The Nexican Stand-Off: Social Contracts and Popular Legitimacy in n-Player High Stakes Resource Competition Games

Matthew Scott[1] and Jeremy Pitt[1]

[1]Imperial College London, SW6 2AZ
mss2518@ic.ac.uk

Abstract

In a high-stakes cooperative survival game, self-interested behaviours reward the individual in the short-term but may have a detrimental impact on the collective in the long-term. Such situations can be solved by introducing social contracts between players that reduce the set of possible actions. In the absence of an empowered authority capable of enforcement, however, a player will only uphold such a contract so long as they believe that the other players will do the same. We term this *buy-in*. In this context, we envision a cooperative survival game that extends the scope of the 'conventional' Mexican standoff (a three-player Hawk-Dove game) to n-players, from which we design and implement a self-organising multi-agent system. We devise a set of experiments across varying degrees of initial buy-in and examine its impact on social contracts and the voluntary restriction of self-interest. In particular, we show that there is a cyclical, non-transitive dependency between the three that is both ring-reinforcing and critical for systemic stability.

Introduction

Cooperative survival games refer to a subset of games wherein players must work together to overcome disaster, else suffer the consequences through both personal and communal damage. Common examples of these games are computer games (e.g. Minecraft), board games (e.g. Ravine) and sociological constructs (e.g. Global Warming).

These games are often solvable with the introduction of self-organising mechanisms (Rezaei et al., 2009), such as governance (Ostrom, 1990), where an empowered body can introduce rules to the game. Issues arise, however, in the absence of an external authority where such mechanisms cannot be enforced (Ostrom, 1990), as there cannot be monitoring or sanctioning to dissuade from acting immorally. To this end, self-organising mechanisms designed to simplify such games may be entirely useless, or worse, interact in unexpected, even pernicious, ways (Serugendo et al., 2005; Nafz et al., 2013).

With such an absence in authority, the self-organisation *must* be mutually agreed and players must elect to obey the *Rule of Law*, (Bingham, 2011) where all players are accountable to the same laws that the mechanism enforces. This agreement to obey the law is described as *popular legitimacy* and concerns the extent to which players have confidence in the rules of society, and can directly impact the enforceability of contracts (Davis, 2004).

One example of where popular legitimacy may be used in such cooperative problems is the two-player 'Hawk-Dove' (Smith and Price, 1973) game, where players can choose to attack or defend a shared resource. Both players attacking will result in their mutual destruction, but choosing to defend, while the other attacks, will leave the defender looking like a 'chicken' (hence the colloquial "game of chicken"). This game is extended to three players in the colloquial 'Mexican Stand-Off', in which attacks are fatal and any one player initiating an attack on a second leaves it open to attack by the third, leaving that third player the survivor.

This paper generalises the game, firstly to n players, creating the *Nexican Stand-Off*; secondly, by iteration, giving the opportunity for the stand-off itself to create externalities; and thirdly, by giving players an exogenous allocation of resources per round that can be invested in offence, defence or personal utility, and that can be stolen following a successful attack. This transforms the conventional 2-player, 2-action Hawk-Dove game into a high-stakes n-player, 4-action resource competition game of deterrence, where an attack may be 'rational' in order to survive resource insufficiency but 'irrational' in that it guarantees elimination.

However, the stand-off buys time for the players to self-organise through socially-constructed contracts (externalities) (de Cesare and Geerts, 2012); but, without an external authority for enforcement, each player must decide to maintain these contracts based on how well they perceive other players to uphold them. We term this concept *buy-in*. The question then becomes: "what degree of initial buy-in creates popular legitimacy, and can it be used to introduce social contracts that help defuse otherwise mutually assured destruction?"

Accordingly, this paper is structured as follows: we begin by proposing a scenario that encompasses the aforementioned problem. Then, we introduce the principles of survival games, game theory and potential solution concepts

Downloaded from http://mitp.silverchair.com/isal/proceedings-pdf/isal2023/35/111/2354947/isal_a_00672.pdf by guest on 18 September 2024

for such games in the form of social contracts and popular legitimacy. From this, we develop a multi-agent simulator and subsequently discuss the agent design.

Following the specification of the simulator tool, we carry out a set of experiments that simulate a set of 'survival trials', where agents play the aforementioned game over varying levels of initial buy-in and availability of social contracts. The results of these experiments demonstrate that the problem is solvable given 1.) available social contracts, 2.) contracts that appropriately restrict the action space, and 3.) a sufficiently high degree of initial buy-in. Furthermore, these three variables are shown to be ring-reinforcing and have a cyclical, non-transitive dependency.

Scenario and Background

Scenario

Based on the specification for the problem introduced above, we envision an environment as shown to Figure 1 (a screenshot from the simulator described later). This scenario features n players grouped into k clusters, each facing a Nexican-standoff type situation. This continues in an iterated series of rounds, for an indeterminate number of rounds. Each agent wants to use its resources for personal satisfaction and enjoyment. The dilemma, however, is that every other player would *also* like to use its resources for their own personal satisfaction, but each player is greedy and more resources means more satisfaction, since this avoids resource insufficiency.

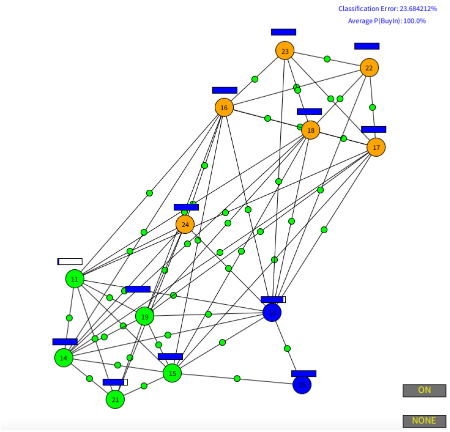
Classification Error: 23.684212%
Average P(BuyIn): 100.0%

Figure 1: Visualisation of n-player Hawk-Dove across multiple clusters using simulator tool described below.

For this reason, players would happily attack/kill another player and steal their resources. To mitigate this, a player can use some of their resources to invest in offence, for a better chance of a successful attack, or some of their resources to invest in defence, for a better chance of *mitigating* the attack. In addition to this, players may form contracts with one another through communication and 'promise' to not attack.

Importantly, all n players don't all play against one-another at any one time. The town that this game occurs in has many different neighbourhoods of players of varying size. Once each day is complete, the players are free to move to other neighbourhoods as they see fit.

Finally, all players are 'genetically coded' with a different social motive (Reinders Folmer, 2016) (such as altruism, competitiveness or narcissism) that influences their strategy for playing against other players. As they are all perfect logicians, this strategy inherits from the conventional Hawk-Dove game, where having a different social motive means that this player will place a different valuation on the payoffs associated with attacking or defending compared to others.

Hawk-Dove Games in Social Systems

Based on the scenario, it is clear that there is a conflict between choosing to attack or defend (or indeed invest in utility). Choosing to defend whilst a player's opponent attacks will yield a net gain in utility for the attacker, and vice versa. Furthermore, a mutual attack will result in damage that may be detrimental to both players. For this reason, we reason that this game is intuitively similar to the Hawk-Dove game, from game theory.

Conventionally, we see the Hawk-Dove Game applied to analysing resource conflicts, such as food in animal kingdoms (Houston and McNamara, 1988) or populations of mating systems (Grafen, 1979; Smith and Price, 1973) and assessing the evolutionary stable state that is reached.

Smith developed the notion of the 'Hawk-Dove game' when trying to mathematically model the logic of animal conflicts (Smith, 1982). This research investigated two ideologies of animal, the 'Hawk' and 'Dove'. A Hawk will "escalate, and continue to do so until injured or opponent retreats" and a Dove will "display, [and] retreat if opponent escalates, before getting injured". The 'fitness' of such animals playing their strategy (the weighted payoffs of applying such a strategy, if it is played with probability p) was subsequently investigated to establish an evolutionary stable strategy, if the game is iterated over multiple turns.

The Hawk-Dove game is also colloquially seen as the "fight or flight" response (Cannon, 1925), where the experience of a threat polarises animals to either fight or flee, based on an innate response of their sympathetic nervous system. This phenomenon has been further investigated to suggest that there may instead be four responses - "fight", "flight", "freeze" and "fawn" - based on a trauma response (Mahaney, 2022; Happe, 2021). We use this extension of the conventional two-action Hawk-Dove game as justification for introducing our own extended action space - boost attack, boost defence, launch attack and boost utility - over the standard "attack/cower" action space.

Downloaded from http://mitp.silverchair.com/isal/proceedings-pdf/isal2023/35/111/2354947/isal_a_00672.pdf by guest on 18 September 2024

Social Contracts in Survival Games

A key problem surrounding cooperative survival games is the issue of self-interested agents acting selfishly (Rovatsos and Lind, 1999). Logical agents will implement calculated strategies that provide personal benefit *instead* of furthering the survivability of the collective (Rovatsos and Lind, 2000), and hence preventing socially desirable outcomes.

Punishment and sanctioning are pre-requisites for a self-governing institution (Ostrom, 1990). Since this problem naturally imposes a polycentric system of governance (there is no centralised authority), these methods of discipline are effectively unenforceable, so social contracts are offered as a *possible* solution.

Davoust (Davoust and Rovatsos, 2020) experiments with the introduction of *social contracts*, as a means of stabilising a system "to produce an optimal outcome in terms of social welfare". This work formalises social contracts as a function of "morality", (f) "payoff" $(\mathbf{u})$ and "strategies", (π) asserting that, for a given game (G), the morality of a given action is sub-optimal:

$$f(\mathbf{u}(\pi(G))) < f(\mathbf{u}(a^*)) \tag{1}$$

The role of a social contract, in this case, is hence to generate a *modified* game (G'), where:

$$f(\mathbf{u}'(\pi(G'))) = f(\mathbf{u}(a^*)) \tag{2}$$

Davoust subsequently proves that the social contract design problem is solvable for "any game and any deterministic, common-knowledge decision procedure" (Davoust and Rovatsos, 2020).

An abstraction of social contracts comes in the form of 'treaties'. Serving both "as a juristic act and as a rule" (Reuter, 1995), treaties act as a form of institutionalised power between players, with society coming to a mutual agreement upon the importance of such a legal device. Ostrom applies the notion of treaties as a means of augmenting the Prisoner's Dilemma (Ostrom, 1990; Pérez-Cirera, 2010), by having players negotiate a contract prior to playing the meta-game. In this formulation, the payoff of the game becomes deterministic, so long as both players can successfully negotiate the contract.

Popular Legitimacy and Buy-In

Although the existence of an optimal social contract is provable, where such a contract is capable of modifying an existing game to yield an optimal strategy, the forced adherence to such a contract without an external authority to reprimand defectors poses a second problem.

Implicitly, social contracts require social construction and hence only yield institutional power so long as there is *popular legitimacy*, where there is popular acceptance of this system of governance. Popular legitimacy is established through the notion of *buy-in* where members of a system must have confidence in the social rules. A lack of popular legitimacy directly impacts the enforceability of contracts (Davis, 2004) and indeed the acceptance of the rule of law (Prado and Trebilcock, 2021), to the extent that the capacity for limiting behaviours in a social system and the degree of popular legitimacy are "mutually reinforcing" (Prado and Trebilcock, 2021).

Stability

In this heterogeneous system, logical players wish to act out of self-interest to maximise their individual utility. In doing so, however, players risk eliminating other players which may lead to systemic collapse. The aim, therefore, is to understand under what conditions such stand-offs are stable, and if conditions lead to a collapse, whether it is possible to learn the conditions which lead to the collapse, and so avoid them in an iterated n-player stand-off. Ultimately, we aim to discover if it is possible for a polycentric system to be stable when the only form of graduated sanctions is 'soft', which is to say that other forms of punishment are effectively unenforceable.

For this, we consider the concept of 'polystability' (Ashby, 1960), which suggests that a system may comprise many equilibria, where the values held by the components in the system may be different to the initial calibration. Inductively, we hypothesise a polystable system with two equilibria - either the treaties are sufficiently powerful to negate any incentive to defect, or all players die due to mutually assured destruction.

Summary

It can hence be reasoned that, given the presence of multiple social motives in the proposed scenario, an absence of centralised governance, the extension of scope to n players and the fact that the game is iterated, a purely analytic approach is not feasible for solving this problem.

Therefore, we need for specification a combination of introspective game theory (i.e, social motives), socially-constructed political meta-games, (i.e, social contracts), and for predictive leverage, rather than a solution concept we need simulation, which is the subject of the following chapter, where we employ a self-organising, multi-agent system to simulate this game.

Simulator Design and Implementation

For this paper, we use a conventional *server-agent architecture*. The server proposes a set of actions to each agent (such as utility investment, treaty formation, etc.) and retrieves the response. This response is then handled by the server to update the environment and pass messages to the other agents.

To address this challenge, we utilise *Processing*, a Java-based language for visualisation, using the paradigm of a **self-organising, multi-agent system**, specifically with Agent-Based Modelling (ABM).

Downloaded from http://mitp.silverchair.com/isal/proceedings-pdf/isal2023/35/111/2354947/isal_a_00672.pdf by guest on 18 September 2024

Simulator Overview

The complete simulator can be abstracted to two main processes which run successively once per turn. Firstly, the agents *move*. Secondly the agents *interact* by 1.) proposing *treaties* and 2.) selecting an *action*. Given these stages, we abstract the full simulation process in Figure 2.

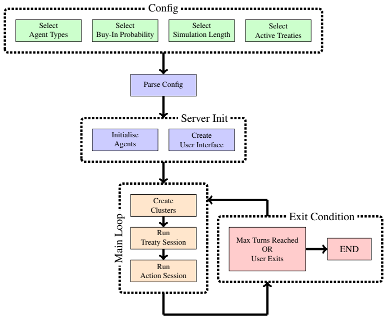

Figure 2: Visualisation of Simulator Control Loop

Treaties as Social Contracts

Treaties serve as social contracts between agents that inhibit certain actions and help develop a social network. Treaties are represented as an agreement between two agents, and the action that it restricts. For this paper, we investigate a single treaty: agents may not invest more than 20 points into offence on a single turn, and agents may not attack on any given turn.

Neighbourhoods

Neighbourhoods are formed using a standard 'K-Means Clustering' algorithm (Hartigan and Wong, 1979), where the number of total clusters is equal to half of the number of alive agents, floored. This allows for neighbourhoods to persist even with a single agent remaining.

Agent Design and Implementation

All agents in this system are provided with the attributes needed to engage in the *treaty session*, *action session* and keep track of its *buy-in*. Therefore, agents have variables to store their *buy-in*, *offence*, *defence*, *utility* and a set of the current treaties they are engaged in.

Agent Types

To introduce heterogeneity into the system, the agent design incorporates a set of eight different social motives. Social motives "reflect the way that people value [others'] interests in relation to their own" (Reinders Folmer, 2016) and have seen prior use in multi-agent systems (Scott et al., 2022). In this implementation, these socially-driven agents make actions based on the ordinality and cardinality of the Hawk-Dove game's payoff matrix. We abstract eight different social motives and summarise them as follows, creating four pairs of 'opposing' ideologies:

Individual: Individual agent types aims to get the best utility for themselves. They hence choose the quadrant that gives the highest payoff.

Martyr: Martyr agent types aims to get the *worst* utility for themselves. They hence choose the quadrant that gives the lowest payoff.

Cooperative: Cooperative agent types aim to strengthen the collective. They hence sum the values of the payoff matrix and choose the highest overall summed payoff.

ICM: ICM agent types aim to sabotage the collective. They hence sum the values of the payoff matrix and choose the lowest overall summed payoff.

Competitive: Competitive agents aim to outperform their rival as much as possible. Hence, this agent type aims to maximise the difference between payoffs.

Equitable: Equitable agents aim to equalise with their rival. They hence aim to minimise the difference between payoffs.

Altruistic: Altruistic agents aim to provide the best result for their opponent. They hence choose the quadrant that gives the highest payoff for the rival.

Aggressive: Aggressive agents aim to sabotage their opponent. They hence choose the quadrant that gives the lowest payoff for the rival.

Identification and Mobility

Through repeated interactions, agents develop a model of the social motive of all other agents in the system. To achieve this, with each interaction an agent will 'poll' its opponent for its social motive, where the accuracy of the response is proportional to the responding agent's buy-in. This means that an agent with 100% buy-in will respond with the appropriate social motive, whereas an agent with 0% buy-in will respond with the 'opposing' social motive.

This also informs how agents move in 2-D space by using the 'boids' algorithm Reynolds (1987). This algorithm comprises three main concepts: *cohesion*, where agents move toward the average position of local flockmates, *alignment*, where agents more towards the average heading of local flockmates and *separation*, where agents move to avoid crowding local flockmates. Flockmates are selected using the adage of "birds of a feather, flock together", where each agent iterates over the other agents in the system, and considers them a flockmate if they share the same social motive as itself. This results in clustering, where each neighbourhood is populated by a homogeneous social motive.

Action Selection

We have previously introduced how the 'attack' and 'stick' actions of the conventional Hawk-Dove game are abstracted in this paper, with the search-space extended to 1.) launch-

Downloaded from http://mitp.silverchair.com/isal/proceedings-pdf/isal2023/35/111/2354947/isal_a_00672.pdf by guest on 18 September 2024

ing an attack, 2.) boosting offence, 3.) boosting defence and 4.) boosting individual utility.

Matrix Generation Agents select their action as a function of three variables. Firstly, they dynamically generate a Hawk-Dove payoff matrix in the form of its evolutionary stable state (ESS) (Smith and Price, 1973), hence taking the form of the game in Figure 3. Each agent generates **two** matrices: one for themselves and one for their opponent, since the numerical payoffs of each matrix vary.

		Opponent	
		Attack	*Defend*
Initiator	*Attack*	$(\frac{V-C}{2},\frac{V-C}{2})$	(V,0)
	Defend	(0,V)	$(\frac{V}{2},\frac{V}{2})$

Figure 3: Numerical payoffs of an ESS Hawk-Dove Game

Here, V and C represent the *value* and *cost* of a contested resource, respectively. Reinterpreting this ESS for the simulator, we take the numerical payoffs for both the initiating agent and opposing agent as follows:

The *initiating agent* takes the value of the resource, V, as the opponent's utility and the cost of the resource, C, as the difference between the opponent's offence and initiator's defence. Conversely, when constructing the payoff matrix for the *opponent*, V takes the value of the initiator's utility and C takes the value as the difference between the opponent's defence and initiator's offence.

The rationale behind this is that, were the two agent both to mutually attack one another, they would suffer damage equivalent to the difference in their defence and their opponent's offence, but stand to gain the opponent's utility, were the attack to succeed.

Preference Order Following the construction of payoff matrices for both the agent selecting the action and their opponent, a preference order for the different quadrants is taken. For an *altruist* agent, say, they wish to play a strategy that results in the largest payoff for the opposing agent. For this reason, they 'augment' the payoff matrix to comprise only the payoffs relative to the opponent and order the quadrants by descending payoffs.

If an agent wanted instead to minimise the payoffs for the opposing agent, the entire augmented matrix is negated before ordering the quadrants by decreasing payoffs. The resulting process is described in Equation 3.

$$\begin{bmatrix}(60,60) & (100,0)\\(0,100) & (50,50)\end{bmatrix} \rightarrow \begin{bmatrix}(60) & (0)\\(100) & (50)\end{bmatrix} \rightarrow 2>0>3>1 \tag{3}$$

where the augmented matrix is calculated for an altruist agent and the quadrants are numbered 0-3 from the top-left to bottom-right. Following this, the agent will also generate a preference order for the *opposing* agent, based on what the initiating agent believes is the opponent's social motive.

Voting Once the two preference orders have been generated, a Borda count vote is applied to determine the 'optimal' quadrant that best appeases both parties. Equation 4 discusses how this is achieved.

$$\sum_{i=1}^{n} n - \mathbf{rank}(a, v_i) + 1 \tag{4}$$

Once this quadrant is obtained, the initiating agent plays the action that will land them in this quadrant; if the quadrant is 0 or 1, the agent will launch an attack, else they will boost their defence.

It is important to note that this strategy doesn't encompass the entire action space, as boosting offence and utility aren't represented in the payoff matrix. For this reason, agents have a 25% chance of boosting utility, a 25% chance of boosting offence and hence a 50% chance of playing this strategy.

Treaty Violations If an action is *forbidden* by a treaty, the agent is faced with a dilemma. They must either abide by the treaty, and choose to boost utility **only** or continue through with the action anyway, violating the treaty. Agents will ignore the restrictions of the treaty with a probability of 100% minus their buy-in, which results in the treaty violation being broadcast to any agents involved.

Upon receiving a notification that the treaty has been violated, agents will choose to uphold the treaty with a probability equal to their buy-in probability, or else sever it entirely. Their overall buy-in will decrease by 5%, however, for each treaty violation.

Combat By electing to launch an attack, agents follow the procedure determined in Algorithm 1 based on the quantity of attack they wish to supply.

Algorithm 1 Combat algorithm

```
$A \leftarrow [0, attacker.offence]$
$D \leftarrow defender.defence$
if $A \geq D$ then
    $damageDealt \leftarrow A - D$
    $defenderHP \leftarrow defenderHP - damageDealt$
    $defender.defence \leftarrow 0$
else
    $defender.defence \leftarrow defender.defence - A$
end if
$attacker.offence \leftarrow attacker.offence - A$
```

Upon receiving an attack, agents will decrease their buy-in proportionally to the damage they receive. This is illustrated in Equation 5.

Downloaded from http://mitp.silverchair.com/isal/proceedings-pdf/isal2023/35/111/2354947/isal_a_00672.pdf by guest on 18 September 2024

$$buyIn = max(0, buyIn - \frac{damageDealt}{100.0}) \quad (5)$$

Experimental Design and Results

By using the aforementioned simulator, the purpose of this research is to construct a set of ‘survival trials’, each with varying initial configurations, that can be used to demonstrate the following set of hypotheses:

> **Hypothesis A** *Higher degrees of initial buy-in permit a longer adherence to contracts.*
>
> **Hypothesis B** *A collective can achieve popular legitimacy in the presence of social contracts, so long as the initial buy-in is sufficiently high.*
>
> **Hypothesis C** *Systemic stability can exist, so long as there is a sufficiently high initial buy-in, and social contracts are allowed.*

Experiments A: Buy-In and Contract Adherence

For this set of experiments, 5 agents of each social motive are simulated across 350 turns with initial buy-in varying from 0 to 1.0 in steps of 0.25. These experiments illustrate the shrinking disparity between the number of treaty proposals (and subsequent acceptances) and treaty breaks due to forbidden actions as the initial buy-in decreases.

From Figure 4, A1, an initial buy-in of 100% yields a system in which there is no treaty dissolution. Agents adhere to treaties once they have been mutually instated, which ensures that the action space is sufficiently restricted. As the initial buy-in is decreased, as with Figure 4, A2-5, it is clear that the disparity between the number of treaty proposals and the number of treaty breaks decreases, resulting that the net change in active treaties plateaus, instead of maintaining a positive gradient.

This illustrates that the total number of active treaties grows at a slower rate with decreasing initial buy-in. Furthermore, we conclude that buy-in has a direct impact on contract enforcement, as decreasing the initial buy-in results in more frequent treaty breaks.

Figure 5, A6 further demonstrates the impact of buy-in on treaty persistence. Clearly, there is a direct correlation between the initial buy-in and the total number of active treaties during simulation. Agents will execute a forbidden action with a probability that is inversely proportional to their buy-in, however uphold a treaty *proportionally* to their buy-in. For this reason, as an agent’s buy-in decreases, it is more likely to act outside of a treaty, and less likely to forgive an agent for breaking a treaty, causing a net decrease in the number of active treaties. This creates a sense of mutual reinforcement, as less frequent treaty breaks allow buy-in to improve, which in turn disincentivises agents from breaking a treaty.

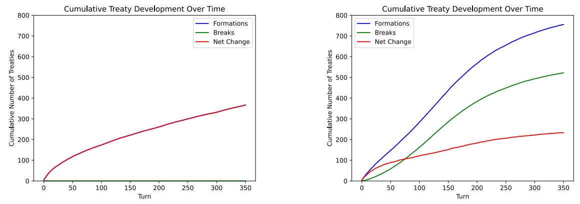

A1: Initial Buy-In 100%   A2: Initial Buy-In 75%

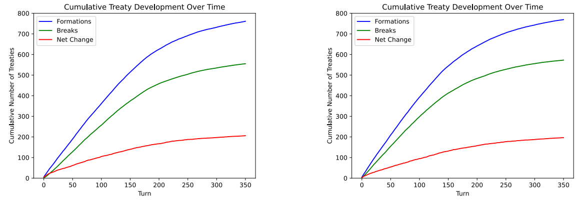

A3: Initial Buy-In 50%   A4: Initial Buy-In 25%

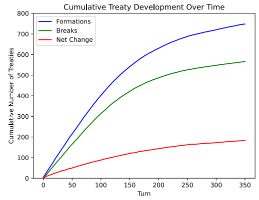

A5: Initial Buy-In 0%

Figure 4: Cumulative treaty proposals, breaks and net change with varying initial buy-in

Experiments B: Buy-In Over Time

This set of experiments are designed to synergise with the experiments performed in Experiments A to greater reveal the degree of mutual reinforcement between restricted actions, popular legitimacy and the adherence to social contracts by investigating how buy-in varies over time. These experiments ensure the same initial configuration as A, however we also investigate an initial buy-in of 0.8 and 0.85 to better illustrate the behaviour around the critical point of stability.

Figure 6, B1 illustrates the direct correlation between the initial and final buy-in when treaty formation is enabled. So long as the initial buy-in is above a threshold of approximately 80%, the effect is self-supportive; the average buy-in over time either persists, or tends to 100%. Conversely, with an initial buy-in below this, the terminal buy-in tends to approximately 55%, although this final value is still (broadly) proportional to the initial level, since an initial buy-in of 0% tends to 45% and an initial buy-in of 80% tends to 60%. Figure B1 also supports the result of Figure 5, A6, as we see that lower degrees of buy-in result in contracts being broken more frequently, which keeps buy-in low.

Figure B2 further investigates the development of average buy-in over time in the absence of treaties. In this case, it is clear that irrespective of initial buy-in, a terminal buy-in of

Downloaded from http://mitp.silverchair.com/isal/proceedings-pdf/isal2023/35/111/2354947/isal_a_00672.pdf by guest on 18 September 2024

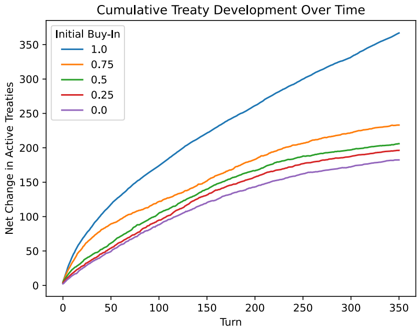

Figure 5: A6: Net change in active treaties over time with varying initial buy-in

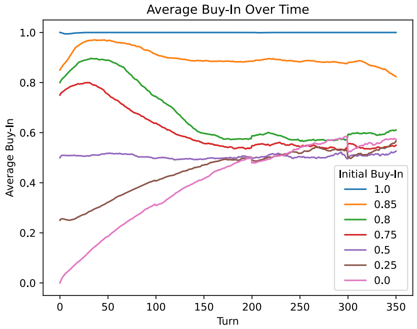

Figure 6: B1: Average buy-in over time with treaty formation *enabled*

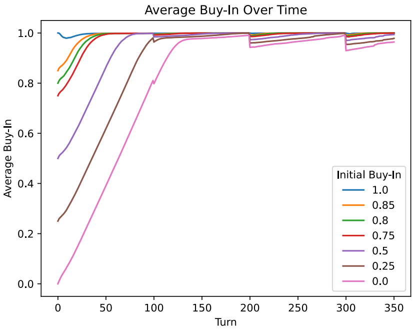

Figure 7: B2: Average buy-in over time with treaty formation *disabled*

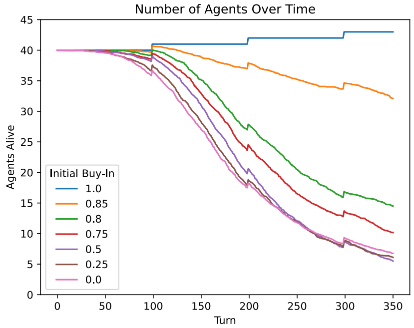

Figure 8: C1: Average agents over time with treaty formation *enabled*

100% can be achieved. We reason that, due to the absence of treaties, there cannot be any treaty breaks and hence the probability of buy-in cannot decrease through this method.

Experiments C: Survivability

This final set of experiments uses the same configuration as Experiments B, and investigates the effects of initial buy-in and availability of contracts on survivability.

Experiment C1 in Figure 8 shows that, following a sufficiently high initial buy-in of 85% and with contracts enabled, it is possible to stabilise the system. We reason that the mutually supportive relationship from Experiments A and B allows for high buy-in to persist and therefore contracts to be upheld. This means that the agents mutually agree not to attack each other, allowing for the heterogeneous system to exist and stabilise, where the terminal number of agents either increases (100% buy-in) or stays constant (85% buy-in).

Removing the contracts however, as with Experiment C2 in Figure 9, results in systemic collapse across all degrees of initial buy-in. Clearly, an absence of contracts means that actions cannot be restricted, and that agents have no incentive not to attack one another. Instead, the number of agents plateaus to a constant distribution, where despite the absence of contracts, all agents can survive. We theorise that this due to the system reconfiguring itself, where the only social motives that remain play a ‘stalemate’ strategy (such as only defending, for example).

Discussion and Further Work

Discussion

From the results of Experiments A and B, we assert that buy-in and contract adherence are mutually supportive. B1 shows that buy-in is able to persist if and only if it is initially above a threshold, and actions are successfully restricted. A6 shows that treaties can only be upheld so long as there is a sufficiently high degree of buy-in. These two factors combine with B2 to show that buy-in can always persist so long as there is no breaking of treaties.

In Experiments C, however, the presence of high buy-in alone is insufficient in keeping the agents alive, as shown in C2. Any degree of buy-in results in systemic collapse, where the number of players tends to a constant less than the initial configuration. Therefore, whilst there may be a mutually supportive relationship between social contracts and buy-in per Experiments A and B, in a high stakes scenario such as

Downloaded from http://mitp.silverchair.com/isal/proceedings-pdf/isal2023/35/111/2354947/isal_a_00672.pdf by guest on 18 September 2024

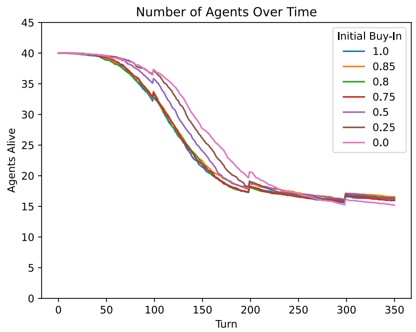

Figure 9: C2: Average agents over time with treaty formation *disabled*

this, there must also be a restricted action space to allow players to survive.

This entails that there is a three-way dependency between the limiting of social actions, the enforcement of social contracts and the degree of popular legitimacy (through individual buy-in): popular legitimacy ensures that social contracts are adhered to, which enforces actions to be restricted, which displays buy-in to the rule of law. Figure 10 is provided to make these interdependencies explicit.

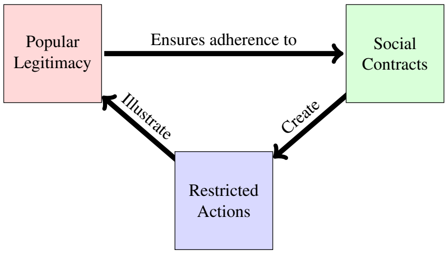

Figure 10: Visualisation of mutual dependencies between buy-in, social contracts and restricted actions

We further assert that the set of interdependencies depicted in Figure 10 is *non-transitive*. Removing any of the blocks in Figure 10 has detrimental effects: whilst popular legitimacy can ascertain adherence to social contracts, it cannot directly restrict actions, as social contracts are required. Furthermore, removing popular legitimacy removes the social adherence to restricted actions, and in effect the rule of law, so social contracts have no power. Finally, removing restricted actions leaves the social contracts with no overall purpose, hindering survivability.

Further Work

Whilst the formation of social contracts are imperative for the *creation* of a social network by serving as 'links', this network could also be used for information transfer. For example, if an agent is confident in their prediction of another agent, the social network could be used to propagate this message to help inform other agents of this knowledge. This would allow for the agent profiles to be developed faster, leading to more rapid clustering.

Furthermore, the simulator tool creates opportunities for a machine learning approach. This could be used to obtain the optimal strategy for agents to survive, both from an individual approach (an 'undefeatable' agent) or a collective approach, where all agents would have the best opportunity to survive.

Summary and Conclusions

In summary, the contributions of this paper are as follows:

- We have proposed a novel scenario that extends the scope of the Hawk-Dove game to n players and allows for the exploration of cooperative survival games with social contracts, social motives and probabilistic buy-in and in the absence of centralised authority. We have also designed, specified and implemented a self-organising multi-agent system to simulate this.
- We have investigated a solution concept for a mixed-mode cooperative survival game and addressed its benefits and pitfalls. This was achieved through a series of 'survival trial' experiments which show that:
  - This game is solvable, providing that there are opportunities for self-organisation, and popular legitimacy is established from the outset.
  - There is a non-transitive, cyclical interdependence between popular legitimacy, social contracts, and restricted actions.

Across all survival trials, it can be seen that an absence of popular legitimacy is detrimental to the survivability of a collective without a centralised authority to provide governance. Therefore, popular legitimacy gives credibility to social contracts that would otherwise be disregarded.

In the absence of social contracts, it is possible to construct a system where all players have total popular legitimacy, as they are never confronted with a reason to question this since no contracts are ever broken. The action space is never restricted, however, so players can justify launching an attack similarly to how they can justify receiving an attack; no rules are ever broken. These are not sufficient conditions for a stable system, however.

In conclusion, in the presence of self-interested players, who would sacrifice others given the opportunity to gain utility, it is imperative to provide a secondary 'meta-game' that is able to reduce this action space and restrict the ability to sacrifice other players. Social contracts are a means of achieving this, however they *must* be underpinned by popular legitimacy when there is no means of punishing players who break a treaty outside of mutually agreed dissolution. Furthermore, these aspects of popular legitimacy, social contracts and restricted action spaces are ring-reinforcing.

Downloaded from http://mitp.silverchair.com/isal/proceedings-pdf/isal2023/35/111/2354947/isal_a_00672.pdf by guest on 18 September 2024

References

Ashby, W. R. (1960). *Design for a Brain: The origin of adaptive behaviour*, pages 44–57. Springer Netherlands.

Bingham, T. (2011). *The rule of law*. Penguin Uk.

Cannon, W. B. (1925). *Bodily changes in pain, hunger, fear and rage: An account of recent researches into the function of emotional excitement*. D. Appleton.

Davis, K. E. (2004). What can the rule of law variable tell us about rule of law reforms? *Mich. J. Int'l L.*

Davoust, A. and Rovatsos, M. (2020). Social contracts for non-cooperative games. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pages 43–49.

de Cesare, S. and Geerts, G. L. (2012). Toward a perdurantist ontology of contracts. In *Advanced Information Systems Engineering Workshops: CAiSE 2012 International Workshops, Gdańsk, Poland, June 25-26, 2012. Proceedings 24*, pages 85–96. Springer.

Grafen, A. (1979). The hawk-dove game played between relatives. *Animal Behaviour*, 27:905–907.

Happe, F. (2021). Fight, flight or fawn: A correspondence with misunderstood women. *TLS. Times Literary Supplement*, pages 22–23.

Hartigan, J. A. and Wong, M. A. (1979). Algorithm as 136: A k-means clustering algorithm. *Journal of the royal statistical society. series c (applied statistics)*, 28(1):100–108.

Houston, A. I. and McNamara, J. M. (1988). Fighting for food: a dynamic version of the hawk-dove game. *Evolutionary Ecology*, 2:51–64.

Mahaney, E. (2022). Trauma, stress, self awareness, managing stress, internal family systems, flooding, fight flight freeze fawn. *Trauma*.

Nafz, F., Steghöfer, H. S. J.-P., Anders, G., and Reif, W. (2013). Constraining self-organisation through corridors of correct behaviour: The restore invariant approach. In *Organic Computing – A Paradigm Shift for Complex Systems. Autonomic Systems, vol 1*, pages 79–93. Berlin, Heidelberg: Springer.

Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge university press.

Prado, M. M. and Trebilcock, M. J. (2021). *Advanced introduction to law and development*, chapter 2. Edward Elgar Publishing.

Pérez-Cirera, V. (2010). *Exploring Game Theory as a Tool for Mapping Strategic Interactions in Common Pool Resource Scenarios*, page 17. Edward Elgar Publishing.

Reinders Folmer, C. (2016). *Social Motives*, pages 886–890. SAGE Publications, Inc.

Reuter, P. (1995). *Introduction to the Law of Treaties*. Routledge.

Reynolds, C. W. (1987). Flocks, herds and schools: A distributed behavioral model. In *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, pages 25–34.

Rezaei, G., Kirley, M., and Pfau, J. (2009). Evolving cooperation in the n-player prisoner's dilemma: A social network model. In *Artificial Life: Borrowing from Biology*, pages 43–52, Berlin, Heidelberg. Springer Berlin Heidelberg.

Rovatsos, M. and Lind, J. (1999). Learning cooperation in repeated games. In *Proceedings of the Workshop on Agents learning about, from and with other Agents (IJCAI-99), Stockholm, Sweden*. Citeseer.

Rovatsos, M. and Lind, J. (2000). Hierarchical common-sense interaction learning. In *Proceedings Fourth International Conference on MultiAgent Systems*, pages 239–246. IEEE.

Scott, M., Dubied, M., and Pitt, J. (2022). Social motives and social contracts in cooperative survival games. In *Coordination, Organizations, Institutions, Norms, and Ethics for Governance of Multi-Agent Systems XV: International Workshop, COINE 2022, Virtual Event, May 9, 2022, Revised Selected Papers*, pages 148–166. Springer.

Serugendo, G. D. M., Gleizes, M.-P., and Karageorgos, A. (2005). Self-organization in multi-agent systems. *The Knowledge engineering review*, 20(2):165–189.

Smith, J. M. (1982). *Evolution and the Theory of Games*. Cambridge university press.

Smith, J. M. and Price, G. R. (1973). The logic of animal conflict. *Nature*, 246(5427):15–18.

Downloaded from http://mitp.silverchair.com/isal/proceedings-pdf/isal2023/35/111/2354947/isal_a_00672.pdf by guest on 18 September 2024