

Strategies for Resource Exploitation

October 25, 2006

Markus Brede^a Fabio Boschetti^b and David McDonald^c

CSIRO CMAR

^a Markus.Brede@csiro.au CSIRO Marine and Atmospheric Research, Bellenden Street
ACT, Australia

^b Fabio.Boschetti@csiro.au CSIRO Marine and Atmospheric Research, Floreat WA,
Australia

^b David.McDonald@csiro.au CSIRO Marine and Atmospheric Research, Aspendale
VIC, Australia

1 Abstract

In a mixed strategy, game-theoretical scenario mimicking the behaviour of fishing vessels competing for a limited renewable resource, agents following either a Collective Intelligence or a purely selfish strategy quickly outperform fully-cooperative teams as well as agents not planning for future action by acting randomly. The stable balance between fully selfish agents and the Collective Intelligence depends subtly on the ratio of instantaneous demand to instantaneously-available resource as well as on the dynamics of the resource itself. This suggests use of ratio of strategies as an indicator of the level of resource exploitation. The Collective Intelligence performance proves to be extremely robust to uncertain information, especially when longer records of historical catch are accounted for.

Key words: Minority Game, Collective Intelligence, Game Theory, Optimisation.

2 Introduction

The growing impact of human activity on the natural environment is making it increasingly difficult for policy makers and managers to achieve sustainable-use objectives. At the root of this difficulty is a poor understanding not only of both renewable resource dynamics and the dynamics of human demand but also the nonlinear, and often counter-intuitive, interplay between human demand and renewable resource response.

2 Introduction

We assert that an improved understanding of these things will be gained by examining the social-ecosystem as a whole, drawing from some of the lessons of complex systems science.

Technically, devising sustainable resource-use policies can be seen as a mathematical problem, and a very difficult one for at least three reasons: a) it is complex, b) it is an inverse problem and c) it is ill-posed. It is a complex problem because the action taken by one agent affects the future behaviour of all other agents and, in turn, the future of the original agent itself. This self-referentiality results in complex cycles of positive and negative feedback loops that make predicting the global outcome of a given policy particularly difficult. It is an inverse problem because for a given global management goal, no method is available to directly determine the action each player needs to take in order to achieve that goal. Finally, this problem is ill-posed because the final management goal is rarely unique; rather it often represents a compromise between several competing sub-goals. In the jargon of applied mathematics, this is a complex, multi-objective optimisation problem.

As a consequence of these difficulties this kind of problem rarely, if ever, has a closed-form analytical solution; rather a considerable amount of trial and error, helped by experience and specific expertise, is often needed. A reasonable way to address this trial-and-error approach is via numerical modelling. This allows one to simulate the counter-intuitive global outcomes arising from the agents' interaction, thereby providing a tool to explore, and hopefully to understand, its inherent complexity. Once a model is built, experiments are very easy to perform, effectively making available to the manager a virtual laboratory in which several policy scenarios can be tested quickly to determine not only the policies which are most suitable for achieving a specific required outcome, but also their resilience and stability to external perturbation.

The present paper describes one such model and is aimed at studying the way a renewable resource is harvested in a competitive human society.

Given renewable resources which are distributed over different locations, and a number of agents with the same potential harvesting capability, we should expect that the best individual harvests are obtained by the agents who access the least-exploited locations; that is, the areas where competition is lowest. This is a generalised version of what is called a Minority Game in the Econophysics literature [1]. Basically we cast the exploitation of a renewable resource within a game-theoretical framework, in which agents aim to predict which areas will be least exploited at the next iteration. Despite its apparent simplicity, the Minority Game displays a number of complex features [2]. The fundamental driver of this complexity is the following self-referential and self-defeating loop [3]: a) a strategy is a winning one if it leads an agent to be in a minority group (a group which accesses the resource in an underexploited area) b) winning strategies are reinforced and are adopted by more and more agents c) agents cease to be in the minority and form the majority group once a winning strategy is adopted by a large enough number of players and d) at that stage the strategy becomes a losing strategy.

In a previous work Boschetti [4] showed that a Collective Intelligence [5] approach can achieve almost optimal resource exploitation even under self-defeating scenarios, in which fully competitive or fully collaborative strategies fail to achieve policy goals. Most important, the Collective Intelligence (COIN) approach allows the team of agents

3 Problem Setting

to adapt to resource variability faster than other methods. The crucial feature of COIN is that agents try to optimise their individual return not directly but indirectly: that is, by approximating how their action increases the return of the entire community. This subtle referential change results in the agents implicitly seeking strategies which account for the overall team performance as well as for their own individual return, which depends on community performance. With a slight abuse of terminology, there is a sense for which a COIN approach attempts to solve a problem (optimal resource exploitation) accounting for both global (community) and local (individual) requirements. The end result is that the agent community manages, very quickly, to spread its harvesting effort proportionally to the resource distribution, despite having no direct information about it, by allowing an equal share of the resource for each agent and a consequent maximum (or optimal) global exploitation.

Boschetti [4] left two main sets of questions unanswered; first, how robust is COIN performance? Is it also successful in a scenario where other strategies are present as well? Is it sensitive to agents' cheating (that is behaving selfishly)? How does noise affect its performance? Second, how does COIN behave under a more realistic (not fully renewable) resource dynamics? In the present paper we address the first set of questions. In particular we allow agents to competitively choose what strategy to adopt and check whether COIN can co-exist with other strategies in an evolutionary stable configuration. As a by-product of this analysis, we suggest that the final balance of strategies can be used as a rough indicator of the ratio between resource availability and resource demand. This indicator has the potential to be of considerable use for policy making.

3 Problem Setting

Following the modelling framework employed in [4] we mimic a fishery scenario in which a fleet composed of a number of competitive vessels accesses a fish stock at different locations. The model can be extended to the exploitation of different renewable resources in a straightforward manner.

We consider our fishery to be divided into a number of fishing zones $i = 1, \dots, z$, each containing a population of Z_i units of fish. N fishing vessels exploit these zones in the following way. Let us assume that $n_i(t)$ of the $N = \sum_i n_i(t)$ vessels exploit zone i at time t and let each of the vessels have a limited transport capacity C_{\max} : that is, C_{\max} limits the maximum catch for a vessel. In this setting, the catch of vessel j that fished in zone i at time step t is given by

$$C_i^j(t) = \delta_i^j \text{Min}(Z_i/n_i(t), C_{\max}), \quad (1)$$

where $\delta_i^j = 1$ if the vessel fished in zone i and $\delta_i^j = 0$ otherwise. This means that all the vessels accessing one zone share the available catch equally. According to Eq. (1) the total catch of all vessels is $C_{\text{tot}}(t) = \sum_j C^j(t)$. The fish stock in each zone is a fully renewable resource. Here, we assume that the fish population always fully recovers before the next fishing haul. From a population dynamics perspective this equates to modelling a large external migration which replenishes the fishery at each time step. Clearly, this is not a realistic assumption and we adopt it here only to discriminate

3 Problem Setting

the effect of the vessels' behaviour from that of the resource dynamics. More realistic resource dynamics will be included in future work.

The fleet's catch depends crucially on how the vessels distribute their effort over the fishing zones. For example, if all vessels access the same zone, they will not be likely to catch at their maximum capacity, as the limited stock of fish available in that zone is shared by the whole fleet. On the other hand, if the vessels distribute themselves proportionally to the stock of fish available in each zone, the maximum allowable catch can be harvested by each vessel. This 'optimal' distribution is difficult to achieve for two reasons. First, the fish resource is unknown. Second, there is no global coordination among the vessels; each vessel must choose where to fish using only information about its previous catch; no information sharing, nor centralised decision making exists which processes the global information and directs the vessels to a pre-assigned fishing location.

As in [4], a vessel i 's decision on where to fish next is stochastic and is based on a probability table $W_j^i(t)$, which represents the vessel's own perception of the world, that is, the perception of the fishing potential of each zone. $W_j^i(t)$ is proportional to the catch vessel j expects to obtain should it access zone i in the next time step. By construction, one has $\sum_i W_j^i(t) = 1$.

A vessel's perception is not fixed in time, but will change according to its past catches. We call a strategy the approach a vessel uses to change perception; that is, the way it accounts for past catches. In the present paper, we examine four different strategies:

- MG (minority game): a fully competitive and selfish approach. This strategy is equivalent to the one commonly used in the traditional minority game (cf. [6]), from which the name comes; a vessel uses only information about its own previous catch to determine its next fishing zone and aims solely to maximise its own return. Information about the past T_{mem}^1 time periods is discounted linearly with time and summarised as (cf. [4])

$$W_j^i(t) = \frac{\sum_{t'=0}^{T_{\text{mem}}-1} C_j^i(t-t') \frac{T_{\text{mem}}-t'}{T_{\text{mem}}}}{\sum_j \sum_{t'=0}^{T_{\text{mem}}-1} C_j^i(t) \frac{T-t'}{T_{\text{mem}}}}. \quad (2)$$

- TG (team game): a vessel bases its decision making on the past overall performance of its team. We assume that a vessel cooperates with the n other vessels following the TG strategy. The catch of the team is shared equally by the members of the team. The basis of modifying the probability table is then not the individual catch, but the discounted team catch per team player.
- COIN (Collective Intelligence): a vessel bases its decisions on the impacts of its historical catches. A vessel's impact is the difference between the overall catch in the zone accessed and the overall catch in the same zone that one would expect had the vessel not gone fishing:

$$E_j^i(t) = \delta_j^i (n_j \text{Min}(Z_j/n_j, C_{\text{max}}) - (n_j - 1) \text{Min}(Z_j/(n_j - 1), C_{\text{max}})). \quad (3)$$

¹In most of the following experiments we set $T_{\text{mem}} = 20$.

4 Competing Strategies

Because part of the catch of vessel j could have been caught by other vessels in the vicinity, the catch of the fleet without vessel j is not necessarily equal to the overall catch of the fleet minus the catch of vessel j (we refer the reader to [4] for the derivation of the expression). Basically this cost function discourages vessels from fishing in areas in which their catch could be caught by others (areas which are either overexploited or already close to maximum exploitation) and instead prompts them to choose areas which would otherwise be underexploited. In COIN the discounted *impact* rather than the discounted *catch* is the basis for the modification of the probability table.

- RAND: a vessel chooses its next fishing zone randomly, without any regard to its historical catches. We implement this by setting

$$W_j^i = \omega_j^i / \sum_j \omega_j^i, \quad (4)$$

where the ω_j^i 's are random numbers chosen from a uniform distribution over $[0, 1]$. On average, the vessel will select each of the available fishing zones with equal probability. This strategy is chiefly used as a baseline to compare with the other strategies.

Previous results [4] have been reported for a scenario in which all vessels follow just one of the above four strategies. These results show that a community of vessels using COIN has two major advantages. First, the total catch C_{tot} can be optimized without negative impacts on the individual catches. Second, the process of adaptation to the fish population distribution is markedly more efficient than for communities using the other strategies. This appears especially important when changes in the distribution of fish over the zones occur.

The scenario described in [4], in which all vessels follow just one of the above strategies, is built on a simplified assumption, which will only hold in a highly regulated world. More realistically, the agents' behaviour will be heterogeneous, with some vessels following COIN, some being team players, others behaving utterly selfishly and, possibly, some moving in a random way. It is thus important to test whether the results in [4] are stable under these conditions. Going one step further, one might also inquire how the competition of agents employing different strategies affects the distribution of strategies in the agent population. Both problems will be addressed in the following section by using approaches from evolutionary game theory ([7, 8, 9]), which have already been applied to the minority game in different contexts (see [10] for an example). The results in [4] augment previous results in [10], and here we explore more in depth the interplay of different strategies and the process which leads the agents to distribute their effort over the different fishing zones.

4 Competing Strategies

4.1 Interactions of Different Strategies

As in [4], we consider a scenario, where 250 units of fish are distributed over $z = 4$ fishing zones. Three of the zones host a total stock of $n_0 = n_1 = n_2 = 50$, whereas

4 Competing Strategies

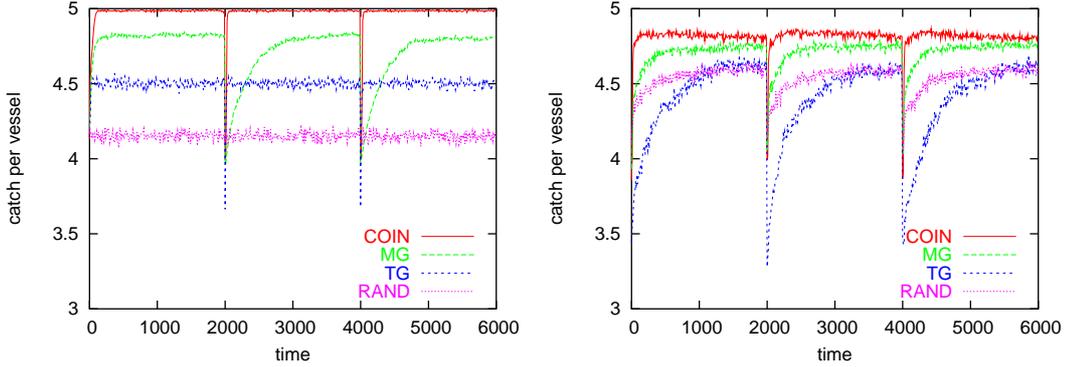


Figure 1: $Z_{z \neq b} = 50, Z_b = 100, b$ changing randomly every $T_{\text{change}} = 2000$ timesteps. Other parameters are $N = 50$ and $C_{\text{max}} = 5$. **(left)** Homogeneous population of vessels following just one strategy. **(right)** Equal proportion of vessels follow the strategies COIN, MG, TG, and RAND. Mutual interactions influence the performance of the strategies. Averaged over 100 independent runs.

a fourth hosts $n_3 = 100$ units of fish. These resources are exploited by 50 vessels, each of which has a maximum fishing capacity of $C_{\text{max}} = 5$ units at each timestep. After every T_{change} iterations, 50 units of fish migrate from the most populated fishing zone to another, randomly selected, zone. The migration timescale T_{change} allows us to control the dynamics of fish migrations. A large T_{change} approximates an essentially static scenario, a small T_{change} describes a faster changing environment for the fishing fleet. We will employ this basic setup for all of the following experiments. Tests of other scenarios indicate that our main results are robust and are not affected by this specific choice.

In Figure 1, we show some simulation data for a migration timescale of $T_{\text{change}} = 2000$. On the right hand side, we show a mixed population with equal numbers of vessels following each strategy and compare it to previous results obtained with a homogeneous population in [4]. From the simulation data, one clearly recognizes the effect of the mixing of strategies. COIN no longer results in optimum catches (as a result of the fact that it no longer distributes the vessels optimally over the fishing zones), but still adapts faster than the other strategies to the migrations. Over time, however, its advantage is reduced as the other strategies gain in relative performance². Similarly, the selfish MG strategy still performs as second best. However, unlike with the homogeneous fish populations, RAND outperforms TG when fish stocks differ across zones. A random assignment of fishing locations results in only a relatively small yield loss after a migration (due to the other strategies over-exploiting some zones), and eventually performs much better than in the scenario with random-players only, displayed in the left hand figure. This is due to the fact that, in this mixed-strategy scenario, it is less likely for randomly-behaving vessels to overcrowd a fishing zone because the other vessels adapt to the changing resource landscape.

²In fact, we find that, however different the transients, for this ratio of strategies asymptotically TG, COIN and MG approach the same catch (data not shown)

4 Competing Strategies

The worst affected by the presence of other strategies is the team game strategy. In the presence of other strategies, team players tend to focus on the zone with highest resource stocks. Consequently, they suffer most heavily from changes in the fish distribution following a migration (cf. Fig. 1 (right)). Also, they no longer react rapidly to a new scenario, only reaching their stationary value after all other players have adapted to the new environment.

Some conclusions can be drawn from this first experiment. First, the interaction of different strategies changes their performance. For example, RAND performs very poorly if it is the only strategy used, but it performs much better when competing against other strategies. Random players force other players to try to adapt to their erratic behaviour, thereby gaining on them. TG vessels adapt only after the other players have settled in and even then barely outcompete RAND players.

Second, despite the ranking of the two best performing strategies, COIN and MG, not being affected by interactions with other strategies, their relative performance is altered to the advantage of MG and COIN no longer achieves optimum performance.

So far, we have investigated the interaction of strategies in fixed proportions (i.e. each strategy is adopted by one quarter of the vessels). However, in a real life situation, a poorly performing vessel would most likely attempt to improve its performance — probably by adopting a different strategy. Pressures to do so could be economic in nature (poorly performing players are forced to exit the market and are replaced by new players) or might result from social interactions and peer pressure (seeing a ‘neighbour’ with a different strategy perform better can induce a player to copy it). In essence this means that the mix of strategies would change over time.

4.2 Dynamics of the Strategy Mix

To investigate the effect of changes in the dynamics of players following each strategy, we allow agents to change their strategies over time. There are several ways of doing this; here we follow an approach developed in ‘Evolutionary Game Theory’ [7, 8, 9]. Each player is assigned a fitness value from its past performance. Independent of whether we model an economic or a socially driven process, we assume that agents are able to spread their strategies proportionally to their fitness values. Better performing strategies will thus spread faster and outcompete those that perform worse. This results in an evolutionary process operating at timescales longer than harvest-dynamics timescales.

More precisely, we consider the following evolutionary process:

1. In a strategy mix $\{n_i\}$ (i.e., n_i players follow strategy i , where we label the four strategies corresponding to the order in which we introduced them), perform $T \gg T_{\text{change}}$ periods of resource harvesting. Each vessel is assigned a fitness, proportional to its overall catch.
2. Each vessel can generate offspring in proportion to its fitness.
3. With a small probability p_{inv} , strategies can ‘reinvade’. That is, the strategy of a randomly-selected vessel is replaced with a randomly-selected new strategy.
4. Iterate 1 to 3 until the achievement of stationary dynamics.

4 Competing Strategies

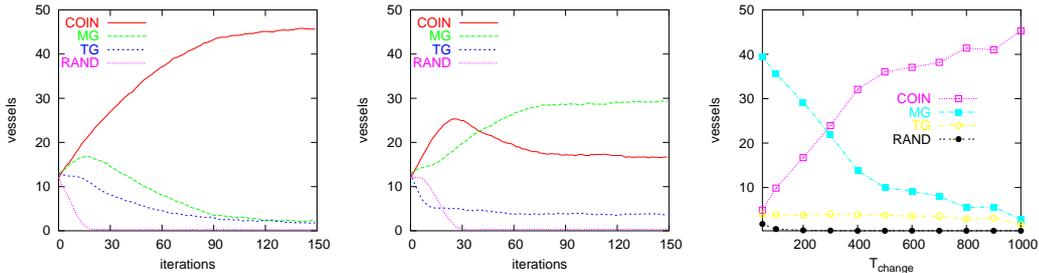


Figure 2: Evolution of interacting strategies. $Z_{z \neq b} = 50, Z_b = 100, N = 50$. Initially an equal proportion of vessels share the four strategies COIN, MG, TG, and RAND. In the first two panels, the x-axis gives the number of evolutionary steps (each corresponding to $10 \times T_{\text{change}}$ fishing time periods), the y-axis gives the numbers of players following each strategy. The left panel represents a static environment with no fish migration. COIN quickly replaces most other strategies, only a small proportion of MG and TG occasionally reinvade, whereas RAND is completely abandoned. The middle panel displays a very dynamic scenario where fish migrate frequently between zones ($T_{\text{change}} = 200$), resulting in a balance between COIN and MG vessels and a few TG. The faster the rate of change, the more MG players, while TG stays approximately constant at around 4. In the right hand panel average numbers of vessels following each strategy depends on the migration timescale. The fitness of individuals was determined by averages over 10000 time steps

An important point to note is that the dynamics work at the players' level, not at a strategy level. In the latter case, the strategy with the best average performance would always win. In the case we model, however, differences among individual vessels play an important role. A strategy might result in a very good average fitness, from a mix of very successful and unsuccessful vessels. A competing strategy, may attract players of similar fitness with an average performance slightly worse than the first strategy. Players following the second strategy, while outcompeted by superior performers of the first strategy, can still outperform poor performers of the first strategy. The result is a balance between numbers of players adopting both strategies. The first strategy can dominate only if its worst performer can outperform the best performer of the second strategy.

Another important point is that the re-invasion probability (this is, the probability that a strategy may be re-adopted after it has been abandoned) defines the level at which populations are statistically meaningful. In our simulations with $N = 50$ agents we used $p_{\text{inv}} = 1/50$. That means, on average at each step of the evolutionary dynamics, one player will be set to a new randomly-selected strategy. We should thus expect that even a strategy that is outcompeted by all others will be followed by a baseline average number of players $n_{\text{min}} = 1/4$ when the dynamics are stationary. Larger values $n_i > n_{\text{min}}$ imply that strategy i can at least occasionally beat one other strategy.

In Figure 2 we summarize the results of these experiments. We started with simulations of a stationary fish distribution (no migration, $T_{\text{change}} = \infty$). As can be seen

4 Competing Strategies

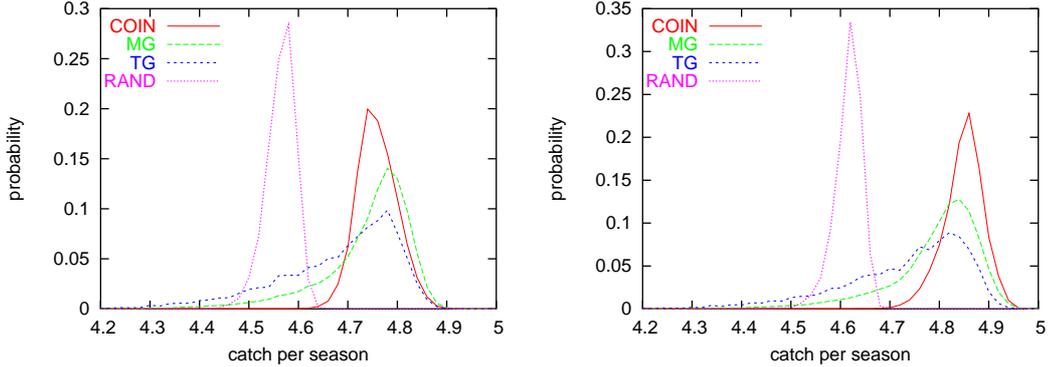


Figure 3: $Z_{z \neq b} = 50, Z_b = 100$. Histograms of the performance of the strategies for two migration timescales $T_{\text{change}} = 200$ (left) and $T_{\text{change}} = 500$ (right). The data were obtained modelling 15 COIN, 30 MG, 4 TG and 1 RAND vessels (optimal setup for $T_{\text{change}} = 200$, cf. Fig. 2) and averaged catches over 2000 time steps. The histograms are constructed from 5000 runs.

in Fig. 2(left), COIN outcompetes the other strategies. Whereas players abandon the random strategy completely, a small number of selfish MG players and a small team of four TG players survives³. This result is interesting in the light of the observation in section 4.1 that in a mixed-strategy setting with equal numbers of players following each strategy, RAND can outcompete TG. This seems to confirm the notion that TG performances improve for smaller teams [11], making the 4-player team more efficient than the 12-player team in section 4.1.

In Fig. 3 we display the average performance of the strategies (in terms of catch distributions) for two migration timescales. In both cases, it is noteworthy that TG and MG players perform very unevenly. For both strategies there is a long tail of worse than average players. In both scenarios —apart from a few strongly underperforming TG and MG players— RAND is outcompeted by all other strategies. Note the shift of the maxima of COIN and TG players from $T_{\text{change}} = 200$ to $T_{\text{change}} = 500$ that marks the transition from MG to COIN dominance.

We also find that the final stationary strategy mix depends on the timescale of fish migrations T_{change} . Figure 2(middle) shows the evolutionary dynamics for a fast migration timescale $T_{\text{change}} = 200$. In this scenario MG players dominate and relegate COIN players to a second place. Finally, in Fig. 2(right) the average populations for all the strategies for $T_{\text{change}} = 50 \dots 1000$ are given. We essentially find that COIN dominates in very slowly changing scenarios whereas MG performs best for short timescales.

The experiments presented so far were all carried out for a balanced resource situation, over which an optimal distribution of vessels would guarantee the maximum possible catch. The performance of the strategies, however, also depends on the ratio of the total vessel capacity NC_{max} to the available catch $\sum_i Z_i$, i.e. the exploitation

³The size of this team does not scale with the number of players, i.e. teams of approximately four TG players form also in a scenario with twice the stock of resources and twice the number of players

4 Competing Strategies

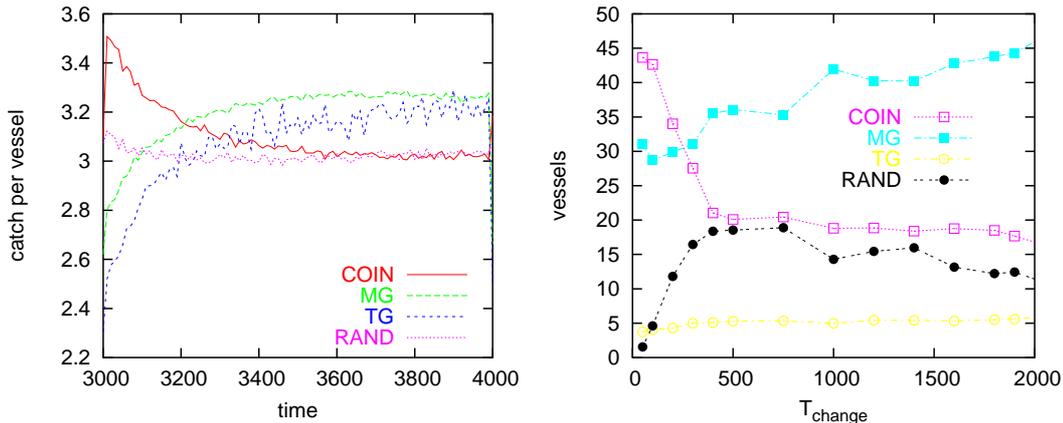


Figure 4: $Z_{z \neq b} = 50$, $Z_b = 100$, $N = 80$ and $T_{\text{change}} = 1000$. In the left hand panel time evolution of the average catch per vessel is displayed for 43 COIN agents, 17 MG agents, 4 TG agents and 16 RAND agents. Note that COIN does not start at a peak after the migration at $T = 3000$, but benefits from very fast adaptation. The panel on the right displays the proportion of COIN, TG, MG and RAND as a function of T_{change} .

ratio

$$r = \frac{NC_{\max}}{\sum_i Z_i}. \quad (5)$$

This ratio defines the amount of vessel space per unit of available fish, given an optimal distribution of the vessels over the fishing zones. Clearly, for very large ratios, most fishing zones will almost always be overcrowded. Under these conditions, individual contributions very seldom lead to an impact on the global catch of the fleet and the behaviour of COIN is expected to become more erratic.

In Figure 4 we summarize some simulation results for a fleet of $N = 80$ vessels over the same fishery scenario displayed above.

The average catch per vessel for different strategies differs from the previous scenario. In this overcrowded scenario (the ratio is $r = 80 \times 5/250$) COIN can dominate only by adapting faster than all the other strategies. COIN adapts very quickly to a migration, thereby benefitting from it compared to the other strategies. With time, however, vessels with different strategies gradually adapt to the new fish distribution and COIN vessels are pushed away (all resources tend to be overcrowded). With longer migration time scales, thus with less migrations to take advantage of, COIN vessels cannot achieve high impacts and their performance approaches that of a vessel adopting RAND. Also, depending on the migration, RAND has the potential to perform well in comparison to TG.

In Figure 5 we show the results of simulations in which we systematically altered the exploitation ratio. This was done by keeping the available units of fish in the zones fixed, but changing the number of vessels exploiting them. In strongly underexploited scenarios, the importance of a balanced arrangement of the vessels over zones

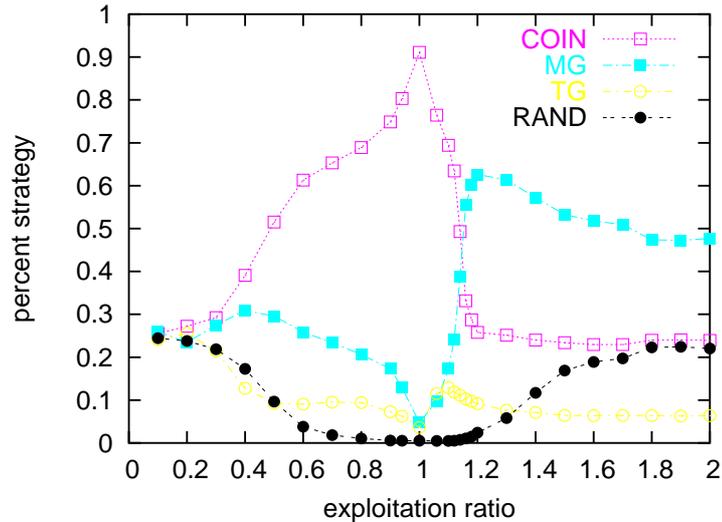


Figure 5: Dependence of the mix of strategies on the exploitation ratio. $T_{\text{change}} = 1000$ and the fish distribution is modelled as above. Every data point is averaged over 100 runs.

is low. Many otherwise suboptimal configurations yield satisfactory average catches; all strategies fare rather evenly. As the resource availability gets tighter and the actual arrangement of players over the zones becomes more important, the fraction of COIN players rises until it reaches its maximum for an exploitation ratio of around one. Past this point competition for scarcer and scarcer resources increases. As a result COIN players are driven away, because they are unable to achieve an impact on the global catches. Finally, for very large exploitation ratios, fluctuations, i.e. movements of vessels from one zone to another that are unrelated to changes in the resource distribution, become more and more frequent. Because COIN players can no longer achieve impacts, their behaviour approaches that of RAND players; at this stage, the fraction of COIN plus RAND players almost equals the fraction of those adopting the MG strategy.

5 COIN and Uncertain Information

As observed in the introduction, an agent adopting the COIN strategy requires more information about its environment than an agent that behaves completely selfish (MG). Whereas an MG agent only requires a record of its own historical catches to determine its next action, a COIN agent needs to know about its impact history. Calculating the impact, however, requires information on how many other vessels accessed the same fishing zone (n_i) and the size of the fish-stock available in that zone (Z_i)⁴. This additional information requirement could leave the COIN strategy vulnerable to false information.

⁴Actually, in our somewhat artificial model setup $Z_{i \neq b} = 50, Z_b = 100, N = 50$, an agent can deduce its impact from its own catch. This, however, is not true for more general fish distributions.

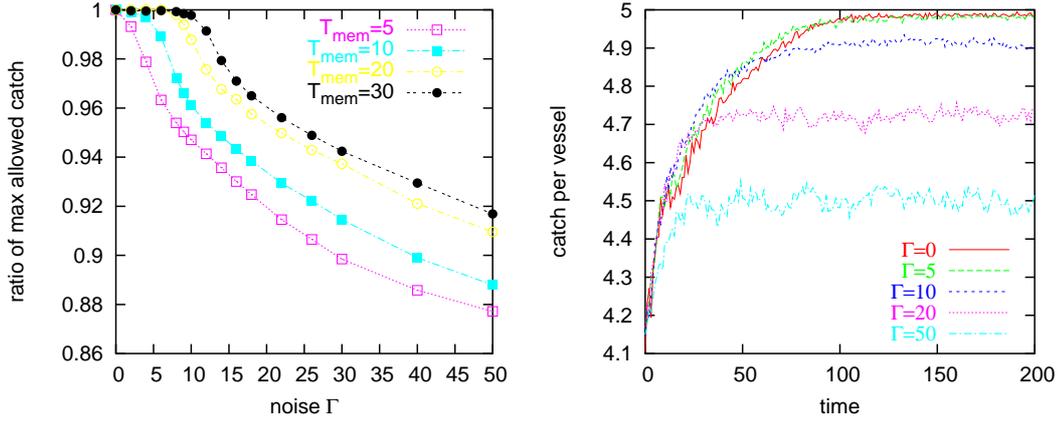


Figure 6: The panel on the left displays the effects of uncertainty on the performance of COIN. Different curves correspond to different lengths of the history used for decision making (T_{mem}). Remarkably, the performance of COIN is not affected strongly, even for noise levels comparable to the maximum allowed catch (no impact larger than the maximum catch is possible). The panel on the right displays averaged time series for COIN and different levels of uncertainty for $T_{\text{mem}} = 20$. For comparison, with the same setup but no uncertainty, MG yields $4.77 \pm .06$, TG 4.5 ± 0.013 and RAND 4.15 ± 0.018 .

Uncertainty in the information required by COIN stems from two sources:

1. Uncertainty in the amount Z_i of resource available in the fishing zone. Information about Z_i is hard to obtain for vessels in real life scenarios and will mostly be based on a very error prone estimate.
2. Uncertainty about the number of other vessels in the same zone. Reasons for the lack of such information might be bad weather or communication problems.

In the following we include uncertainty into our model without discriminating its sources. After determining a vessel's impact as described in section 3 according to Eq. (3), we add to it a random number Ξ_t drawn from a uniform distribution over $[-\Gamma, \Gamma]$, i.e.

$$E_j^i(t) \rightarrow \text{Max}(E_j^i(t) + \Xi_t, 0), \quad (6)$$

where the parameter Γ allows us to tune the magnitude of the uncertainty. Next we calculate the average catches of a fleet comprised of only COIN vessels, see Fig. 6.

We find that COIN is remarkably robust to noise. Even by applying uncertainty levels much larger than the noise-free maximum possible impact, the COIN performance is not strongly affected. Only for noise levels $\Gamma \approx 2C_{\text{max}}$ a gradual deterioration sets in. However, even noise levels $\Gamma \approx 10C_{\text{max}}$ still allow COIN players attain much better performance than RAND players.

Why is COIN so remarkably robust? The reason appears to reside in the strong impact of history in the decision making process. When choosing its next fishing zone,

an agent effectively assesses a weighted average over its last $T_{\text{mem}} = 20$ catches. This means that, even though individual pieces of information might be heavily corrupted, the fluctuations are smoothed out in the final calculation, cf. Eq. 2. Only when the uncertainty is large or biased enough to impact the averaging process does noise considerably affect COIN’s performance, cf. Fig. 6 (left).

The inclusion of uncertainty in the model affects COIN’s optimal memory length. Without noise, memory length needs to compromise two different trends. It has to be long enough to enable COIN to ‘learn’, but not too long, in order to allow for fast adaptation to migration. Large uncertainty shifts the balance towards longer memory requirements.

6 Summary and Discussion

Originally, COIN was conceived as an numerical optimisation tool to address engineering problems involving the coordination among a large number of components [5]. In [4] we showed that the COIN algorithm is simple enough (and we simplified it further) to be used by real human agents using only primary-school level calculations. Clearly, the real world adoption of a technique depends not only on its ease of implementation but also on its robustness. In this application, robustness needs to be assessed against uncertain or insufficient information and against explicit and illicit (cheating) competition.

In this paper, we addressed these challenges by testing COIN against three different approaches in a mixed strategy environment, under different scenarios of resource variability and with increasing levels of data uncertainty.

First, we learned that the performance of a strategy changes considerably in a mixed strategy scenario: that is, when it needs to compete directly against other strategies. In these cases, we found that random behaviour is not viable and that only very small teams of cooperating players can survive. The dominant strategies are always either the Collective Intelligence (COIN) or a fully selfish one (minority game, MG). Which of these performs better depends on two parameters: a) the exploitation ratio; that is, the ratio between available resource and demand and b) the frequency of variation in the resource distribution. In balanced scenarios with an exploitation ratio of around one, COIN dominates for intermediate and large migration timescales, while MG prevails in fast-changing scenarios. For overexploited systems, the converse is found to hold.

Our second main result is that COIN proved to be extremely resistant to incorrect information, which we modelled as noise in the determination of vessels’ impacts on the overall community performance. The performance of COIN subject to noise is strongly dependent on the length of the record of past catches used in the decision making process. In the context of the model, accounting for past information corresponds to a weighted average over past results, hence a longer memory tends to provide stronger resistance to noise. The problem-dependent optimal history length is determined by a trade-off between a) the minimum length of time required for COIN to learn, b) the requirement for fast adaptation to change in the resource distribution and c) the need to account for incorrect information. In a very uncertain environment, where impacts can only be determined approximately, allowing for long histories of catches is expected

to achieve a better result.

As a by-product, our experiments suggest that the above mentioned dependency of the stable balance between COIN and MG on the exploitation ratio, may be employed as an indication of the exploitation ratio itself. Poor performance of vessels following a COIN strategy, effectively behaving randomly, may be an indicator of an overexploited resource. This flags the possibility that COIN vessels could not only improve resource exploitation, under favourable circumstances, but also act as a monitor of the efficiency of a management policy. We plan to explore this option further in our future work.

References

- [1] Y., Zhang, “Modeling Market Mechanism with Evolutionary Games,” *Europhys. News*, **29**, 51, (1998).
- [2] R., Savit, R., Manuca, and R., Riolo, “Adaptive Competition, Market Efficiency, and Phase Transitions,” *Physical Review Letters*, **82** (10) (1990) 2203–2206.
- [3] D., Batten, “Are some human ecosystems self-defeating?,” *Environmental Modelling and Software*, (in press)
- [4] F. Boschetti, “Improving resource exploitation via Collective Intelligence by assessing agents’ impact on community outcome,” *submitted to Ecological Modelling?* (2006).
- [5] D. Wolpert, K. Wheeler and K. Tumer, “Collective Intelligence for Control of Distributed Dynamical Systems,” *Europhys. Lett.* **49**(6) (2000) 708.
- [6] D. Challet and Y.-C. Zhang, “Emergence of Cooperation and Organization in an Evolutionary Game,” *Physica A* **246** (1997) 407-418.
- [7] M. A. Nowak and R. M. May, “Evolutionary Games and Spatial Chaos,” *Nature* **359** (1992) 826–829.
- [8] M. A. Nowak and K. Sigmund, “Chaos and the Evolution of Cooperation,” *Proc. Natl. Sci. USA* **90** (1993) 5091–5094.
- [9] R. M. Axelrod, *The Evolution of Cooperation* (Basic Books, New York, 1984).
- [10] N. F. Johnson, P. M. Hui and T. S. Lo, “Self-Organized Segregation within an Evolving Population,” *Phys. Rev. Lett.* **82** (1999) 3360.
- [11] D. Wolpert and K. Tumer, “Optimal Payoff Functions for Members of Collectives,” *Advances in Complex Systems* **4**(2/3) (2001) 265-279.
- [12] D. Wolpert, K. Tumer and D. Bandari, “Improving Search Algorithms by Using Intelligent Coordinates,” *Phys. Rev. E* **69** (2004).