues(i.neighbours);
    31 i.UpdateStrategy(nc);
    32 end
  33 end

```

Algorithm 1. Simulation process.

5.24 The simulation is run for 20000 time steps^[10], and each time step is structured in the following way:

1. According to the environmental situation, each agent has some needs to satisfy and the possibility of providing services (lines 2-14 in Algorithm 1). The needs are only assigned to an agent when a service is created, and that is why we assign the needs after the creation of the service. The new service is created with a random type and a random capacity.^[11]
2. Depending on its most urgent needs (i.e., those with a shorter deadline), each agent asks the system^[12] for a needed service (line 18 in Alg. 1), thus receiving a list containing potential service providers. The server list has the following form $\langle a_0, d_0 \rangle, \langle a_1, d_1 \rangle, \dots, \langle a_k, d_k \rangle$, where a_i represents the agent identity, and d_i represents its distance from the requester.
3. The selected service-provider is the one nearest to the agent (line 19 in Alg. 1) and it is requested for the service (as it can be seen in the schematic communication protocol shown in Figure 3) with an *urgent_event* parameter (line 20 in Alg. 1).
4. As mentioned before, service providers can take any of these three decisions: "Share_requested_service", "Not_allowed_to_share" or "Do_not_to_share".
5. If the response by the service provider (received in line 20 of Alg. 1) is positive and the quality of the received service is at least equal to the expected one, the requesting agent receives the service. This agent may be requested to provide a service in the following round, itself becoming a service-provider (line 21-25 in Alg. 1). Both the original service-provider and the requester are occupied for a number of time steps (representing the transaction time specified by *contact_distance* in Alg. 1). If the quality of the service is insufficient, the requester will continue its search.
6. Having received the potential service provider's response, depending on the response and the requester's normative goal (specified by norms salience), the requester can decide (in line 26 in Alg. 1) to punish or sanction the service-provider. As written in section 5.2, punishment works by imposing a fine on the target, modifying the cost-to-benefit ratio of norm compliance and violation, while in sanctioning the fine is accompanied by a normative message, making explicit the existence and violation of the norm to the target (potentially to the audience also)^[13].
7. At the end of each time step (line 29-32 in Alg. 1), agents observe the interactions that have occurred around them, this way checking the amount of norm compliance, violation, punishment, and sanction, and subsequently update the norm's salience (as explained in Sec. 4.1). Our agents only have access to local information; they can record only the normative information of their direct neighbors in the social network, although the information that they receive is completely accurate.

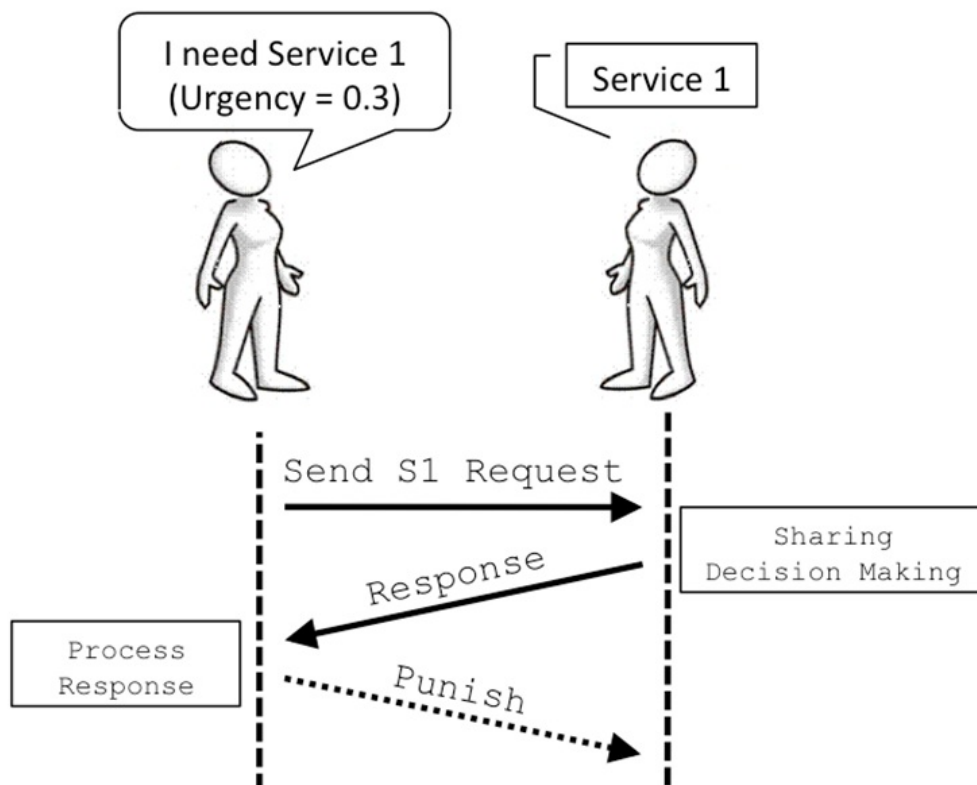


Figure 3. The Interaction Protocol Between Agents

5.25 While the present simulation model is a simplification of a real system, we think that it is endowed with the necessary ingredients to study the process of norm internalization in a more realistic scenario than a classical Prisoner's Dilemma. Another simplification of the present model concerns the normative knowledge held by the agents: they know which actions are compliant with or violate which norms. This assumption dispenses the need to solve problems associated with ontology alignment, inter-agent communication, and norm-representation languages, allowing our specific object of study to be focused.

Agent Architectures: EMIL-I-As vs IUMAs

5.26 In environments where the designer has perfect knowledge of both the system's behavior and the agents' needs, hardwired strategies are the best option: the system designer can use scheduling algorithms to synchronize the usage of the services amongst the agents to obtain the optimal distribution. However, in this work, we are interested in scenarios where the dynamics of the system are unpredictable and unfold at runtime. Therefore, we need agents that are able to adapt to environmental changes. We present two different types of agents whose behaviors are contrasted: IUMAs and EMIL-I-As.

IUMAs decision-making

5.27 The decision-making of the IUMAs is modeled using a classical reinforcement learning approach (as in Villatoro et al. 2009; Sen & Airiau 2007); in this type of architecture, the cognitive load of an action has no intrinsic value other than the payoff that it conveys. Therefore, IUMAs are unaware of, or ignore, their norm-compliance or violations and only know which actions return higher payoffs. Aiming to maximize their instantaneous utility, IUMAs share their services only if the probability of being punished is sufficiently high. Moreover, these agents never punish as punishing is costly and reduces their utility, and there is no possibility of punishing this second-order defection.

5.28 The IUMAs' decision-making for the first stage decision is modeled with a classic Q-Learning algorithm (as in Villatoro et al. 2009; Sen & Airiau 2007). The learning algorithm used here is a simplified version of the Q-Learning one (Watkins & Dayan 1992).

5.29 The Q-Update function for estimating the utility of a specific action is:

$$Q(a) \leftarrow (1 - \alpha) \times Q^t(a) + \alpha \times \text{reward} \quad (1)$$

where *reward* is the payoff received from the current interaction and $Q^t(a)$ is the utility estimate of action *a* after selecting it *t* times. When agents decide not to explore, they will choose the action with the highest Q value. The reward used in the learning process is the one obtained from interaction, considering also the amount of punishment received. In order to follow the methodology established in previous work on social learning (Sen & Airiau 2007; Mukherjee et al. 2007; Villatoro et al. 2009), in all the experiments presented in this article the exploration rate has been fixed at 25%, i.e., one-fourth of the actions are chosen randomly.

5.30 The IUMAs' decision-making for the second stage is also modeled with Q-Learning. However, as there are no potential risks to not punishing (as in the meta-punishment situation presented by Axelrod (1986), agents will always prefer not to punish.

EMIL-I-As decision-making

5.31 The second type of agent is EMIL-I-A. Its decision-making works differently before and after the internalization process. As described in Sec.4.1, EMIL-I-As are endowed with the norm recognition module, allowing them to acquire the normative belief that there is a norm and the enforcement normative belief that punishment is consequent to norm violation. Upon recognizing a norm, they calculate the convenience of complying with it or not. For example, if a norm is intensively defended through the application of punishments or sanctions, EMIL-I-As will observe it, otherwise they will violate it.

5.32 However, they do not know beforehand what the surveillance rates of the norm are. During the simulation, agents update (with their own direct experience and observed normative social information) the perceived probability of being punished. Before internalizing a norm, EMIL-I-As' decision-making is also sensitive to a *risk tolerance* rate: when the perceived punishment probability is below their *risk tolerance* threshold, agents will decide to violate the norm; otherwise, they will observe it. Although this process may provide agents with the highest possible benefits, it also imposes the computational cost of evaluating each of the options at every every time step.

5.33 Once the norm has been transformed into an internalized normative goal, EMIL-I-As can avoid this cost-benefit calculation: they will observe it in an *automatic* way. Nevertheless, the salience mechanism is still active, and it is continuously updated. This way, if necessary, agents are able to unblock the automatic action and to restore cost-benefit analysis in order to decide whether to comply or not.

5.34 Initially, only *norm-holders* know about the norms governing their environment. Any EMIL-I-A can (or cannot) be a norm-holder, depending on whether it has been initialized with knowledge about norms and their respective salience. How the initial number of norm-holders affects the performance of a group is tested in section 5.8.

5.35 EMIL-I-As take two decisions: (a) whether to share a service or not, and (b) whether to punish or sanction norm violators. Concerning the first choice (see Alg.2), to share a service or not, if no norm has yet been internalized, the agents evaluate all the available actions and they choose the one returning the highest payoff (lines 13-32 in Alg. 2). Otherwise, if the norm has been internalized, the agents behave in an automatic way: the decision to comply is no longer evaluated, but simply executed (lines 1-9 in Alg. 2). Agents automatically perform the internalized norm compliant action unless an urgent event request arrives (lines 10-12 in Alg. 2). In Alg. 2, *Cooperation*(*Coop*) refers to the norm compliant action and *Defection*(*Def*) to the norm defecting action.

5.36 In our scenario, an urgent event can only happen with respect to the second norm ("Share only high quality services", which in case of an urgent event is violated sharing services with any quality). We can observe the management of [urgent events in lines 1-12 of Alg. 2.^[14]


```

Input : Set of norms  $N$ , Required Service  $s$ , Requested urgent event  $u$ , Current time  $t$ 
Data :  $R_{Def} = 3$  which specifies Potential Reward for Defection.
Data :  $R_{Coop} = 0$  which specifies Potential Reward for Cooperation.
Data :  $Norm\_Activation\_Value$  which specifies the minimum required salience of a norm in order to be activated
/* If there is a norm that has been internalized and that has a salience higher than
the salience of the urgent event (if there is an urgent event), then apply that
norm
*/
/* NOTE:  $Sal_i^u = 0$  if there is no urgent event active
*/
1 forall the  $n$  in  $N$  do
2   if  $internalized(n) \wedge Sal_i^u \leq Sal_i^n$  then
3     if  $n.ancecedents\ hold$  then
4       /* If the antecedents hold perform the associated action
5       */
6       return  $n.action\_hold$ ;
7     else
8       /* If the antecedents fail perform the associated action
9       */
10      return  $n.action\_fail$ ;
11    end
12  end
13 end
/* If there is no norm internalized that has a salience higher than the salience of
the urgent event, apply the action associated to the urgent event (if it exists)
*/
14 if  $Sal_i^u > 0$  then
15   return  $u.action$ ;
16 end
/* If there is no urgent event and there are no internalized norms, do a benefit-cost
calculation to see if it is worth it to observe any norm
*/
17 updateBeingPunishedProb();
/* First, reduce the reward for defection if there is at least a norm that has a
perceived punishment probability above the risk tolerance level of the agent
*/
18 forall the  $n$  in  $N$  do
19   if  $riskTolerance \leq perceivedPunishmentProb(n)$  then
20      $R_{Def} \leftarrow$  Deduce Punishment Cost;
21     break;
22   end
23 end
/* Second, if the reward for defection is higher than the reward for cooperation,
defect. Otherwise cooperate.
*/
24 if  $R_{Def} \geq R_{Coop}$  then
25   /* In our example, the defect action is Do_not_share
26   */
27   return  $defectAction$ 
28 else
29   /* Apply the first norm that has a salience bigger than the Norm activation
30   threshold
31   */
32   forall the  $n$  in  $N$  do
33     if  $Sal_i^n \geq Norm\_Activation\_Value$  then
34       if  $n.ancecedents\ hold$  then
35         return  $n.action\_hold$ 
36       else
37         return  $n.action\_fail$ 
38       end
39     end
40   end
41 end
42 end

```

Algorithm 2. Resource Sharing decision-making.

5.37 In Alg. 2, $internalized(n)$ is the function that returns a boolean value specifying whether the agent has internalized the norm n and $perceivedPunishmentProb(n)$ is the function that returns the agent's perceived probability of being punished when not complying with the norm n .

5.38 If a violation is detected, a second decision, whether to punish, sanction, or do nothing, has to be taken^[5] The result of this decision depends on the salience of the violated norm. This decision-making process is domain specific and the process for our scenario is described in Algorithm 3. The cost of punishment is fixed to 1:4, meaning that punishment costs the punisher 1 unit while reducing the target's endowment by 4 units (the 1:4 punishment ratio is used because it has been found (Nikiforakis & Normann 2008) to be more effective in promoting cooperation). The decision-making of EMIL-I-As is not affected by the cost of applying a punishment/sanction to the target (as if this expense was to be considered we could fall into a second order dilemma, as expending resources is never rational in these type of scenarios); however, the costs of being punished/sanctioned affects the agents' decision to cooperate or defect in the next rounds.

```

Input : Requesting Agent i, Requested Agent j, Required Service s, Requested Quality  $q_{req}$ , Requested urgent
event u, response, Obtained Quality  $q_{rec}$ 
Data : Norm_Activation_Value which specifies the Initial Saliency Value of a Norm when Activated.
/* Response against an Excused Service Denial */
1 if response == Not_allowed_to_Share then
    /* If low quality was requested... */
    2 if  $q_{req} < 1$  then
        /* And N1 is more important, it is considered as a Punishable Violation */
        3 if  $Sal_i^{N1} > Sal_i^{N2}$  then
            if  $Sal_i^{N1} \geq Norm\_Activation\_Value$  then
                4 | i.sanction(j, N1);
            else
                5 | i.punish(j, N1);
            end
        end
    end
23 end
/* Response against a Service Denial */
11 if response == Do_Not_Share then
    /* If low quality was requested... */
    12 if  $q_{req} < 1$  then
        /* And any norm is active, it is considered as a Punishable Violation */
        13 if  $Sal_i^{N1} \geq Norm\_Activation\_Value$  then
            14 | i.sanction(j, N1);
        else
            15 | i.punish(j, N1);
        end
        17 else
            18 if  $Sal_i^{N2} \geq Norm\_Activation\_Value$  then
                19 | i.sanction(j, N2);
            else
                20 | i.punish(j, N2);
            end
        end
    end
23 end
/* Response against a Low Quality Sharing */
24 if response == Share_requested_service  $\wedge$   $q_{rec} < q_{req}$  then
    /* If it is not an urgent event... */
    25 if  $u < Sal_i^{N2}$  then
        /* And N2 is active, it is considered as a Punishable Violation */
        26 if  $Sal_i^{N2} \geq Norm\_Activation\_Value$  then
            27 | i.sanction(j, N2);
        else
            28 | i.punish(j, N2);
        end
    end
30 end

```


Experimental Design and Research Hypotheses

- 5.39 As previously stated, the number of variables in this simulation scenario are minimized and we change only those that can directly affect the decision making of the agents. While more complex simulation scenarios would result in more realistic settings, they also demand further developed agent architecture and interaction protocols.
- 5.40 To reduce the search space, the service and task assignment have been fixed to a constant rate of 10%: the services are created at an average rate of 1 every 10 time steps, and at that time step, 10% of the population will be assigned with tasks that need that service. Tasks are assigned only when services are created.
- 5.41 Agents are located in a scale-free network (that represents the social networks (Newman 2003; Albert & Barabasi 2002). This topology restricts the agents' observation window (with respect to the social and normative information) to their direct neighbors, but they can potentially interact (ask and provide services) with any other agent in the network.
- 5.42 The results presented in this section are the average results of 25 simulations.
- 5.43 All non-IUMAs agents are initialized with a constant propensity to violate norms and the perceived probability of being punished/sanctioned is equal to or lower than 30%. At the beginning of a simulation, some agents (the specific quantity varies from experiment to experiment) know the existence of the two norms, N1 *Always Share When Asked For a Service* and N2 *Share Only High-Quality Services*. We refer to these agents as *norm holders* and the initial saliency of their norms is set to 0.9. The simulations are populated with a constant number of 50 agents, with a variable number of EMIL-I-A (from 0 to 50) (whether or not norm holders) and IUMA (from 0 to 50) agents.
- 5.44 Before proceeding to present and discuss the experimental findings, we summarize the experimental hypotheses and the variables that are manipulated during the experiments. The main hypothesis we test is that in a social dilemma situation, EMIL-I-As achieve higher level of cooperation than agents whose decision-making is based on purely utility calculation (IUMAs) (see Sec. 5.7). Beside testing it, we are also interested in exploring how the presence of norm holders, agents that know which norms govern the environment from the start of the simulation, affects the norm internalization process (see Sec.5.8). Finally, we want to check the flexibility of the EMIL-I-A architecture with respect to its ability to (1) internalize the correct norm and de-internalize it if its saliency decreases (as shown in Sec.5.6), attending also to local conditions (as shown in Sec.5.9) and (2) handle emergencies (as shown in Sec.5.10).
- 5.45 In the tragedy of the digital commons scenario, the hypothesis that internalizers outperform self-interested agents in obtaining cooperation is tested with respect to the number of the successful transactions achieved, as shown in Sec.5.7. In a population of 50 agents, varying the relative numbers of EMIL-I-As (from 0 to 50) and IUMAs (from 0 to 50), it is possible to observe how the percentage of successful transactions changes accordingly.
- 5.46 The effect of norm holders on the norm internalization process is tested by varying the number of initial norm holders (from 0 to 50) and observing how this change affects norms' saliency, i.e., one of the conditions required for norm internalization to take place.
- 5.47 Finally, the capability of EMIL-I-As to adapt to a dynamically changing environment is tested in a scenario in which both the services capacity (from infinite to 3) and the decided-quality (for low to high and vice versa) of the services vary during the simulation (see Fig.4(e)).


Experiment 1: Adapting to Environmental Conditions

- 5.48 One key feature of our internalizers is their ability to dynamically adapt to unpredictable changes. To test this ability, we designed several dynamic situations in which the providers' capacities and the requesters' desires vary. From the beginning of the simulation, the environment is fully populated with EMIL-I-As that already have the two highly salient norms stored in their minds; they start as norm holders.^[16]


[Treatment 1: Infinite services' capacity and expected low quality of the services (meaning that agents need to have the service without paying attention to its quality)]




[Treatment 2: Restricted services' capacity (restricted to 3 clients) and expected low quality of the services]



[Treatment 3: Restricted services' capacity (restricted to 3 clients) and expected high quality of the services]



[Treatment 4: Services' capacity is variable along the simulation: the services' capacity is infinite during the first 2000 time-steps, and after that moment, it linearly decreases from 20 to 1, reaching the service exploitation point (where they cannot longer offer high quality services) after time-step 8000. Agents are assigned tasks to be fulfilled at high quality level along the entire simulation.]



[Treatment 5 (The Complex Loop): Before the time-step 2000, the capacity of the services is infinite, but it is set to 3 after that. Before time-step 5000 and after the time-step 15000, agents are assigned needs that are satisfied also with low quality services. During the rest of the execution, the expected quality of the services is high.]




Figure 4. Different Combinations of Resources Capacities Distribution and Expected Quality Distributions. On the x-axis the time-steps of the simulation are shown, on the y-axis the amount of EMIL-I-As with each internalized norm is indicated.

- 5.49 First, in Fig. 4a we show the dynamics of a basic scenario, where the services are endowed with *infinite* capacity to attend clients, and the requesters do not care about the quality of the services they receive (meaning that also services with a low quality satisfy their needs). This type of experimental setting represents an hypothetical world with infinite services and low maintenance users.
- 5.50 We observe that after a number of time-steps, the internalization process starts occurring, leading to an increased number of agents that internalize norm 1 (Always Share When Asked For a Service). Why is norm 1 spreading and being internalized rather than norm 2 (Share Only High-Quality Services)? This happens because in this experimental condition, agents looking for a service are not interested in its quality thus even though they receive low quality services their needs are satisfied and they will not punish/sanction the service provider. Sharing a service, whatever its value, is interpreted as an action compliant with norm 1 and since instances of this action exceed the number of actions in which high quality services are shared, the salience of this norm is higher than the other's. This experiment confirms that our internalization architecture works in selecting the correct (by experimental design) norm to internalize.
- 5.51 The second treatment is a slight but important variant of the previous one. This time, the services' capacity is restricted to a maximum 3 clients. Results in Fig.4b show that agents correctly internalize the norm that rules the situation (i.e., norm 1), although the internalization process is slower than in the previous treatment. The reason for this delay lies in the dynamics taking place in this treatment: no clear information is given to the agents for deciding which norm, norm 1 or norm 2, is governing the system. EMIL-I-As are initialized with both norms at high salience. When the service providers start sharing, they provide requesters with high-quality services, even though the requesters' needs would also be satisfied by low-quality services. Agents interpret high-quality transactions as actions compliant with norm 2, thus increasing its salience value. After a while, agents start sharing services of lower quality, because the list of requesters exceeds the amount of high-quality services. Low-quality service transactions make the salience level of norm 2 decrease and the salience level of norm 1 increase to its maximum value (a necessary condition for norm internalization as shown in Sec.4.1.4). We can observe how all these micro-dynamics in the individuals affect the global behavior of the system producing the delay shown in Fig. 4b.
- 5.52 In the third treatment (see Fig.4c), we test the agents' capacity to internalize norm 2: service capacity is restricted to 3 clients and the requesters' needs are satisfied only when receiving high-quality services. In general, we observe that the agents that internalize norm 2 do not reach a majority. When they do, however, they reach a majority quickly. This is because until the number of services is high enough to satisfy all the requests with their expected quality, being compliant with norm 2 can be interpreted also as an action compliant with norm 1, thus maintaining the salience of norm 1 at a high level.^[17]
- 5.53 In the fourth treatment (see Fig.4d), a dynamic change in the environment is included. The system is programmed in such a way that for the first 2000 time-steps the capacity to provide services is infinite; after that time, during the simulation, the capacity linearly decreases from 20 to 1. Agents' needs are satisfied only when they receive high-quality services. We observe that after time-step 8000, the capacity of the services significantly decreases and agents slowly start to de-internalize norm 1 and substitute it with norm 2. This treatment shows how our agents are able to adapt to this dynamic situations.
- 5.54 In order to test speed of their adaptation, we designed a final treatment. In Fig.4e, we show the results of the treatment named *The Complex Loop*, in which we can observe that EMIL-I-As perform efficiently even in a complex situation, where not only the environment (the services capacity), but also the agents' preferences (desired quality of the services) change during the simulation.
- 5.55 The *Complex Loop* treatment shows that when the services' capacity is infinite, meaning that high-quality services are always available, and the agents' needs are satisfied by high-quality services, norm 1 is internalized. At time-step 2000, an abrupt change occurs in the environment, making the services' capacity drop to 3, and agents switch to internalizing norm 2. After time step 5000, agents' preferences change (starting now to be satisfied also with low-quality services), producing another change in the internalization dynamics. Agents return to internalizing norm 1. Finally, agents' preferences change again at time step 15000, preferring high-quality services and internalizing norm 2 again.
- 5.56 We can conclude that a population of EMIL-I-As can adapt to sudden and unexpected environmental changes in a flexible manner, selecting the correct norm to internalize.

Experiment 2: EMIL-I-As vs IUAMs

- 5.57 In this experiment, we compare the performance of EMIL-I-As and IUAMs using a number of simulations and treatments, each time with 50 agents. In our treatments, we vary the proportion of EMIL-I-As and IUAMs in intervals of 10.^[18]




Figure 5. Percentage of successful transactions with different proportions of IUMAs and EMIL-I-As.

- 5.58 Figure 5 shows the percentage of successful transactions^{19]} that occur in the system during an identical situation to the one presented in the "Complex Loop" experiment, except, here, the population is a combination of EMIL-I-As and IUMAs.
- 5.59 The results show that the number of successful transactions is proportional to the number of EMIL-I-As. Internalization proves effective in promoting cooperation even in complex situations, i.e., when both the environment (the service capacity) and agents' preferences (desired-quality of the services) change. The explanation of this result can be found in the different motivations behind the decisions of IUMAs and EMIL-I-As to comply with norms. When IUMAs share a service, their payoff is reduced, as sharing prevents them from taking other actions for a number of time-steps, and therefore they decide to free-ride. In other words, they always ask for services, but they never share.
- 5.60 On the other hand, EMIL-I-As decide to comply or violate the norm, not only to maximize their material payoffs, but also because they have recognized the existence of a norm ruling the scenario and then they generate the consequent normative goal. Besides mere self-interest, the normative goal provides EMIL-I-A agents with an additional reason to comply with the norm. Consequently, when the environment changes and agents need to find out the new norm, EMIL-I-As perform a certain number of unsuccessful transactions before properly managing the new situation.
- 5.61 The results of this experiment show that norms are important for solving a tragedy of the commons type situation. Moreover, they also illustrate that to be effective the process of norm internalization needs to be flexible enough for agents to adapt to highly unpredictable and rapidly changing environments.

Experiment 3: Effect of Initial Norm Holders

- 5.62 This experiment aims to analyze how the initial number of norm holders in a group fully populated by EMIL-I-A agents affects the performance of that group. As explained in Sec. 4.1, EMIL-I-As do not know which norms are in force in the group at the onset of the simulation, therefore the presence of norm-holders who explicitly communicate the norms is necessary for triggering the process of norm recognition.




Figure 6. Complex Loop Experiment Average Saliency. On the x-axis the simulation time is shown, while on the y-axis the number of initial norm holders (randomly located in the network) is indicated

- 5.63 Focusing again on *The Complex Loop* experiment (presented in Sec. 5.6 and shown in Fig. 4e), in Fig. 6 the average norm's saliency per agent is shown. We obtained similar results in the rest of experiments performed in Sec. 5.6, but we decided to focus on this situation because of its complexity and completeness (it contains all the possible dynamics of our scenario with respect to norm internalization and de-internalization). Figure 6a represents the saliency dynamics of norm 1 and Figure 6b that of norm 2. At any time-step in the experiment, only one of the two norms prescribes how to ideally behave. We can observe that the saliency of both norms changes according to the changes in the environment, obtaining a higher saliency in the norm that currently applies to the environmental conditions.
- 5.64 In the first 2000 time-steps of the simulation, the number of norm holders has a substantial effect on norm saliency. As the number of norm holders increases, so does the norm's saliency. Special attention should be paid to the situation with no initial norm-holders; here norms are never recognized as no explicit norm elicitation occurs (by explicit norm invocation or by sanction). Even a small number of initial norm-holders (10% of the population in this case) allows agents to recognize the norms and internalize them. Once recognized and stored according to their degree of saliency, agents will start complying with, and defending, the norms. Thus, norm holders are necessary for triggering a virtuous circle from compliance with the norms to their enforcement. After the two norms have been recognized (around time-step 5000), the number of initial norm holders no longer affects the successive dynamics. After time-step 15000 (and depending on the treatment conditions), the service's capacities and agents' needs change again, also leading to changes in the saliency of norms.

Experiment 4: Testing Locality: Norm Coexistence

- 5.65 Given that each agent's norm saliency is updated using information obtained from their social network, network topology and locality, the latter determined by agent's needs and services conditions, affect norm saliency. Consequently, different dilemmas could appear in different parts of the social network, depending (again) on the population's preferences and environmental conditions. We find that EMIL-I-As are able to cope with the locality of the norms, allowing the coexistence of different competitive norms in the same social network. To observe this result we performed the following experiment.

Figure 7. Different norms coexisting in the same social environment.

5.66 We placed 60 agents on a one-dimensional lattice, with neighborhood size set to 6 (i.e., each agent has a constant number of 6 neighbors). In Fig. 7, agents in the top right and in the bottom left quarters of the ring are assigned with infinite capacity services and low-quality fasted services and the rest of the agents are given limited capacity services and high-quality desired tasks. The color of the nodes indicates the internalized norms (light color standing for norm 1, dark color for norm 2, and white stands for norms not internalized). The self-adaptive capability of EMIL-I-As shows a good performance in the designed environment: two norms can coexist in different areas of the same social network.

Experiment 5: Dealing with Emergencies: Selective Norm Violation

- 5.67 One of our claims is that EMIL-I-As are able to unblock a norm not only in situations where its salience is very low, meaning that the norm is disappearing, but also in emergency situations.
- 5.68 While IUMAs cannot handle normative requests with different levels of urgent event (because they are not endowed with normative architectures that can keep track of norms' salience, and compare them with explicit urgent requests), EMIL-I-As can (thanks to the salience module). In the ambulance case, the siren announces an emergency that is more salient than the observance of the norm *stop at the red traffic light*. In our distributed web-service scenario, we can imagine the following situation. One of our agents is writing a scientific paper that includes a number of important calculations. Unfortunately, the day before the deadline, the agent realizes it needs to rerun some simulation experiments. In order to obtain the results in time, it needs to use calculation clusters that it does not have access to directly, and instead, are distributed amongst its peers. It makes requests to its peers accordingly and, in response, the requested agents can choose to follow a "first-in, first-out" norm or, given the urgent situation, make an exception and execute the calculations of our last-minute scientist. It is easy to see how the violation of the first-in, first-out norm by its colleague-neighbors, benefits our agent.
- 5.69 Experimental results show that the EMIL-I-As successfully handle emergencies when these are explicitly specified (as the siren in the ambulance case). However, there is a cost associated to this advantage. An emergency is interpreted by the audience as a non-punished norm-violation, leading to a reduction of the norm's salience (stop at the red traffic light). If agents face the same emergency too often, the salience of the norm "normally" regulating the situation will decrease and this reduction will unblock the internalization process, leading the agent back to the normative (benefit-to-cost calculation) phase. We can then affirm that emergencies can be managed by internalizers, but only when they occur sporadically, otherwise they are interpreted as a change in the salience of the norm governing the environment.

Summary and conclusions

6.1 In this work, we have presented EMIL-I-A, an agent architecture for internalizing norms that is designed for adaptive systems. EMIL-I-A allows for a balanced trade-off between calculation time and efficiency. The norm salience mechanism plays a key role in the internalization process. It works as the conductor orchestrating the process of norm internalization in scenarios where more than one norm applies. In other words, norm salience allows norm compliance to synchronize with fluctuating environmental conditions and agents' tasks, thus guaranteeing a successful adaptation.

Advances on the state of the art

- 6.2 What is the value added of the work presented so far with regard to the current understanding of the process of norm internalization? First, we presented an *operational* model of norm internalization, which can be reproduced and computationally tested. Second, we presented a *dynamic* model of norm internalization as a multistep process, demonstrating the relationships between norm internalization and related phenomena, such as (semi)automated behaviors and routines or habits. Third, we showed by means of *simulation* to what extent and under what conditions norm internalization positively affects cooperation in a group of agents.
- 6.3 We tested EMIL-I-As in a tragedy of the digital commons scenario, where the emergence of cooperation is a difficult task to achieve and they were able to successfully deal with this task. The simulation data show that when facing situations in which the environment can rapidly change, internalizers are able to adapt. We also observe that the amount of initial norm holders and EMIL-I-As speeds up the process of norm convergence, even when the scenarios are dynamic and different norms have to be internalized and de-internalized.
- 6.4 A fourth interesting result achieved concerns the subjective character of norm salience, which allows the coexistence of different norms in the same social environment: depending on the agents' interests and the environmental conditions, different norms can emerge in the same environment.
- 6.5 Last, the experiments have shown that our internalization architecture, EMIL-I-A, is flexible enough to handle emergencies and decide to violate an already internalized norm. We claim that internalization provides agents with self-policing capabilities that are very useful in settings where social control is unfeasible and expensive.

Future Work

- 7.1 As Axelrod Axelrod (1986) already envisioned in his work, internalization is a key mechanism for norm maintenance. In his work, he explored the *metanorm* mechanism in a simulated scenario, leaving seven other mechanisms (dominance, internalization, deterrence, social proof, membership, law and reputation) experimentally unexplored.
- 7.2 The experiments presented in this work do not allow us to claim that the internalization process makes agents save calculation and execution time. We need to test this hypothesis in a different scenario. Nonetheless, our intuition is that the automated behavior obtained via the internalization process results in a faster execution and decision-making time with respect to other classical deliberative architectures.
- 7.3 As one of the short term tasks for future work, a new scenario in which the norms do not subsume each other needs to be found. In such a scenario, we would be able to obtain clearer results about how the adaptation process is improved thanks to the salience module and the internalization process.
- 7.4 The experimental results presented in this work show that internalization performs satisfactorily in self-regulated web service provision. So far, however, we have designed and run experiments only with agents that all desire the same level of quality for the requested services. In future versions of this work, we will address heterogeneous situations (where agent requests are different, and therefore different norms apply), thus tackling *dominance*, another mechanism pointed out by Axelrod (1986). A less immediate but no less valuable direction for future work is the modeling of internalization as a continuum from deliberate action to automatic response, in which interesting combinations of both types of behavior might be explored and compared in terms of both individual and social utility.

Acknowledgements

This work was supported by the Spanish Education and Science Ministry [AEI project TIN2006-15662-C02-01, AT project CONSOLIDER CSD2007-0022, INGENIO 2010; Proyecto CBIT [TIN2010-16306]; Proyecto Intramural de Frontera MacNorms [PIFCOO-08-00017]; SINTELNET FET-OPEN Coordinated Action; the Generalitat de Catalunya [2009-SGR-1434]; the European University Institute; the Complexity-Net Project SEMIRA; the EC Seven Framework Programme project GLODERS; the EC Sixth Framework Programme and European project COST Action IC0801 Agreement Technologies. We also thank the CESGA and Rede Galega de Bioinformatica for the technical support.

Notes

¹We refer to *goals* as internal representations triggering-and-guiding action at once: they represent the state of the world that agents want to reach by means of action and that they monitor while executing the action. In particular, a normative goal is a goal that an agent happens to have as long as she has a normative belief about *a*. In other words, an agent has a normative goal only if she believes herself to be subject to a norm (Conte & Castelfranchi 1995; Conte 2009; Miller et al. 1960).

²When agents recognize a social input as a norm, its initial salience is fixed to what we refer to as the *Norm Activation Value*. Then, this starting value is dynamically updated step by step depending on the normative information gathered from the social environment. As it is shown in Sec. 4.1.2, the salience value of any norm is within the range [0, 1]. Therefore, the specific Norm Activation Value within the simulation is 0.5, the middle point between the minimum and the maximum salience value.

³The resulting salience measure (Sal ? [0,1], 0 representing the minimum salience and 1 the maximum salience) depends on the social information that each agent gathers from the environment.

⁴These values have also been checked through a cross-methodological study including human subjects, see Andrighetto et al. (2013).

⁵This value is generated by calculating the maximum and minimum value of the weighted normative cues.

⁶We are aware that associating *urgent events* with a pre-defined salience value is a short-cut and agents should learn them instead. How agents learn what has to be considered as urgent will be explored in future work.

⁷These conditions equally apply to all agents, however, each agent fulfills them at different moments, because norm salience is subjective and updated with local information (C,O,NPD,P,S,E), which might be different for each agent.

⁸The main intuition for the incorporation of this transaction time is to study the effects of homophily (Val et al. 2012) in our distributed web service scenario. However, we decided to postpone that

research and in this work to focus only on the effects of norm internalization.

⁹The purpose of this specification is to cover the necessities of this example while keeping it as simple as possible. We are aware of the extensive literature about norm specification but this is out of the scope of this article.

¹⁰This amount of time-steps has been chosen according to the experimenter criteria in order to show the effects of the internalization. Because of agent's decision making and simulation design, these time-steps do not have any correlation with real world phenomena timings.

¹¹In order to reduce the space of the parameters, the *ServiceAssignmentProbability* and the *needAssignmentProbability* have been fixed to 0.1. In future versions of this work, real data will be extracted from available data sets in order to obtain a more realistic simulation model.

¹²This system works similarly to the BitTorrent (Zohar & Rosenschein 2009) where the central server maintains an updated list of service providers.

¹³As the communication of norms amongst agents is out of the scope of the present research, we have assumed perfect communication. This implies that when one agent decides to sanction another, the normative message that accompanies the monetary punishment is equally understood by both parts as the explicit communication of a certain norm without possibility of misunderstanding.

¹⁴Without the Saliency Control Module, that provides saliency measures for different norms, the urgent event management module would not properly work as there would exist no way to compare the importance of an urgent event with respect to the norms.

¹⁵As we said in Section 5.2, only agents that have recognized that there is a norm regulating their group will sanction, otherwise they will just use punishment.

¹⁶As we will see in Sec. 5.8, the internalization process speed is proportional to the initial amount of norm holders.

¹⁷We understand that in this scenario there is a relationship between the compliance of Norm 2 and Norm 1 which might difficult the analysis of the experimental results, although this exemplifying scenarios should be use for the understanding of the dynamics of norm internalization and de-internalization.

¹⁸As in the previous one, in this scenario, EML-I-As are also norm-holders.

¹⁹We consider transactions as successful when a service provider offers a service of the requested quality.



References

- ADAR, E. & Huberman, B. A. (2000). Free riding on gnutella. *First Monday*. [doi:10.5210/fm.v5i10.792]
- ALBERT, R. & Barabasi, A.-L. (2002). Statistical mechanics of complex networks. *Review of Modern Physics* 74(1), 47–97. [doi:10.1103/RevModPhys.74.47]
- ANDRIGHETTO, G., Brandts, J., Conte, R., Sabater, J., Solaz, H. & Villatoro, D. (2013). Punish and voice: Punishment enhances cooperation when combined with norm-signalling. *PLoS ONE* 8(6), e64941. [doi:10.1371/journal.pone.0064941]
- ANDRIGHETTO, G., Campenni, M., Cecconi, F. & Conte, R. (2010a). The complex loop of norm emergence: A simulation model. In: *Simulating Interacting Agents and Social Phenomena* (Deguchi, H. & et al., eds.), vol. 7 of *Agent-Based Social Systems*. Springer Japan, pp. 19–35.
- ANDRIGHETTO, G. & Castelfranchi, C. (2013). Norm compliance. the prescriptive power of normative actions. *Paradigma* 2.
- ANDRIGHETTO, G., Villatoro, D. & Conte, R. (2010b). Norm internalization in artificial societies. *AI Communications* 23, 325-339.
- ANDROUTSELLIS-THEOTOKIS, S. & Spinellis, D. (2004). A survey of peer-to-peer content distribution technologies. *ACM Comput. Surv.* 36, 335–371. [doi:10.1145/1041680.1041681]
- AXELROD, R. (1986). An evolutionary approach to norms. *American Political Science Review* 80(4), 1095–1111. [doi:10.2307/1960858]
- BICCHIERI, C. (1990). Norms of Cooperation. *Ethics* 100(4), 838–861. [doi:10.1086/293237]
- BICCHIERI, C. (2006). *The Grammar of Society: The nature and Dynamics of Social Norms* Cambridge University Press.
- BLANC, A., Liu, Y.-K. & Vahdat, A. (2005). Designing incentives for peer-to-peer routing. In: *INFOCOM*. [doi:10.1109/infcom.2005.1497907]
- CAMPENNI, M., Andrighetto, G., Cecconi, F. & Conte, R. (2009). Normal = normative? the role of intelligent agents in norm innovation. *Mind & Society* 8, 153–172. [doi:10.1007/s11299-009-0063-4]
- CHUDEK, M. & Henrich, J. (2011). Culture-gene coevolution, norm-psychology and the emergence of human prosociality. *Trends in cognitive sciences* 15(5), 218–226. [doi:10.1016/j.tics.2011.03.003]
- CIALDINI, R. B., Reno, R. R. & Kallgren, C. A. (1990). 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. [doi:10.1037/0022-3514.58.6.1015]
- CONTE, R. (2008). Review of the book generative social science: Studies in agent-based computational modeling. *Jasss - the Journal of Artificial Societies and Social Simulation* 10(4).
- CONTE, R. (2009). Rational, goal governed agents. In: *Springer Encyclopedia of Complexity and System Science*. (Meyers, R. A., ed.), Springer. [doi:10.1007/978-0-387-30440-3_445]
- CONTE, R., Andrighetto, G. & Campenni, M. (2013). *Minding Norms: Mechanisms and Dynamics of Social Order in Agent Societies*. Oxford University Press. [doi:10.1093/acprof:oso/9780199812677.001.0001]
- CONTE, R. & Castelfranchi, C. (1995). *Artificial Societies: The Computer Simulation of Social Life*, chap. Understanding the functions of norms in social groups through simulation. UCL Press, pp. 74–118.
- CONTE, R. & Castelfranchi, C. (2006). The mental path of norms. *Ratio Juris* 19(4), 501–517. [doi:10.1111/j.1467-9337.2006.00342.x]
- CRIADO, N., Argente, E. & Botti, V. (2010). Normative Deliberation in Graded BDI Agents. In: *Eighth German Conference on Multi-Agent System Technologies (MATES-10)*, vol. 6251 of *LNAI*. Springer. [doi:10.1007/978-3-642-16178-0_7]
- CURBERA, F., Duffler, M., Khalaf, R., Nagy, W., Mukhi, N. & Weerawarana, S. (2002). Unraveling the web services web: An introduction to soap, wsdl, and uddi. *IEEE Internet Computing* 6, 86–93. [doi:10.1109/4236.991449]
- DE PINNINCK, A. P., Sierra, C. & Schorlemmer, M. (2007). Friends no more: norm enforcement in multiagent systems. In: *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems, AAMAS '07*. New York, NY, USA: ACM. [doi:10.1145/1329125.1329238]
- DECKER, K., Sycara, K. P. & Williamson, M. (1997). Middle-agents for the internet. In: *JCAI (1)*.
- EPSTEIN, J. (2006). *Generative Social Science. Studies in Agent-Based Computational Modeling*. Princeton-New York: Princeton University Press.
- FAILLO, M., Grieco, D. & Zarrì, L. (2012). Legitimate punishment, feedback, and the enforcement of cooperation. *Games and Economic Behavior*
- GALBIATI, R. & D'Antoni, M. (2007). A signalling theory of non monetary sanctions. *International Review of Law and Economics* 27, 204–218. [doi:10.1016/j.irl.2007.06.008]
- GIARDINI, F., Andrighetto, G. & Conte, R. (2010). A cognitive model of punishment. In: *Proceedings of the 32nd Annual Conference of the Cognitive Science Society Austin, TX: Cognitive Science Society* (Ohlsson, S. & Catrambone, R., eds.). Portland, Oregon.
- GINTIS, H. (2004a). The genetic side of gene-culture coevolution: internalization of norms and prosocial emotions. *Journal of Economic Behavior and Organization* 53, 57–67. [doi:10.1016/S0167-2681(03)00104-5]
- GINTIS, H. (2004b). The genetic side of gene-culture coevolution: internalization of norms and prosocial emotions. *Journal of economic behavior & organization* 53(1), 57–67. [doi:10.1016/S0167-2681(03)00104-5]
- HASSIN, R. R., Bargh, J. A. & Zimmerman, S. (2009). Automatic and flexible: the case of nonconscious goal pursuit. *Social Cognition* 27(1), 20–36. [doi:10.1521/soco.2009.27.1.20]
- HOUSER, D. & Xiao, E. (2010). Understanding context effects. *Journal of Economic Behavior & Organization* 73(1), 58–61. <http://ideas.repec.org/a/eee/jeborg/v73y2010i1p58-61.html> [doi:10.1016/j.jebo.2008.01.006]

- KENNEDY, W. G. & Bugajska, M. (2007). Integrating fast and slow cognitive processes. In: *Proceedings of the International Conference on Cognitive Modeling (ICCM 2010)* (Salvucci, D. & Gunzelmann, G., eds.). Philadelphia.
- KENNEDY, W. G. & Trafton, J. G. (2007). Long-term symbolic learning. *Cognitive Systems Research* 8(3), 237–247. [doi:10.1016/j.cogsys.2007.06.005]
- KINGSLEY, D. (1949). *Human Society*. New York: The Macmillan Co.
- KOHLBERG, L. (1981). *The Philosophy of Moral Development: Moral Stages and the Idea of Justice (Essays on Moral Development, Volume 1)* Harper & Row, 1st ed.
- MASCLET, D., Noussair, C., Tucker, S. & Villeval, M.-C. (2003). Monetary and nonmonetary punishment in the voluntary contributions mechanism. *American Economic Review* 93(1), 366–380. [doi:10.1257/000282803321455359]
- MILLER, G., Galanter, E. & Pribram, K. H. (1960). *Plans and the Structure of Behavior*. New York: Henry Holt. [doi:10.1037/10039-000]
- MUKHERJEE, P., Sen, S. & Airiau, S. (2007). Norm emergence in spatially constrained interactions. In: *Proceedings of ALAg-07*. Honolulu, Hawaii, USA.
- NAPSTER (2006). Napster. <http://www.napster.com/>.
- NEWMAN, M. E. J. (2003). The structure and function of complex networks. *SIAM REVIEW* 45, 167–256. [doi:10.1137/S003614450342480]
- NIKIFORAKIS, N. & Normann, H.-T. (2008). A comparative statics analysis of punishment in public-good experiments. *Experimental Economics* 11(4), 358–369. [doi:10.1007/s10683-007-9171-3]
- PIAGET, J. (1965). *Moral Judgement of the Child* Free Press.
- REYKOWSKY, J. (1982). Development of prosocial motivation: A dialectic process. In: *The Development of Prosocial Behavior* (Eisenberg, N., ed.). New York: Academic Press. [doi:10.1016/B978-0-12-234980-5.50020-1]
- SEN, S. & Airiau, S. (2007). Emergence of norms through social learning. In: *Proceedings of the 20th International Joint Conference on Artificial Intelligence IJCAI'07*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- SUNSTEIN, C. R. (1996). Social norms and social roles. *Columbia Law Review* 96(4), 903–968. [doi:10.2307/1123430]
- TOBIAS, R. (2009). Changing behavior by memory aids: a social psychological model of prospective memory and habit development tested with dynamic field data. *Psychological Review* 116, 408–438. [doi:10.1037/a0015512]
- ULLMAN-MARGALIT, E. (1977). *The Emergence of Norms*. Clarendon Press, Oxford.
- VAL, E. D., Rebollo, M. & Botti, V. (2012). Self-Organized Service Management in Social Systems. In: *45 Hawaii International Conference On System Sciences*.
- VILLATORO, D., Sen, S. & Sabater-Mir, J. (2009). Topology and memory effect on convention emergence. In: *Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology – Volume 02, WI-IAT '09*. Washington, DC, USA: IEEE Computer Society. [doi:10.1109/WI-IAT.2009.155]
- WATKINS, C. & Dayan, P. (1992). Q-learning. *Machine Learning* 8(3–4), 279–292. [doi:10.1007/BF00992698]
- ZOHAR, A. & Rosenschein, J. S. (2009). Adding incentives to file-sharing systems. In: *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems – Volume 2, AAMAS '09*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.