

Generalising Axelrod’s metanorm game through the use of explicit domain-specific norms

Abira Sengupta¹, Stephen Cranefield¹, and Jeremy Pitt²

¹ University of Otago, Dunedin, New Zealand

² Imperial College London

abira.sengupta@postgrad.otago.ac.nz, stephen.cranefield@otago.ac.nz,
j.pitt@imperial.ac.uk

Abstract. Achieving social order in societies of self-interested autonomous agents is a difficult problem due to lack of trust in the actions of others and the temptation to seek rewards at the expense of others. In human society, social norms play a strong role in fostering cooperative behaviour—as long as the value of cooperation and the cost of defection are understood by a large proportion of society. Prior work has shown the importance of both norms and metanorms (requiring punishment of defection) to produce and maintain norm-compliant behaviour in a society, e.g. as in Axelrod’s approach of learning of individual behavioural characteristics of boldness and vengefulness. However, much of this work (including Axelrod’s) uses simplified simulation scenarios in which norms are implicit in the code or are represented as simple bit strings, which limits the practical application of these methods for agents that interact across a range of real-world scenarios with complex norms. This work presents a generalisation of Axelrod’s approach in which norms are explicitly represented and agents can choose their actions after performing *what-if* reasoning using a version of the event calculus that tracks the creation, fulfilment and violation of expectations. This approach allows agents to continually learn and apply their boldness and vengefulness parameters across multiple scenarios with differing norms. The approach is illustrated using Axelrod’s scenario as well as a social dilemma from the behavioural game theory literature.

Keywords: Norms, Expectation Event Calculus, Metanorms game

1 INTRODUCTION

The conflict between social benefit and an individual’s self-interest is a central challenge in all social relationships as individuals may put their own interests ahead of those of the society as a whole, often leading to a suboptimal outcome for all—a situation known as a *social dilemma*.

Understanding how societies can solve this conflict and achieve cooperation toward the collective good is essential. In a fishing scenario, for example, it is profitable for a fisherman to catch as many fish as possible, but if everyone is selfishly doing the same thing, the fishery will eventually run out of fish.

The question of why people often do cooperate with others remains, despite the fact that individuals may benefit more from defecting than from cooperating. The classical game theory literature from the last few decades models social dilemmas using payoff matrices or trees and solution concepts such as the Nash equilibrium, while making the assumption that all agents are perfectly rational and well-informed. However, this becomes intractable to reason about for a large number of agents

Human society suggests that cooperation can occur due to internal and social motivations such as altruism, rational expectations (e.g. focal points), social choice mechanisms (e.g. voting and bargaining) and social norms[8].

One body of prior work has focused on the role of norm, providing evidence that social norms play an important role in fostering cooperation. Social norms imply that members of society should comply with prescribed behaviour while avoiding proscribed behaviours [14]. Bicchieri explains why agents adhere to social norms, claiming that a social norm emerges as a result of our expectations of others and beliefs about their expectations [3]. Axelrod’s use of the norm and metanorm game illustrates how agents adopt normative behaviour after learning individual parameters of boldness and vengefulness, where their boldness represents their propensity to violate norms and vengefulness represents their inclination to punish others for violating norms [2]. However, his evolutionary study is based on an implicit representation of norms, which limits its practical use for agents that interact in a variety of real-world situations with a range of different norms.

In this work, we propose a generalisation of Axelrod’s method where norms are represented explicitly and agents can choose their course of action after engaging in *what-if* reasoning to compare the normative outcomes of alternative actions. This approach is significant because it enables agents to continuously learn and apply their boldness and vengefulness parameters across a variety of scenarios with various norms.

The following is the structure of the paper. Axelrod’s norm and metanorm game-based scenario is discussed in Section 2. Section 3 emphasises the use of explicit norms to encode Axelrod’s mechanism. Section 4 depicts the results of generalisation of Axelrod’s norm and metanorm games, as well as the use of boldness and vengefulness in other scenarios. The prior event calculus models of norms are described in Section 5. Section 6 discusses the conclusion and future work.

2 Background of Axelrod’s model

Axelrod’s states that a norm is determined by “how often the action is taken and just how often someone is punished for not taking it”. To represent norms, [2] develops a game in which players learn the parameters of bravery and vengefulness over generations of the population and can choose to deviate from the norms and metanorms, receiving punishment for their violations. The goal of the *norms* investigation is to see how *cooperation* emerges from norms.

To explain the players' strategies Axelrod uses an evolutionary approach which selects the strategies that are working well for players, while the strategies that are not working well are discarded. The difference between game theory and evolutionary approaches is that in game theory, players use strategies to determine payoffs based on their own and others' choices. In the case of the evolutionary approach, no rational calculation is required to determine the best strategy. New strategies are viewed as mutated versions of old ones in this

2.1 The norms game

Axelrod's norm game follows an evolutionary model in which successful agent strategies propagate over generations (Figure 1(a)). A strategy is a pair of values representing the agent's boldness and vengefulness. Each agent has the option of defecting by violating a norm, and there is a chance of being observed by other agents with the probability S , which is drawn individually for each agent from a uniform distribution ranging . Each of the agents has two decisions to make.

- Agents must decide whether to defect or cooperate based on their boldness value (B). A defecting agent (when $S < B$) is one who receives a Temptation payoff ($T = 3$) while other agents receive a Hurt payoff ($H = -1$). If agents decide to cooperate, no one's payoff will change.
- If the agent observes others defecting (as determined by the S value), the agent must decide whether to punish those defectors based on its vengefulness (a probability of punishment). Punishers incur an enforcement cost ($E = -2$) every time they punish ($p = -9$) a defector.

Axelrod's simulated the norm game five times with 100 generations each with twenty agents³. Following the utilities of each agent are used to evolve the population of agents. Agents with scores greater than to the average population score plus one standard deviation are reproduced twice in a new generation. Agents with a score less than the average population score minus one standard deviation are never reproduced in a new generation. Other agents are only reproduced once⁴. The initial values of B and V are chosen at random from a uniform distribution of eight values ranging from $0/7$ to $7/7$, implemented as a 3 bit string. During reproduction each bit has a 1% chance of being flipped as a mutation.

2.2 The Metanorm game

Axelrod found that norms alone were not sufficient to sustain norm compliance in the society. Axelrod therefore introduced the concept of a meta-norm, is the

³ Axelrod uses five runs for a hundred generations to simulate the norm and metanorm game. However, we follow the recommendation of [6] and use 100 runs.

⁴ Axelrod does not state how he maintains a fixed population size after applying these reproduction rules. We follow the approach of [6] involving random sampling when the new population is too large, and random replication when the population is too small.

encouragement for enforcing a norm based on punishment (Figure 1(b)). The metanorm game includes punishment for those agents who do not punish defectors after observing them defect. The chance of observing an agent who did not punish observed defectors by each of the other 18 agents (excluding the defectors). Metapunishers receive Meta-enforcement cost (-2) every time they meta-punish (-9).

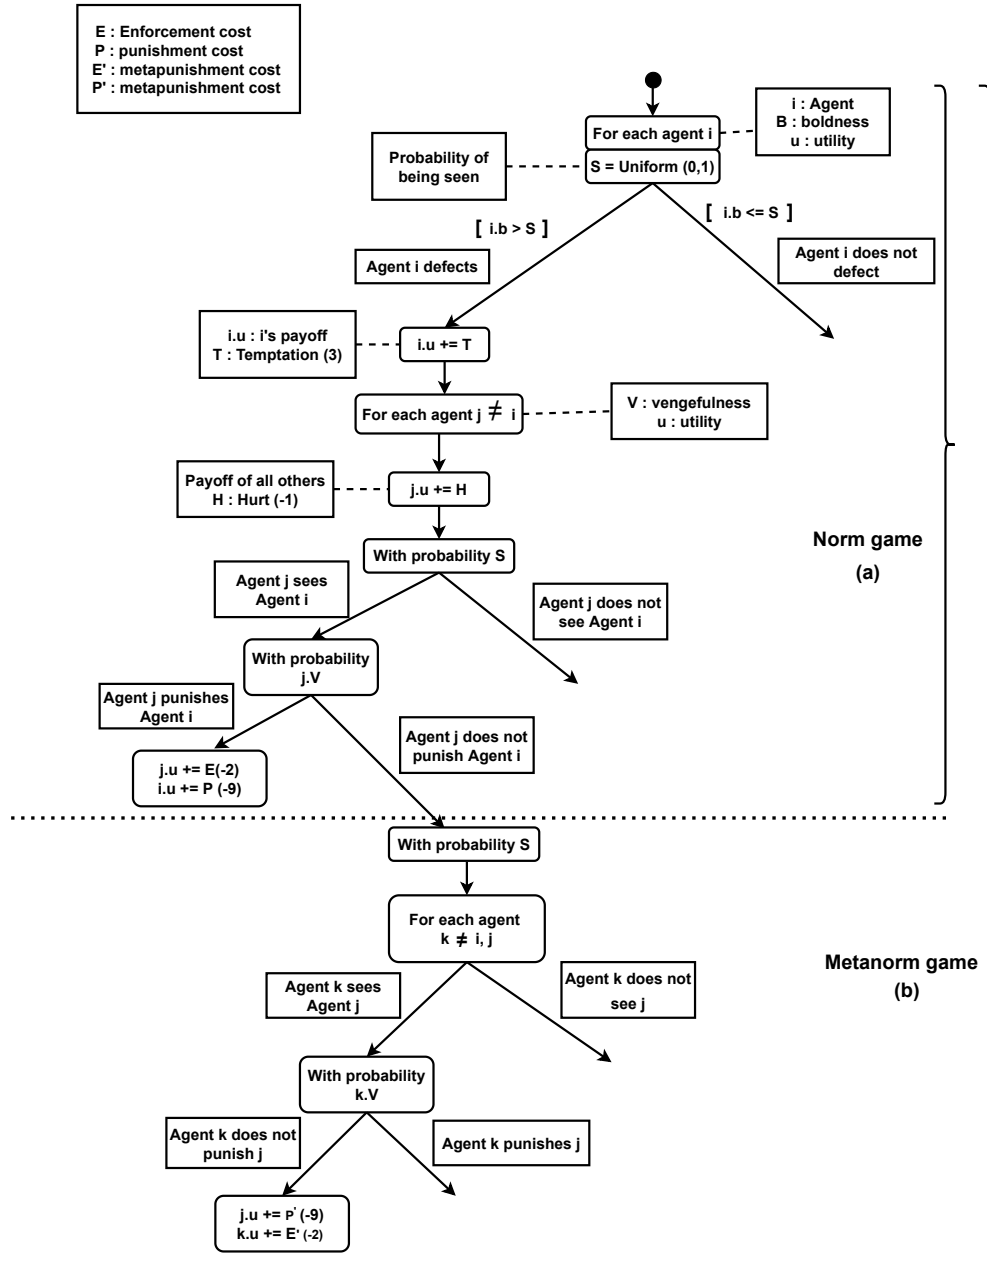


Fig. 1: (a) Axelrod's evolutionary study's norm game says if agent i is bold enough, defect; otherwise, cooperate. If agent j is vengeful enough and sees defection, punish agent i else do not punish. (b) **The metanorm game** adds each observing agent k . Agent k punishes agent j if agent j does not punish agent i (defector), k sees this, and agent k is vengeful; otherwise, agent k does not punish.

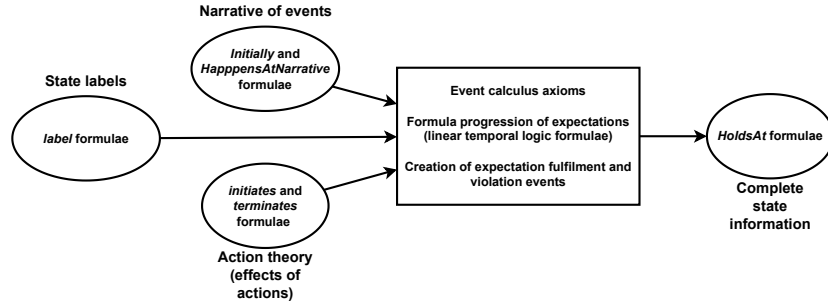


Fig. 2: Overview of reasoning in the Expectation Event Calculus

3 ENCODING AXELROD’S MECHANISM USING EXPLICIT NORMS

To model and generalise Axelrod’s mechanism, we provide an explicit representation of norms and a mechanism that can compare alternative actions to determine which will lead to a norm violation. The expectation event calculus (EEC), a discrete event calculus extension, provides this capability.

3.1 The Expectation Event Calculus

The event calculus (EC) consists of a set of predicates that are used to encode information about the occurrence of events and dynamic properties of the state of the world (known as fluents), as well as a set of axioms that relate these predicates [13]. This logical language supports various types of reasoning. In this work, we use it for *temporal projection*. This takes as input a narrative of events that are known to occur (expressed using $happensAt(E, T)$, where E is an event and T is a time point) and a domain-specific set of clauses defining the conditions under which events will *initiate* and *terminate* fluents (expressed using the predicates $initiates(E, F, T)$ and $terminates(E, F, T)$). The EC axioms are then used to infer what fluents hold at each time point. By default, fluents are assumed to have *inertia*, i.e. they hold until explicitly terminated by an event.

The EC, in general, assumes that time is dense, and time points are ordered using explicit $<$ constraints. In this work, we use the *discrete event calculus* (DEC), which assumes that time points are discrete and identified by integers [10].

The expectation event calculus (EEC) [4] is an extension of the DEC that includes the concepts of expectation, fulfilment, and violation. Expectations are constraints on the future, expressed in a form of linear temporal logic, that the agent wishes to monitor. Expectations are free from inertia and instead are automatically *progressed* from one state to the next, which means they are partially evaluated and re-expressed in terms of the next time point. During progression, if they evaluate to true or false, a fulfilment or violation is generated.

Figure 2 illustrates temporal projection in the EEC. In addition to the standard features of the DEC, there are two special kinds of fluents: *exp_rule* and *exp*. A conditional rule to create expectations is expressed by an *exp_rule(Cond, Exp)* fluent. Here, *Cond* is a condition on the past and/or present, while *Exp* represents the future expectation. *Exp* will be expected if *Cond* holds, in which case an *exp(Exp)* fluent is created. In our implementation of the EEC, the condition can test for fluents holding, the occurrence of events (expressed using *happ(E)*), the presence of a symbolic label *L* in a state (using the expression *@L*) and complex expressions involving conjunctions and linear temporal logic operators such as *next*, *eventually*, *always* and *never* can be used. Labels are associated with time points using the *label(L, T)* declarations, and are not required to be unique. To distinguish between basic events in the narrative and the inferred fulfilment and violation events, we use the predicate *happensAtNarrative* to declare the narrative events.

In contrast to the approach of [4], we represent fulfilments and violations as events rather than fluents, denoted *fulf(Cond, Exp, T, Res)* and *viol(Cond, Exp, T, Res)*, where *Cond* and *Exp* are the condition and expectation of an expectation rule that was triggered at time *T* to create the expectation, and *Res* is the residual expectation (after being progressed zero or more times since its creation) at the time of its fulfilment or violation⁵.

Our EEC implementation includes a *what-if* predicate that accepts two alternative event lists (*E*₁ and *E*₂) as arguments and infers the fluents that would hold and the events (including violation and fulfilment events) that would occur if each list of events were performed at the current time point. It returns the fluents and events that would occur if the events in *E*₁ are performed but not those in *E*₂, and those that would occur if the events in *E*₂ occur but not those in *E*₁. This can be used as a basic form of look-ahead to assist an agent in deciding between two actions, and especially in this work, to compare which (if any) of two events will cause expectation violations.

3.2 Modelling Axelrod’s scenario with the EEC

We model time as a repeated cycle of steps and associate an EEC label with each step. We use the event calculus *initiates* and *terminates* clauses to define the effects of events that update an agent’s *S* value, give payoff to an agent as the outcome of all agents’ cooperate or defect actions, and punish and metapunish agents.

We use the EEC within a simulation platform [5] that integrates Repast Symphony [11] with the EEC through queries to SWI Prolog. This includes an institutional model in which agents take on roles by asserting to the EEC narrative that certain institution-related events have occurred such as joining an institu-

⁵ There is also an extended version of the *exp* fluent with these four arguments—the version used in this paper has only the residual expectation as its argument.

tion and adding a role⁶. Each role has an associated set of conditional rules. When the rule engine is run at the start of a simulation step (Figure 3), these rules recommend the actions that are relevant to the agent’s current role based on queries to the EEC, e. g. to check the current step’s label and the fluents that currently hold. Then the agent can run scenario-specific code to select one of the actions to perform.⁷

In contrast to Axelrod’s implicit representation of a norm and metanorm, our explicit representation of a norm implies that three norms are at task. In the metanorm game, each action choice is governed by a norm. As there are three choice-points, there are three norms that are modelled using *exp_rule* fluents.

First Order Norm

```
initially(
  exp_rule(member(A, society),
    never(happ(defect(A))))).
```

This *initially* clause creates an expectation rule (*exp_rule*) expressing the *first-order norm*, which states that no defection should occur for any agent who is a member of the society.

As the *first-order norm* described above is likely to be insufficient to motivate selfish agents to follow the norm and cooperate with others, a *second-order norm* is required.

Second Order Norm

```
initially(
  exp_rule(and([sawViolation(B,A,R,_),
    pl(contains_term(
      defect(A),R))] ),
    happ(punish(B,A)))).
```

The *second-order norm* requirement states that if the *first-order norm* is violated by an agent, and another agent who observes the violation should punish the *first-order norm* defector.

This norm (*second-order norm*) is triggered by a *sawViolation* fluent, which is created when a violation of the *first-order norm* occurs and defectors observed. The following *initiates* clause creates this fluent.

```
initiates(viol(_,_ ,ResidualExp),
  sawViolation(B,A,ResidualExp,T),
  T :-
  responsible(ResidualExp, A),
  agent(B),
```

⁶ <https://github.com/maxant/rules>

⁷ At present we assume there are no more than two relevant actions.


```

B \== A,
holdsAt(s(A,S), T),
random(R),
R < S.

```

The condition for this clause first determines which agent is responsible for the unfulfilled expectation, then generates a possible observer different from the violator and compares the agent's S value with a random number to determine whether or not the violation has been observed.

In our application, the violated expectation will include an instantiation of one of the action terms `defect(A)`, `punish(A)` or `metapunish(A)`, and we can use these to identify the responsible agent. We therefore define the `responsible` predicate in Prolog as follows.

```

responsible(Expectation, A) :-
    contains_term(defect(A),
                  Expectation).

responsible(Expectation, A) :-
    contains_term(punish(A,_),
                  Expectation).

responsible(Expectation, A) :-
    contains_term(metapunish(A,_),
                  Expectation).

```

When there are *second-order norm* violators, a new norm, *third-order norm*, is required.

Third Order Norm

```

initially(
    exp_rule(and([sawViolation(B,A,R,_),
                  pl(contains_term(
                      punish(A,_),R))] ,
                  happ(metapunish(B,A))))).

```

According to this EEC rule, observer agents must punish the violators of the *second-order norm* when the violating agents fail to punish the *first-order norm* defector after observing their defection. The `pl` term in the rule's condition contains a goal to be evaluated using Prolog.

Figure 3(a) and (b) illustrate the differences between our implementation of the metanorm game with implicit and explicit norms. Figure 3(a) makes hard-coded action choices, following Axelrod's algorithm. However, in Figure 3(b), whenever there is an action choice to be made, the two action choices are compared using *what-if* reasoning that is informed by one of the three norms. If one of the choices will cause violation

The EEC label “Defect or cooperate” indicates a timestep in which an agent must decide whether to defect or cooperate based on whether they have a high boldness value or not (Figure 3(a)). Figure 3(b) shows the first-order norm has been created by its `exp_rule` in the initial time step and so is active already. After learning the boldness parameter, agents can use the *what-if* mechanism to compare the outcomes of two alternative actions, cooperation or defection. If the agent’s boldness value exceeds S , the agent will defect or violate the first order norm; otherwise, the agent will cooperate.

Implicit norm simulation at the “punishment” label, demonstrates how agents are punishing each observed defector with the probability given by the vengefulness parameter (Figure 3(a)). Whereas the explicit norm representation simulation cycle demonstrates *second-order norm* implementation within the “punishment” step and how the agent can use the *what-if* mechanism to punish or do not punish (Figure 3(b)).

If an agent violates with high boldness at the “punishment” label of the implicit norm simulation, then an agent with high vengefulness meta-punishes the violator (Figure 3(a)). However, in our work, we can see how the EEC implements the *third-order norm* at “meta_punishment” label and agents use the *what-if* mechanism to decide whether or not to meta-punish the defectors (Figure 3(b)).

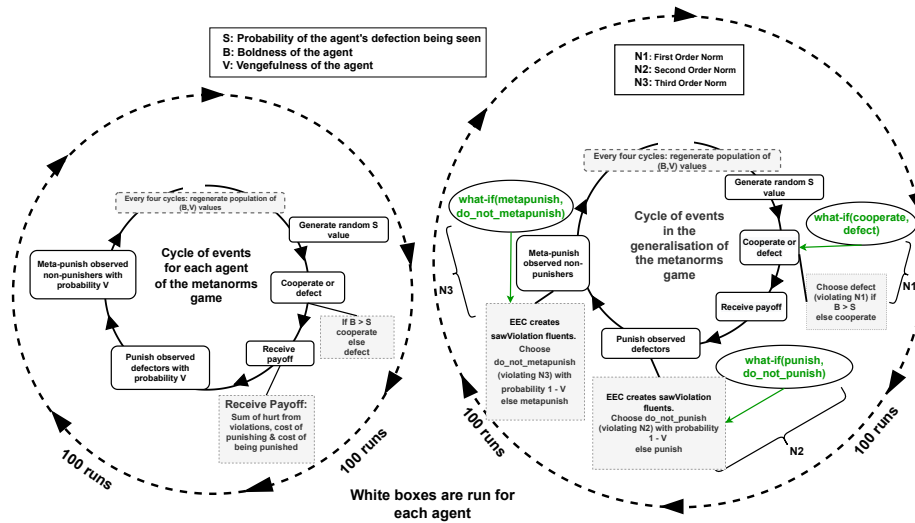


Fig. 3: The distinction between implicit and explicit metanorms. It depicts the cycle of events for Axelrod’s metanorms game (a) in its original form with implicit norms, and (b) in our generalised form with explicit norms.

Table 1: Roles and their possible actions

Label	Role	Possible Action
cooperate or defect	Temptation Role	Cooperate or Defect
punishment	Possible punisher Role	Punish or Do not punish
metapunishment	Possible punisher Role	metaPunish or Do not meta-punish

Figure 4 depicts in more detail our use of explicit norms with agents that are aware of norm violations. We have three norms represented rules in different times they triggered and created expectations. The EEC *initially* clauses generate the *exp_rule* fluents for the *first-order norm* (N_1), *second-order norm* (N_2), and *third-order norm* (N_3). Each agent has two roles: *Temptation role* and *Possible punisher role*. Table 1 shows what actions an agent can take in the simulation when assigned to a specific role for each label. Agents with the *Temptation role* can choose to cooperate or defect at the “cooperate or defect” label. While at “punishment” label an agent with *possible punisher role* can choose to punish or do not punish, and an agent can choose to metapunish or do not metapunish at the “metapunishment” label with same “Possible punisher role”. At time tick = 1 of the simulation, both roles (*Temptation role* and *Possible punisher role*) are activated.

The EEC *what-if* predicate is used to consider two options: defect or cooperate, punish or do not punish, meta-punish or do not meta-punish (depending on the current step in the simulation cycle), and determine whether one rule produces a violation while the other does not. The non-violating option is then chosen (or a random choice if there is no violation). If both options result in a violation, the cost of each violation is calculated (using domain-specific knowledge) and the less expensive option is chosen. If the costs are the same, a random selection is made.

The role *Temptation role* contains rules to suggest the actions (Table 1) when the probability of a defection being seen by someone (S value) is less than agent’s boldness (B). Where, the role, *possible punisher*, contains different rules at different labels (i. e. “punishment” and “meta_punishment”) to suggest the corresponding actions when agent has high vengefulness value (V).

At the simulation’s final label, “regenerate_population”, successful agents are mutated and form a new generation with the same size [6]. In this simulation, folded outlined arrows represent iteration: one for the three norms and their corresponding expectations within one generation, and the other for 100 generations of simulation.

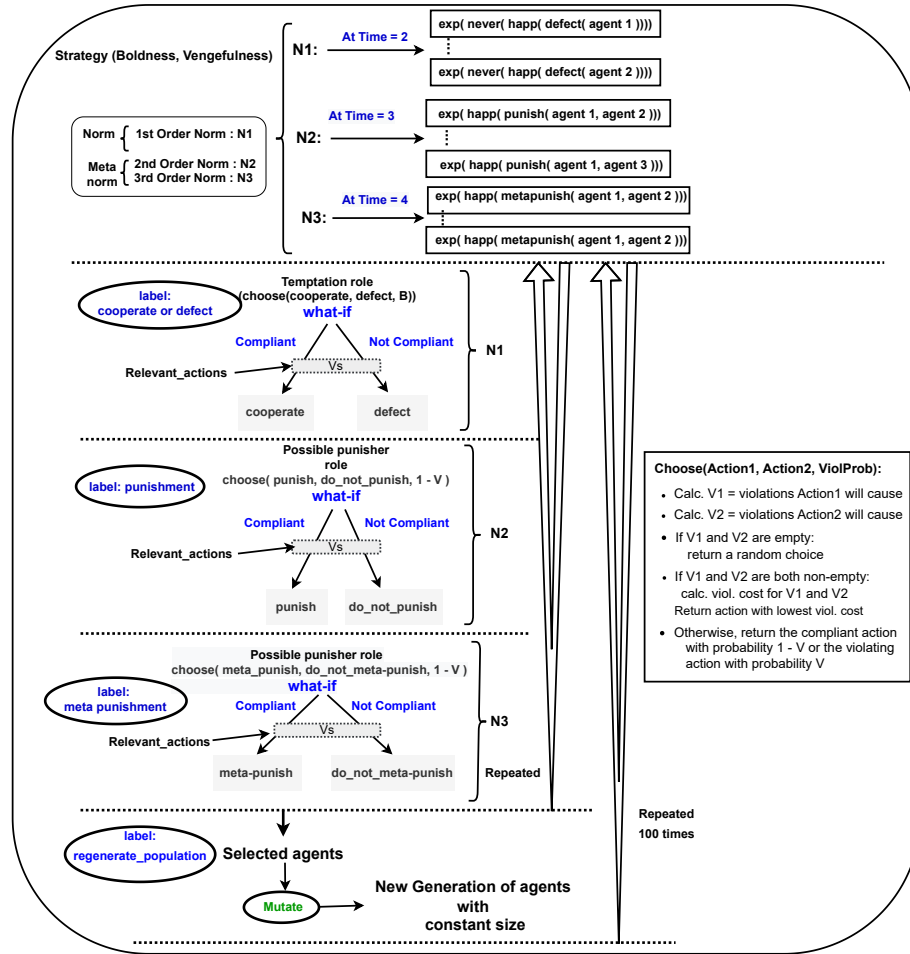


Fig. 4: The generalisation of Axelrod's approach in which norms are explicitly represented and agents can choose their actions based on *what-if* reasoning using expectation event calculus, which tracks the creation, fulfilment, and violation of expectations.

4 Results

A. Experiment 1: Generalisation of Axelrod's metanorm game The simulation of the experiment depicts what happens when agents violate the second order norm and are metapunished by vengeful agents adhering to the *third order norm*. Figure 5(a) shows a scatter plot representation demonstrating that boldness is always low and vengefulness ranges from high to average. We use twenty agents, 100 generations, and 100 runs to simulate Axelrod's generalisation

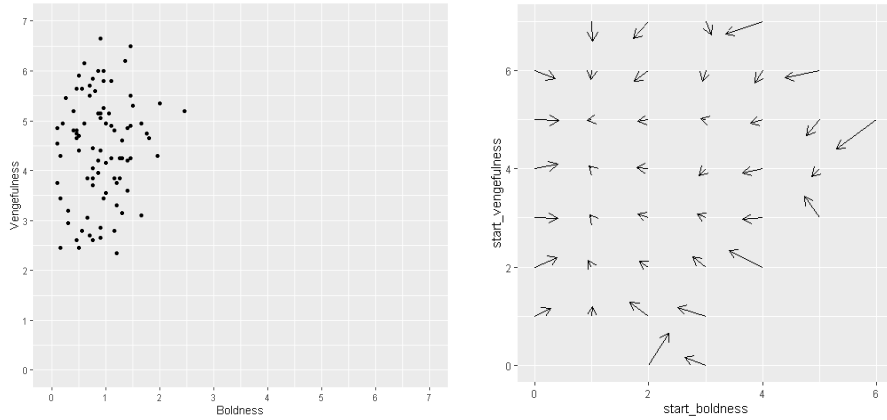


Fig. 5: (a) **Scatter plot** of mean boldness along x-axis wise and mean vengefulness along y-axis wise of generalisation of Axelrod’s study with all three norms. (b) **Vector plot representation** of mean boldness and vengefulness.

of explicit norms. The scatter plots generated by the mean value of boldness and vengefulness in each of 100 runs.

Figure 5(b) depicts the vector representation of the same data set ⁸. Vectors show how boldness and vengefulness change across generations in the population. The results show that *what-if* reasoning with explicit norms causes the first-order norm to be largely upheld in the society due to low boldness being maintained.

However, compared to Axelrod’s results with implicit norms, there remains some moderately low vengefulness values. This may be explained by a limitation of the rule engine we use: only a single string predetermined string can be returned as the result of a rule, e. g. “punish”. Thus, when the rule’s condition compare’s an agent’s vengefulness with a random number, the outcome is either to punish all or none of the observed violators. In contrast, in Axelrod’s algorithm, each individual punishment is the result of a different randomise decision. Further investigation is needed to determine whether this explains the ability for agents with lower vengefulness to remain in the population.

B. Experiment 2: Using the boldness and vengefulness in other Scenario [9] introduced a scenario that we refer to as the plain-plateau scenario in our previous work. The scenario depicts a society in which people have the option of living on a river plain with easy access to water, otherwise, they can live on a plateau. Flooding is a risk for river-plain residents. When the government has complete discretionary power, it is in the government’s best interests to

⁸ Populations with similar average levels of boldness and vengefulness are grouped together to create each vector. The end point of each arrow shows the average levels of these features one generation later.

compensate citizens whose homes have been flooded by taxing citizens who live on the plateau, creating a prisoner’s dilemma situation. In previous work [12], we experimented with the use of social norm-based expectations to achieve coordination where citizen agents are hard-coded to prefer actions that will result in no violation.

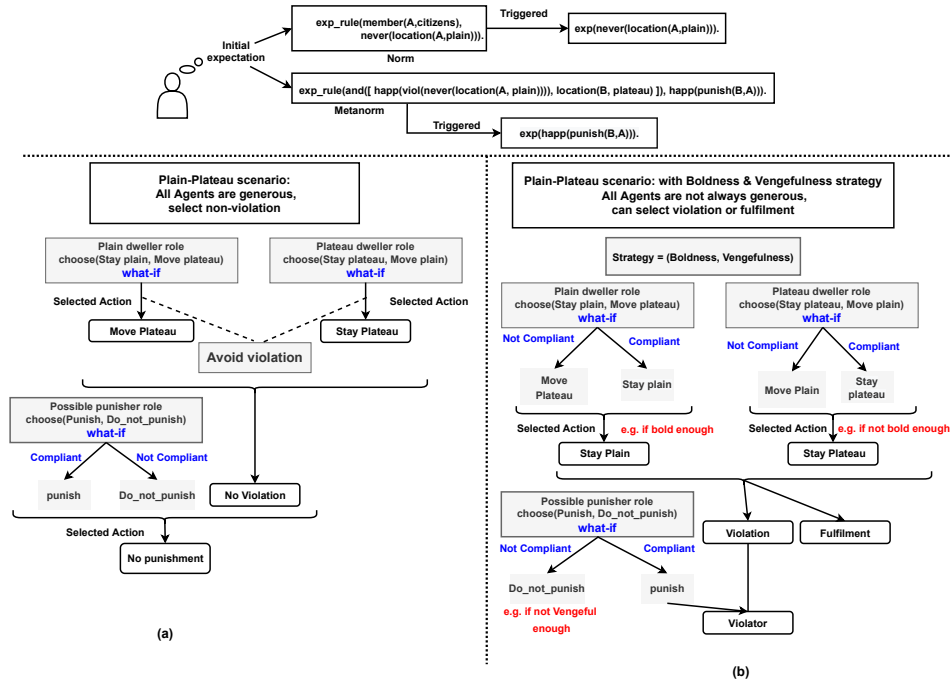


Fig. 6: (a) **The plain-plateau scenario** in which agents are hard-coded to always choose the non-violating actions. (b) **The plain-plateau scenario** as an application of our generalised Metanorms game.

Figure 6(a) illustrates this prior work. Each agent has either a plain dweller or a plateau dweller role, and in each simulation cycle there are two choices: agents can stay on the plateau or move to the plain. In this scenario, we assume there exists a norm that no one should live in the plain and a metanorm stating that plateau dwellers should punish those who live on the plain.

Figure 6(b) illustrate the application of our generalisation of the metanorms game to this scenario⁹. We import boldness and vengefulness parameters from a run of the experiment 1. This simulates an agent finding itself in a new scenario after having already evolved its personality with respect to norms, and illustrates

⁹ Currently, we do not include a level 3 norm for the plain-plateau scenario.

the generality of our approach using explicit norms and *what-if* reasoning. In the plain-dweller role, the EEC *what-if* predicate is used to consider two options: move to the plateau or stay in the plain; similarly, in the plateau-dweller role, the *what-if* mechanism is used to consider either move plain or stay plateau. When agents with high boldness in the plain-dweller role choose to stay in the plain, they violate the norm. When other vengeful agents observe violators, they punish them unless insufficient vengefulness causes them to violate the metanorm.

5 Prior event calculus models of norms

This section of the paper reviews some research on the use of event calculus in autonomous agent reasoning to examine the effects of norms.

Alrawagfeh [1] suggests formalising prohibition and obligation norms using event calculus and offers a method for BDI agents to reason about their behaviour at runtime while taking into account the norms in effect at the time and previous actions. Norms are represented by EC rules that initiate fluents with special meanings. The introduced fluents represent punishments for breaking a prohibition norm or failing to fulfil obligation norms, or the rewards for fulfilling obligation norms. The normative reasoning strategy assists agents in selecting the most profitable plan by temporarily asserting to the event calculus the actions that each plan would generate and considering the punishments and/or rewards it would trigger.

In Alrawagfeh’s work, norms cannot be changed dynamically without changing the event calculus rule base, because they are defined by EC initiates clauses. In contrast, in our approach, EC rules can be instantiated automatically from *exp_rule* fluents, which can be changed dynamically by events.

Alrawagfeh has no representation of active normss, violations or fulfilments: only punishments and rewards. In our work, expectation creation, fulfilment, and violation are represented as events, and the *what-if* predicate compares alternative events to track expectation creation, fulfilment, and violation. We do not assume that rewards and/or punishments will always follow violations and fulfilments; these could be defined by separate *exp_rule* or EC indicates clauses.

Hashmi et al. [7] proposed a new EC predicate for modelling all types of legal norms in order to capture the deontic effects of obligations. They do this to ensure that a norm enters into force at the same time as the triggering event occurs. In contrast, our approach does not necessitate the introduction of a new type of EC predicate in order to initiate a deontic predicate. An *exp_rule* or an expectation can be created with a standard initiates clause and an *exp* fluent is created by an *exp_rule* in the state where the condition of the rule becomes true. However, the EEC includes additional axioms to deal with the progression of expectations.

Alrawagfeh and Hashmi et al. both use standard EC, whereas we use discrete EC because this work involves discrete time simulations.

6 Conclusion and future work

In previous work, we used the EEC *what-if* mechanism for choosing actions in the presence of expectations how agents adhere to norms and choose actions using the *what-if* mechanism. However, we assumed that all agents are hard coded to avoid expectation-violating actions.

The current work builds on this previous work to remove this assumption, but it also makes the following significant standalone contribution. It generalises Axelrod's study in which norms are explicitly represented. This allows the metanorm game to be used across multiple scenarios.

There are several possibilities for future work. Axelrod's simulated hundreds of generations, and we are proposing this mechanism in real-time agents with different scenarios, implying that the learning mechanism must be different. Population evolution would take place in this case over long periods of time, so it's not useful. Instead we could use a pairwise comparison approach where an agent may adopt another agent's strategy based on a comparison of their respective fitnesses, using the Fermi equation [15].

References

1. Alrawagfeh, W.: Norm representation and reasoning: a formalization in event calculus. In: International Conference on Principles and Practice of Multi-Agent Systems. pp. 5–20. Springer (2013)
2. Axelrod, R.: An evolutionary approach to norms. *American political science review* **80**(4), 1095–1111 (1986)
3. Bicchieri, C.: *The grammar of society: The nature and dynamics of social norms.* Cambridge University Press (2005)
4. Cranefield, S.: Agents and expectations. In: International Workshop on Coordination, Organizations, Institutions, and Norms in Agent Systems. pp. 234–255. Springer (2013)
5. Cranefield, S., Clark-Younger, H., Hay, G.: A collective action simulation platform. In: Multi-Agent-Based Simulation XX: 20th International Workshop, MABS 2019, Montreal, QC, Canada, May 13, 2019, Revised Selected Papers 20. pp. 69–80. Springer (2020)
6. Galan, J.M., Izquierdo, L.R.: Appearances can be deceiving: Lessons learned re-implementing axelrod’s evolutionary approach to norms. *Journal of Artificial Societies and Social Simulation* **8**(3) (2005)
7. Hashmi, M., Governatori, G., Wynn, M.T.: Modeling obligations with event-calculus. In: Rules on the Web. From Theory to Applications: 8th International Symposium, RuleML 2014, Co-located with the 21st European Conference on Artificial Intelligence, ECAI 2014, Prague, Czech Republic, August 18-20, 2014. Proceedings 8. pp. 296–310. Springer (2014)
8. Holzinger, K.: The problems of collective action: A new approach. MPI Collective Goods Preprint No. 2003/2, SSRN (2003), doi:10.2139/ssrn.399140
9. Klein, D.B.: The microfoundations of rules vs. discretion. *Constitutional Political Economy* **1**(3), 1–19 (1990)
10. Mueller, E.T.: *Commonsense Reasoning.* Morgan Kaufmann (2006)
11. North, M.J., Collier, N.T., Ozik, J., Tatara, E.R., Macal, C.M., Bragen, M., Sydelko, P.: Complex adaptive systems modeling with repast simphony. *Complex Adaptive Systems Modeling* **1**(1), 3 (2013)
12. Sengupta, A., Cranefield, S., Pitt, J.: Solving social dilemmas by reasoning about expectations. In: Coordination, Organizations, Institutions, Norms, and Ethics for Governance of Multi-Agent Systems XIV: International Workshop, COINE 2021, London, UK, May 3, 2021, Revised Selected Papers. pp. 143–159. Springer (2022)
13. Shanahan, M.: The event calculus explained. In: *Artificial Intelligence Today*, pp. 409–430. Springer (1999)
14. Thøgersen, J.: Social norms and cooperation in real-life social dilemmas. *Journal of economic psychology* **29**(4), 458–472 (2008)
15. Wu, B., Bauer, B., Galla, T., Traulsen, A.: Fitness-based models and pairwise comparison models of evolutionary games are typically different—even in unstructured populations. *New Journal of Physics* **17**(2), 023043 (2015)