

Noe: Norms Emergence and Robustness Based on Emotions in Multiagent Systems

Sz-Ting Tzeng¹, Nirav Ajmeri², and Munindar P. Singh¹

¹ North Carolina State University, Raleigh NC 27695, USA
{stzeng,mpsingh}@ncsu.edu

² University of Bristol, Bristol, BS8 1UB, UK
nirav.ajmeri@bristol.ac.uk

Abstract. Social norms characterize collective and acceptable group conducts in human society. Furthermore, some social norms emerge from interactions of agents or humans. To achieve agent autonomy and make norm satisfaction explainable, we include emotions into the normative reasoning process, which evaluate whether to comply or violate a norm. Specifically, before selecting an action to execute, an agent observes the environment and infer the state and consequences with its internal states after norm satisfaction or violation of a social norm. Both norm satisfaction and violation provoke further emotions, and the subsequent emotions affect norm enforcement. This paper investigates how modeling emotions affect the emergence and robustness of social norms via social simulation experiments. We find that an ability in agents to consider emotional responses to the outcomes of norm satisfaction and violation (1) promote norm compliance; and (2) improve societal welfare.

1 Introduction

Humans, in daily life, face many choices at many moments, and each selection brings positive and negative payoffs. In psychology, decision-making [26] is a cognitive process that selects a belief or a series of actions based on values, preferences, and beliefs to achieve specific goals. The final choice could be rational or irrational in terms of utility theory since emotions often influence decision-making [25]. Emotions, the responses to internal or external events or objects, can involve the decision-making process and provide extra information in communication [16,25]. Social norms regulate behaviors in a human society [27,23], but humans and agents have to deviate from norms in certain contexts. For instance, people do handshaking at the normal time. Yet, people have to deviate from the social norm during a pandemic.

An agent that models the emotions of its users and other humans can potentially behave in a more realistic and trustworthy manner. The decision-making process for humans or agents involves evaluating possible consequences of available actions and choosing the action that maximizes the expected utility [11]. Without considering emotions or other affective characteristics, such as personality or mood, some compliance seems irrational [4]. Humans' compliance does show hints on rational planning over their objectives [16]. Including emotion or personality in normative reasoning makes

these compliance behaviors explainable. Norms either are defined in a top-down manner or emerge in a bottom-up manner [23]. Works on norms include norm emergence based on the prior outcome of norms, automated run-time revision of sanctions [10], or considering various aspects during reasoning [2,1]. However, sanctions in the real world are more subtle instead of harsh punishment. For instance, the sanctions could be trust update or emotional expression and might change one's behavior [21,5]. Kalia et al. [15] considered norm outcome with respect to emotions and trust and goals. Modeling and reasoning about emotions and other affective characteristics in an agent then become important in decision making and would help the agent enforce and internalize norms, which boost agent autonomy.

At the basic level, emotions help individuals understand others' preferences, beliefs, and intentions and therefore coordinate social interactions. Furthermore, emotions serve as motivations or deterrents for others' social behavior and therefore play an essential role in learning. At the cultural level, emotions facilitate cultural identities and help individuals learn the norms and the values of their culture [16]. Based on this understanding of emotions and norms, we propose an agent architecture that integrates the BDI architecture [7] with a normative model [4,27] and emotional model [3,17].

Accordingly, we investigate the following research question.

RQ_{emotion}. How does modeling the emotional responses of agents to the outcomes of interactions affect the norm emergence and social welfare in an agent society?

To address RQ_{emotion}, we refine the abstract normative emotional agent architecture [4] and investigate the interplay of norms and emotions. We propose a framework *Noe* based on BDI architecture [7], norm life-cycle [23,12,4], and emotion life-cycle [3,17]. To evaluate *Noe*, we design a simulation experiment with various agent societies. We investigate how norms emerge and how emotions in normative agents influence social welfare.

To make the problem tractable, we apply one social norm in our simulation and simplify the emotional expression to reduce the complexity. Specifically, our *Noe* agents process emotions by appraising norm outcomes. For the emotion model, we adopt the OCC model of emotions [22] in which we consider both valence and intensity of emotions and assume violation of norms yields negative emotions.

Organization. The rest of the paper is structured as follows. Section 2 discusses the relevant related works. Section 3 describes *Noe*, including the symbolic representation and the decision-making in *Noe*. Section 4 details the simulation experiments we conduct to evaluate *Noe* and describes the experimental results. Section 5 presents the conclusions and the future directions.

2 Related Works

Savarimuthu and Cranefield [23] proposed a life-cycles model for norms and discussed varied mechanisms of norm study. Broersen et al. [8] introduced the so-called Beliefs-Obligations-Intentions-Desires (BOID) architecture on top of the Beliefs-Intentions-Desires (BDI) architecture [7], which further include obligation and conflict resolution.

Ortony et al. [22] model emotions based on events, action, and objects. Marsella and Gratch [17] proposed a computational model of emotion to model appraisal in perceptual, cognitive, and behavioral processes. Moerland et al. [20] surveyed on emotions with reinforcement learning. Bourgaïs et al. [6] reviewed different approaches that modelers used to integrate emotions into social simulations. Keltner and Haidt [16] differentiate the functional approaches and research of emotions by four-level analysis: individual, dyadic, groups, and cultural. Briefly, emotions provide some information for agents or people to coordinate social interactions. We take inspiration from these works.

Argente et al. [4] propose an abstract normative emotional agent architecture, an extension of Belief-Desire-Intention (BDI) architecture, which combines emotional, normative, and BDI component. Argente et al. claim that different tasks require agents with different characteristics. As for mood, a positive mood may lead to forgiveness of an insult. Argente et al. defined four types of relationships between emotions and norms: (1) emotion in the process of normative reasoning, (2) emotion generation with norm satisfaction or violation, (3) emotions as a way to enforce norms, (4) anticipation of emotions promotes internalization and compliance of social norms. Yet, Argente et al. do not validate the interplay between emotions and norms with their proposed architecture.

Bourgaïs et al. [5] present an agent architecture that integrates cognition, emotions, emotion contagion, personality, norms, and social relations to simulate humans and ensure explainable behaviors. In the proposed architecture, agents work in four steps: (i) agents perceive the environment and link the environment and its knowledge; (ii) agent updates its cognitive mental state, generates emotions based on its knowledge, and updates the social relation with others; (iii) agents make decisions via a cognitive engine and a normative engine; (iv) the architecture generates a temporal dynamic by degrading the cognitive mental states and emotion's intensity and updates norm status. However, emotions serve as activation of norms but do not enforce social norms in this work.

Von Scheve et al. [24] allow emotion generation with norm satisfaction or violation. Specifically, an observer agent perceives the transgression of a norm of another, its strong negative emotions (e.g., contempt, disdain, detestation, or disgust) constitute negative sanctioning of the violator. The negative sanctioning then leads to negative emotions (e.g., shame, guilt, or embarrassment) in the violator. Besides, compliance with the social norms can stem from the fear of emotional-driven sanctions, which would lead to negative emotions in the violator. Such fear enforces social norms. Yet, emotions are not part of the decision-making process in this work.

3 Noe

We now describe the architecture, norm formal model, and decision-making.

3.1 Architecture

Noe integrates the BDI architecture [7] with a normative model [23,12,4] and an emotional model [3,17]. A Noe agent assesses the environment, including other agents'

explicit emotions, its cognitive mental states, and infer possible outcomes to make a decision. Figure 1 shows the three components of *Noe*.




Fig. 1. *Noe* Architecture: rectangles represent processes; circles represent data instances; arrows represent information flows.

The normative component of *Noe* includes the following processes:

- Identification: the agent recognize norms from its norm base based on its beliefs
- Instantiation: activate norms related to the agent
- Normative reasoning process: the reasoning process makes decisions based on the beliefs, current intention, self-emotions, other-emotions, active norms, and how the norm satisfaction or violation influences the world and itself. The *Noe* agents then update the intention based on the results of normative reasoning.

- Norm fulfillment process: check if a norm has been fulfilled or violated based on the selected action.

The BDI component includes the following parts:

- Beliefs: update beliefs based on perceptions.
 - Desires: generate desires based on the beliefs
 - Intention: the highest priority of desires to achieve based on the beliefs
 - Action: select action based on the current intention, emotions, possible outcomes, and the evaluation of violating or complying with norms, if any
- The beliefs, desires, and intentions are mental states of Noe agents.

The emotional component includes the following processes:

- Appraisal: calculate the appraisal value based on the beliefs, desires, and the elicited emotion triggered by norm satisfaction or violation of a norm. In this work, we consider only the elicited emotion.
- Emotion: generate emotion based on the appraisal value [17]

We consider both valence and intensity of emotions and assume violation of norms yields negative emotions.

Figure 2 illustrates the interactions between agents in our simulation scenario.

3.2 Norm formal model

Social norms describe the interactions between agents in a sociotechnical system. We adopt the representation in Singh [27] where a social norm can be formalized as $Norm(subject, object, antecedent, consequence)$ where the subject and object are agents. This representation describes a norm activated by the subject towards the object when the antecedent holds, and the consequent indicates if the norm was satisfied or violated).

Following Singh [27], we consider three types of norms in Noe.

- Commitment (C): the subject commits to the object to bring out the consequence if the antecedent holds. Consider Alice and Bob are queuing up in a grocery store. Alice and Bob commit to each other to keep social distance during the pandemic, represented as $C(Alice, Bob, during = pandemic, social_distance)$.
- Prohibition (P): the object prohibits the subject from the consequence if the antecedent holds. Caleb, the grocery store manager, prohibits Bob from jumping the queue while lining up in that store, represented as $P(Bob, Caleb, when = line-up; at = grocery_store, stay_in_queue)$.
- Sanction (S): same as commitment or prohibition, yet the consequence would be the sanctions. Sanctions could be positive, negative, or neutral reactions to any norm satisfaction or violation [21]. If Bob breaks the queue, he receives negative sanctions from Alice, represented as $S(Bob, Alice, jump, negative_sanctions)$. Negative sanctions could be physical actions, e.g., scold, or emotional expression, e.g., expressions of disdain, scowl, or disgust.

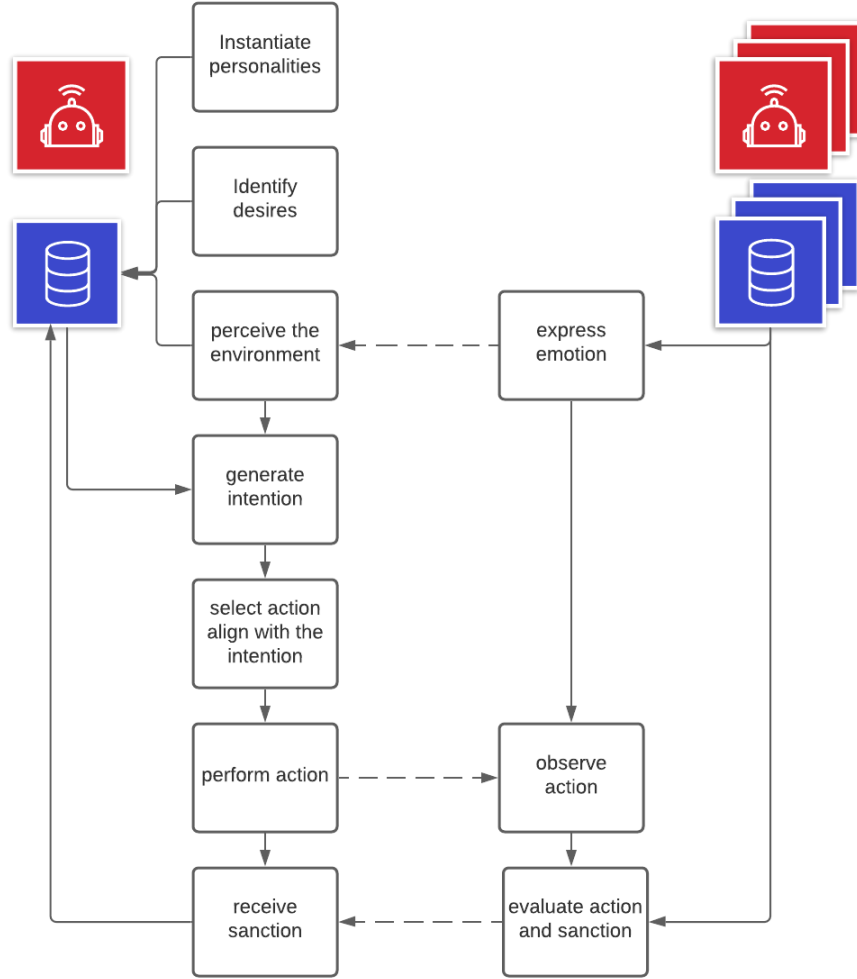


Fig. 2. The interaction between Noe agents.

To simulate the norm emergence and enforcement in human society, we include emotions into the decision-making process since, by nature, humans do not always act rationally in terms of utility theory. Here we formalize emotions with $E_i(target, intensity, decay)$ indicating agent a_i has emotion e toward the target with intensity and decay value. A target in this representation can be an object or an agent. An example of the prohibition case would be, Bob would not jump the queue if Alice is angry, represented as $P(Bob, Alice, Bob \succ Alice \wedge E_{Alice} = angry, stay)$.

We model the emotional response of agents with triggered emotions from norm satisfaction or violation [4]. Here we represent the elicited emotions with

$Elem_{name}(A_{expect}, A_{real}, Em_1, Em_2) | Em_1, Em_2 \in E; A_{expect}, A_{real} \in A$ where A is a set of actions. E is a set of emotions, and Em_1 and Em_2 are the emotions triggered by norm satisfaction and violation accordingly. If the A_{expect} is equal to the A_{real} , a norm has been fulfilled, and Em_1 was elicited. $Ap(beliefs, desires, Elem)$ represents the appraisal function.

3.3 Decision-Making

Schwarz [25] addresses the influence of moods and emotions at decision making and discussed the interplay of emotion, cognition, and decision making. Specifically, the aspects include pre-decision affect, post-decision affect, anticipated affect, and memories of past affect. In our model, we include the pre-decision affect into the decision-making process. Specifically, people recall information from memories that match their current affect [25]. This recalled information affects people’s decision-making process. Moreover, people in a sad mood particularly tend to overestimate adverse outcomes and events.

In our model, emotions serve as mental objects, which we represent with simple values where positive values indicate positive emotions and larger values indicate stronger emotions. We simplify from the OCC model [22] and consider certain emotions in this work, i.e., fear, shame, reproach, and admiration. Specifically, we express emotions with numerical values where positive values indicate positive emotions and the value of emotions indicates its intensity. We ignore mood since mood is hard to detect. We only consider short-term emotions. Noe agents’ appraisal function considers norm satisfaction only. The agents are aware of other agents’ expressed emotions in the same place, their own desires and beliefs, and available actions to achieve their goals. In this work, we assume that agents express true and honest emotions and perceive expressed emotions. In other words, felt emotions are equal to expressed emotions. Another assumption is that emotions are consistent with the notions of rational behavior.

Algorithm 1 displays the decision loop of our model. At the beginning of the simulation, all agents are initialized with certain desires, and during the run, an intention would be generated by prioritizing desires with the agent’s beliefs. When choosing the next move with line 5 in Algorithm 1, the agent chooses the one with maximum utility from all available actions. Algorithm 2 details the action selection. The decision takes the agent’s beliefs, current intention, and possible consequences into accounts. While norms are activated with the beliefs, the agent would further consider emotions and cost and possible consequences with norms at line 9 in Algorithm 2. For instance, if people violate some social norms, they may be isolated from society. Regarding the influence of emotions, people may overestimate the negative outcomes when they are in the negative emotion and tend to comply with the norms.

4 Evaluation

We evaluate Noe via a line-up environment where agents form queues to receive service. We detail the environment in Section 4.1.

Algorithm 1: Decision loop of a *Noe* agent

```

1 /* initialize one agent with its desires D */
2 while True do
3   observe the environment and update beliefs B;
4   generate intention I based on B and D;
5   a = ActionSelection(B, I, D, E);
6   execute the selected action a;
7   if action a fulfills a norm then
8     elicit positive or neutral emotions E from agent itself and observer agents;
9     # other-directed emotions
10    E as sanctions to others;
11  else
12    elicit negative or neutral emotions E from agent itself and observer agents;
13    # other-directed emotions
14    E as sanctions to others;
15  end
16  # self-directed emotions
17  E as sanctions to itself;
18 end

```

Algorithm 2: Action selection

```

1 /* choose one action that maximizes utility */
   Input: beliefs, intention, desires, and emotions
   Output: Action a
2 Function Action Selection:
3   for available actions do
4     N = activated norms with beliefs;
5     if  $N = \emptyset$  then
6       selected_action = max_utility(beliefs, intention, possible_action)
7     else
8       for N do
9         # possible result on self and others
10        possible_result = norm_reasoning(N, beliefs, intention, possible_action)
11        * amplifier(emotions)
12        selected_action = max_utility(possible_result)
13      end
14    end
15  return selected_action
16 return

```

4.1 Line-up Environment

Figure 3 shows the line-up environment. We build this line-up environment using Mesa [18], a Python-based framework for building, analyzing, and visualizing agent-based models.




Fig. 3. Simulation details. The left part: accessible locations; The right part: a queue for agents to line up to enter the grocery stores. Only N agents get services at one time.

The line-up environment includes two shared locations—home and grocery stores. The agents move between home and grocery stores to get food. We consider one social norm in the line-up environment — to enter the grocery store, agents are expected to line-up. To simulate real human reactions to norm violations, we refer to a social psychology experiment [19]. In the line-up environment, we model defensive reactions of people in the queue as negative emotions toward those who jump the queue by barging in ahead of someone already in the queue. People show positive emotions toward those who stay in the queue.

We initialize the agents with the following parameter values:

- Health (Integer value from 0–100): When the health value reaches zero, the agent is marked as *deceased* and unable to act. The health value decreases by 1 unit at each step.
- Deceased (Boolean: True or False): the agent’s state; marked as True when an agent runs out of health.
- Emotion (Integer value): simplified with numerical values where positive values indicate positive emotions and negative indicate negative emotions. The emotions come along with a duration. Default at 0.
- Number of food packets owned (Integer value from 0–15): once obtained food from the stores, agents would be able to restore its health value via consuming food anywhere.

- Food expiration day (Integer value from 0–15): once the agent gets food packets, we update the expiration day with 15. The expiration day decreases by 1 unit at each step. Food expires once the expiration day reaches 0. Default at 0.
- Beliefs: the perceived and processed information from the world, including location, other agents' emotions.
- Desires: desired states, including *have food* and *wandering*.
- Intention: the highest priority of desires to achieve at a specific time. When the agent's health is lower than the threshold, 80% of the health, the agent sets its intention as *get food*; otherwise, the agent sets its intention as *wandering*.

When an agent runs low on stock, it has a higher probability of moving to a grocery store. The grocery store can provide food packets to eight agents in one time step. While waiting in line to get food, the agent could either stay in the line or jump ahead in the line to get food with less time. Jumping the line may increase other agents' delay in getting food packets. Those who witness the violation would then cast negative emotions, further interpreted as anger or disdain, triggered by that behavior. To simplify the simulation, we presume the anticipated affects [25] with: (1) receiving negative emotions triggers negative self-directed emotions such as shame and guilt; (2) complying with norms leads to positive or neutral emotions; (3) violating norms leads to negative or neutral emotions. The intensity of emotions triggered each time is fixed but can accumulate. Each triggered emotion lasts 2 steps. At each step, the duration and intensity of emotion decrease by 1 as decay. The values of emotions can add up. A simple assumption here is that people in a bad mood would trigger stronger emotions in response to a non-ideal state. Note that at the beginning of the simulation, we initialize the agent society with health in normal distribution to avoid all agents having the same intention at the same time.

4.2 Agent Types

To answer our research question and evaluate *Noe*, in addition to *Noe*, we define three types of agent societies as baselines. We describe the agents societies below:

Obedient society. Agents in an obedient society always follow norms.

Anarchy society. Agents in an anarchy society jump lines when they cannot get food.

Sanctioning society. Agents in the sanctioning society jump lines considering the previous experience of satisfying or violating a norm. Agents sanction positively or negatively based on norm satisfaction or violations directly and comply with enforced norms.

Noe society. Agents in the *Noe* society jump lines considering the previous experiences of satisfying or violating a norm, current emotional state of the other agents, current self emotional state, and estimated outcome of satisfying or violating a norm. *Noe* agents who observe norm satisfaction or violations would appraise the norm outcomes and trigger emotions to sanction the actor agent.

Table 1 summarizes the characteristics of the agents in the four societies.

Table 1. Characteristics of the various agent societies.

| Agent Type | Violation allowed | Sanctioning | Emotions involved |
|---------------------|-------------------|-------------|-------------------|
| Obedient society | ✗ | ✗ | ✗ |
| Anarchy society | ✓ | ✗ | ✗ |
| Sanctioning society | ✓ | ✓ | ✗ |
| Noe society | ✓ | ✓ | ✓ |

4.3 Hypotheses and Metrics

To address our research question RQ_{emotion} on emotions and norm emergence, we propose two hypotheses:

H₁ (Norm satisfaction): Norm satisfaction in *Noe* agent society is higher compared to the baseline agent societies.

H₂ (Social welfare): *Noe* agent society yields a better social welfare compared to the baseline agent societies.

H₃ (Social experience): *Noe* agent society yields a better social experience compared to the baseline agent societies.

To evaluate H₁ on norm satisfaction, we compute one metric:

M₁ Cohesion: Percentage of norm satisfaction

To evaluate H₂ on social welfare, we compute three metrics:

M₂ Cumulative number of agents deceased

To evaluate H₃ on social experience, we compute three metrics:

M₃ Average waiting time of agents in the queues

M₄ Average health of the agents

To test the statistical significance of H₁, H₂, and H₃, we conduct the independent t-test and measure effect size with Glass's Δ for unrelated societies [14,13]. We adopt Cohen's [9] descriptors to interpret effect size where above 0.2, 0.5, 0.8 indicate small, medium, and large.

4.4 Experimental Setup

We run each simulation with 400 agents and queue size 80 for 3,000 steps. To reduce the simulation time, we choose a relatively small number of agents. Our results are stable for a larger number of agents. The simulation stabilizes at about 1,500 steps of our simulation, but we keep extended simulation steps to have more promising results.

We present the results with an average of a running window of 100 steps. We choose this size of running window to show the temporal behavior change in a small sequence of time. With a larger size, the running window may alleviate the behavior change. To minimize deviation from coincidence, we run each simulation with 10 iterations and compute the mean values.

4.5 Experimental Results

In this section, we describe the simulation results comparing the three baselines and Noe agents. Table 2 summarizes these results. Table 3 lists the value of Glass’s Δ and p -values from the independent t-test.

According to Table 2 and Table 3, we see that Noe generate better cohesion and social welfare than baselines ($p < 0.01$; Glass’s $\Delta > 0.8$). The null hypotheses corresponding to H_1 and H_2 are rejected. Note that we do not consider the cohesion metric for the obedient agent society here since agents in the obedient society are always compliant. However, Noe also yields the worst social experience where low waiting time and high health are desirable states ($p < 0.01$; Glass’s $\Delta > 0.8$). The null hypothesis corresponding to H_3 indicates no significant difference.

Table 2. Comparing Noe agent society with baseline agent societies on various metrics.

| Agent Society | Cohesion | Deceased | Waiting Time | Health |
|---------------|----------|--------------|--------------|--------|
| Obedient | 1 | 55.30 | 8.95 | 79.27 |
| Anarchy | 0.22 | 81.60 | 5.45 | 79.5 |
| Sanctioning | 0.88 | 169.30 | 2.55 | 86.26 |
| Noe | 0.99 | 54.00 | 8.95 | 79.0 |

Table 3. Statistical analysis.

| Agent Society | Glass’s Δ | | | | p -value | | | |
|---------------|------------------|----------|--------------|--------|------------|----------|--------------|--------|
| | Cohesion | Deceased | Waiting time | Health | Cohesion | Deceased | Waiting time | Health |
| Obedient | 0.19 | 0.65 | 0.01 | 0.18 | 0.32 | < 0.01 | 0.98 | 0.52 |
| Anarchy | −102.43 | 3.1 | −40.82 | 0.21 | < 0.01 | < 0.01 | < 0.01 | 0.46 |
| Sanctioning | −13.67 | 15.53 | −76.68 | 3.34 | < 0.01 | < 0.01 | < 0.01 | 8.45 |
| Noe | − | − | − | − | − | − | − | − |

H_1 Norm Satisfaction Figure 4 displays the cohesion, the percentage of norm satisfaction, in the baseline agent societies and the Noe agent society. We find that the percentage of norm satisfaction in the Noe agent society, average at 99% and p -value < 0.01, is constantly higher than the sanctioning agent society, average at 88% and p -value < 0.01 and Glass’s $\Delta > 0.8$. The sanctioning agent society learns to comply with the norm as time goes by. The Noe agent society does sanction as well. Yet, considering emotions and the possible outcome makes Noe agent society enforce the norm faster than the sanctioning agent society. Specifically, Noe agent society enforces the norm at about 100 steps while sanctioning agent society at 1500 steps.




Fig. 4. Simulation result: average cohesion. Comparing average cohesion (M_1) yielded by Noe and baseline agent societies.

H₂ Social Welfare Figure 5 compares the average number of deceased in the obedient, anarchy, sanctioning, and Noe agent societies. Refer to Figure 4, sanctioning agent society soon learns the norm with positive and negative sanctioning from norm satisfaction and violation. However, the agents in that society do not consider the outcome of norm satisfaction to cause compliant agents to die in the queue. When the number of deceased reaches the threshold, the simulation stabilizes. Therefore, no more agent from the sanctioning agent society dies after the threshold. On the contrary, Noe agent society sanctions and considers possible outcomes of norm satisfaction and violation, therefore learning the norm and avoiding unacceptable consequences.




Fig. 5. Simulation result: average number of deceased. Comparing average number of deceased (M_2) in Noe and baseline agent societies.

H₃ Social Experience Figure 6 compares the average duration the agents spend in a queue at the grocery store in the obedient, anarchy, sanctioning, and Noe agent soci-

eties. The *Noe* agent society learns the norm fast and remains the same waiting time in the queue. However, some agents in the sanctioning agent society take advantage of those who learn norms faster than themselves. Therefore, many agents die during the learning process, and the simulation stabilizes. In Figure 6, the obedient agent society shares the same trend with *Noe* agent society since emotions enforce the line-up norm.




Fig. 6. Simulation result: average waiting time of agents in queues. Comparing average waiting time (M_3) in *Noe* and baseline agent societies.

Figure 7 compares the average health of the agents in the obedient, anarchy, sanctioning, and *Noe* agent societies. The sanctioning agent society performs better, average at 86.26, since most agents die in the learning process. The rest of the agents then be able to remain in high health.




Fig. 7. Simulation result: average health value. Comparing average health value (M_4) in *Noe* and baseline agent societies.

Combining the results for H_1 and H_2 and H_3 , we note that while sanctioning enforces norms, a combination of sanctioning and emotions enforce norms better. Specifically, having emotions as amplifiers of outcomes yield higher norm satisfaction compared to our baselines. The results also indicate that, first, sanctioning agents that consider only norm violation or norm satisfaction may bring out worse social welfare compared to Noe that considers both norms and their consequences. Second, although Noe agents remain relatively high waiting time in the queues, the number of deceased is lower than the baselines. Note that the sudden drop of deceased number or increase of health value for sanctioning agents resulted from the stabilization of that society. Third, Noe agents stay in positive emotions during the simulation while sanctioning agents start from negative emotions and finally achieve the expected behaviors.

5 Discussion and Conclusion

We present an agent architecture inspired by norm life-cycle [4], BDI architecture [7], and emotion life-cycle [3,17] to investigate the interactions of norms and emotions. We evaluate the proposed architecture via simulation experiments in an environment where agents queue up to receive service. In our simulations, we consider two characteristics of an agent society: sanctioning and emotions that participate in action selection and arise from evaluating selected action.

As a future extension of current work, we plan to differentiate emotions in Noe instead of general emotions to provide more information. We also consider including a mix of personality in future research to have different appraisal results. In this work, Noe agents are assumed to express true and honest emotions, yet in an adversarial context, emotions can also serve as tools to influence, persuade or deceive others.

Acknowledgments

Thanks for the NSF for partial support under grant IIS-1908374.

References

1. Ajmeri, N., Guo, H., Murukannaiah, P.K., Singh, M.P.: Robust norm emergence by revealing and reasoning about context: Socially intelligent agents for enhancing privacy. In: Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI). pp. 28–34. IJCAI, Stockholm (Jul 2018). <https://doi.org/10.24963/ijcai.2018/4>
2. Ajmeri, N., Guo, H., Murukannaiah, P.K., Singh, M.P.: Elessar: Ethics in norm-aware agents. In: Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS). pp. 16–24. IFAAMAS, Auckland (May 2020). <https://doi.org/10.5555/3398761.3398769>
3. Alfonso Espinosa, B.: Agents with Affective Traits for Decision-Making in Complex Environments. Ph.D. thesis, Universitat Politècnica de València (2017)
4. Argente, E., Del Val, E., Perez-Garcia, D., Botti, V.: Normative emotional agents: A view-point paper. IEEE Transactions on Affective Computing (2020)

5. Bourgaïs, M., Taillandier, P., Vercouter, L.: BEN: an agent architecture for explainable and expressive behavior in social simulation. In: Calvaresi, D., Najjar, A., Schumacher, M., Främling, K. (eds.) *Explainable, Transparent Autonomous Agents and Multi-Agent Systems - First International Workshop, EXTRAAMAS 2019, Montreal, QC, Canada, May 13-14, 2019, Revised Selected Papers. Lecture Notes in Computer Science*, vol. 11763, pp. 147–163. Springer (2019). https://doi.org/10.1007/978-3-030-30391-4_9, https://doi.org/10.1007/978-3-030-30391-4_9
6. Bourgaïs, M., Taillandier, P., Vercouter, L., Adam, C.: Emotion modeling in social simulation: A survey. *Journal of Artificial Societies and Social Simulation* **21**(2) (2018)
7. Bratman, M., et al.: *Intention, Plans, and Practical Reason*, vol. 10. Harvard University Press Cambridge, MA (1987)
8. Broersen, J.M., Dastani, M., Hulstijn, J., Huang, Z., van der Torre, L.W.N.: The BOID architecture: Conflicts between beliefs, obligations, intentions and desires. In: André, E., Sen, S., Frasson, C., Müller, J.P. (eds.) *Proceedings of the Fifth International Conference on Autonomous Agents, AGENTS 2001, Montreal, Canada, May 28 - June 1, 2001*. pp. 9–16. ACM (2001). <https://doi.org/10.1145/375735.375766>, <https://doi.org/10.1145/375735.375766>
9. Cohen, J.: *Statistical Power Analysis for the Behavioral Sciences*. Lawrence Erlbaum Associates, Hillsdale, New Jersey, 2nd edn. (1988)
10. Dell’Anna, D., Dastani, M., Dalpiaz, F.: Runtime revision of norms and sanctions based on agent preferences. In: Elkind, E., Veloso, M., Agmon, N., Taylor, M.E. (eds.) *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS ’19, Montreal, QC, Canada, May 13-17, 2019*. pp. 1609–1617. International Foundation for Autonomous Agents and Multiagent Systems (2019), <http://dl.acm.org/citation.cfm?id=3331881>
11. Edwards, W.: The theory of decision making. *Psychological bulletin* **51**(4), 380 (1954)
12. Frantz, C., Pigozzi, G.: Modeling norm dynamics in multiagent systems. *FLAP* **5**(2), 491–564 (2018), <https://www.collegepublications.co.uk/downloads/ifcolog00022.pdf>
13. Glass, G.V.: Primary, secondary, and meta-analysis of research. *Educational researcher* **5**(10), 3–8 (1976)
14. Grissom, R.J., Kim, J.J.: *Effect Sizes for Research: Univariate and Multivariate Applications*. Routledge, Abingdon-on-Thames (2012)
15. Kalia, A.K., Ajmeri, N., Chan, K., Cho, J.H., Adali, S., Singh, M.P.: The interplay of emotions and norms in multiagent systems. In: *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*. pp. 371–377. IJCAI, Macau (Aug 2019). <https://doi.org/10.24963/ijcai.2019/53>
16. Keltner, D., Haidt, J.: Social functions of emotions at four levels of analysis. *Cognition & Emotion* **13**(5), 505–521 (1999)
17. Marsella, S.C., Gratch, J.: Ema: A process model of appraisal dynamics. *Cognitive Systems Research* **10**(1), 70–90 (2009)
18. Masad, D., Kazil, J.: MESA: An agent-based modeling framework. In: *Proceedings of the 14th PYTHON in Science Conference*. pp. 53–60 (2015)
19. Milgram, S., Liberty, H.J., Toledo, R., Wackenhut, J.: Response to intrusion into waiting lines. *Journal of Personality and Social Psychology* **51**(4), 683 (1986)
20. Moerland, T.M., Broekens, J., Jonker, C.M.: Emotion in reinforcement learning agents and robots: A survey. *Machine Learning* **107**(2), 443–480 (2018)
21. Nardin, L.G., Balke-Visser, T., Ajmeri, N., Kalia, A.K., Sichman, J.S., Singh, M.P.: Classifying sanctions and designing a conceptual sanctioning process model for socio-technical systems. *The Knowledge Engineering Review* **31**(2), 142 (2016)
22. Ortony, A., Clore, G.L., Collins, A.: *The Cognitive Structure of Emotions*. Cambridge University Press, New York (1988)

23. Savarimuthu, B.T.R., Cranefield, S.: Norm creation, spreading and emergence: A survey of simulation models of norms in multi-agent systems. *Multiagent and Grid Systems* **7**(1), 21–54 (2011)
24. von Scheve, C., Moldt, D., Fix, J., von Luede, R.: My agents love to conform: Norms and emotion in the micro-macro link. *Computational & Mathematical Organization Theory* **12**(2-3), 81–100 (2006)
25. Schwarz, N.: Emotion, cognition, and decision making. *Cognition & Emotion* **14**(4), 433–440 (2000)
26. Simon, H.A.: *The New Science of Management Decision*. Harper & Brothers (1960)
27. Singh, M.P.: Norms as a basis for governing sociotechnical systems. *ACM Transactions on Intelligent Systems and Technology (TIST)* **5**(1), 21:1–21:23 (Dec 2013). <https://doi.org/10.1145/2542182.2542203>