An information-based adaptive strategy for resource exploitation in competitive scenarios

Fabio Boschetti and Markus Brede
CSIRO Marine Research - Marine and Atmospheric Research, Australia
Email: [Fabio.Boschetti,Markus.Brede,]@csiro.au

Abstract

Given an exploitation problem, in which a number of agents compete for a limited renewable resource, the optimal harvesting strategy depends on the ratio between resource availability and exploitation effort. For scarce resource a purely competitive, greedy strategy outperforms a more collaborative approach based on the Collective Intelligence, while for more abundant resource the opposite holds. The rationale for this behaviour lies in the amount of information each strategy is able to provide and a combined strategy is possible according to which agents choose dynamically the most informative strategy according to a minimum entropy criteria. This approach, which provides best performance for both under and over exploited scenarios, can be used to monitor the resource status for management purposes and is effective in both centralised and decentralised decision making.

1 Introduction

Several natural phenomena can be cast within the framework of competition: living beings compete for survival, species compete for niches, humans compete for financial resources and, adventuring ourselves into more questionable conjectures, cultures compete for supremacy, ideas compete for our attention [1], and natural selection is even suggested to act the cosmos [2]. These ideas are linked to different versions of Darwinism: *best competitors will survive* and thus dominate in the long run. Competition is thus linked to the concept of optimisation: via competition agents will 'improve' at required tasks.

This framework is intuitively appealing, so much so that it can be easily exported to man-made problems and in particular to engineering and numerical optimisation. Either explicitly, as in the case of Evolutionary Computation and Particle Swarm Optimisation or implicitly, as in standard gradient-based techniques, many methods employed for computer-driven optimisation rely on some sort of competition between components of the optimisation algorithms.

However, competition can express itself at different levels: a short–term winning strategy may fail in the long-term and a locally winning strategy may fail globally. Game theory, as well as everyday experience, shows that the outcomes of competition can be much less intuitive and at times so unintuitive to force us to turn Darwin's conjecture upside down: *the ones who survived must have been the best competitors.* Optimisation practitioners are well aware of these problems, which manifest themselves in the challenge in finding global solutions among local ones.

When the process to optimise can itself respond to the optimisation and adapt to it, the dynamics can be even more complex, in which case modelling may be the only avenue for us to unravel, predict or control the process under analysis.

In recent years we studied one such system. We modelled a fishery including fishing vessels harvesting several fishing zones of constant resources [4]. Vessels compete by aiming at underexploited areas, thereby avoiding sharing the limited resource with the majority of other vessels. This is an example of a Minority Game [3,8] and in the Game Theory literature it is known that this apparently simple process can generate complex dynamics. Next, we included a resource dynamics by simulating population growth in the target species [6,7]; this imposed on the vessels the additional complication of accounting for the evolution of the resource abundance in response to fishing. Finally, we included fleet managers and resource managers, who can impose constraints of the fishery either by regulation [7] or by centralised fleet control [8]. Each step increased considerably the overall complexity of the problem.

For each scenario, we studied how different fishing strategies perform both in a single strategy setting (in which all vessels in the fleet adopt the same strategy) and in an evolutionary economic setting (in which strategies spread in the population according to their past performance). Of particular interest is the comparison between the following strategies: a purely competitive approach, which we call MG in this paper (since this is the strategy normally adopted in Minority Game studies), in which agents aim to optimise their individual return, and a more collaborative approach, called Collective Intelligence (Coin in the rest of the paper), in which agents try to optimise their *impact* on the fleet return. It seems reasonable to consider MG as the 'null hypothesis' against which we test the Coin performance.

One of the most important results of our previous work is that the effectiveness of Coin depends on the ratio between available resource and fishing effort. In particular, the transition from Coin to MG dominance in the fishing fleet coincides with the transition from under-exploited to over-exploited resource status. In [6] we have speculated that this transition could, in principle, be used by a resource manager to detect the level of resource exploitation and decide on possible intervention. The purpose of this work is to explore this idea further and discuss some steps towards a possible implementation.

We start by setting the problem and describing the Coin approach. We then summarise the results from our previous work most relevant to this paper and analyse our new results. We conclude with a discussion of the possible future development of this research.

2 The model

In this paper we report on a model of a simplified, non-spatially explicit fishery, although the model could easily be extended to the exploitation of other renewable resources. We imagine N fishing vessels $n = 1..N$ and Z fishing zones in which an amount $Fish_z, z = 1..Z$ of resource is available. The vessels do not have information about the global distribution of $Fish_z$ and decide where to direct their effort according to the discounted returns of past catches in the different fishing zones (see [4] for details).

Each vessel has a maximum allowed quota (which alternatively can be interpreted as a limited fishing capacity). At each fishing period, a vessel targets a single fishing zone and the resource available in that zone is shared equally among all vessels targeting it. Thus, the catch of a vessel *n* is given by

$$Catch_n = Min(Fish_{zone_n} / Fleet_{zone_n}, Quota) \qquad (1)$$

where $Catch_n$ is the amount of fish caught by vessel *n*, $zone_n$ is the fishing zone chosen by vessel *n*, $Fish_{zone_n}$ is the amount of fish available in $zone_n$, $Fleet_{zone_n}$ is the number of vessels which chose to fish in $zone_n$, with which vessel *n* has to share the available resource. We do not model fishing costs (navigating to the zones, equipment renting/buying, etc) though these could easily be included if needed.

The total catch of the fleet is obviously given by the sum of each vessel's catch,

$$TotalCatch = \sum_n Catch_n \qquad (2)$$

The maximum possible catch of the entire fleet is given by either the total amount of resource in the fishery or the sum of the maximum allowed quota per vessel, if the resource is abundant:

$$MaxCatch_{Fleet} = Min(\sum_z Fish_z, N * Quota) \qquad (3)$$

Notice that, because each vessel has a maximum allowed quota, we have:

$$TotalCatch \leq MaxCatch_{Fleet} \qquad (4)$$

that is, unless the vessels spread their effort wisely, the fleet may not be able to catch to its full capacity.

3 The Minority Game

Given a resource distributed over different zones and a number of vessels with the same fishing capability, we expect that the vessels which access the least-exploited zones will undergo less competition and thus share the local resources with the least number of competitors, catch the most fish and consequently perform the best. Basically the problem can be cast within a game-theoretical framework, in which agents aim to predict which areas will be least exploited at the next iteration. This is a generalised version of what is called a Minority Game in the Econophysics literature [3,9] (see also http://www.unifr.ch/econophysics/minority/).

Despite its apparent simplicity, the Minority Game displays a number of complex features, which arise directly from the following self-referential loop [5]: 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) a winning strategy tends to be imitated and thus is adopted by more and more agents, c) consequently a winning strategy will soon be adopted by the majority of the agents and d) at that stage it becomes a losing strategy.

It is easy to see [4] that, given sufficient resources, optimal resource exploitation in a Minority Game scenario can be achieved by spreading the agents' harvesting effort proportionally to the resource distribution. However, achieving this optimal allocation without centralised control is not trivial. Most of the Minority Game literature assumes fully competitive agents which always try to maximise their return. In this scenario, the self referential loop described above results in oscillations around the optimum distribution; these oscillations correspond to a waste of resource [4] and never dissipate [3].

4 Collective Intelligence

The resource exploitation problem described above can be seen as a decentralized optimisation task, in which we aim to allocate the fishing effort of *N* vessels proportionally to the resource $Fish_z$.

From an optimisation perspective, we could choose to optimise two quantities: a) the 'private' return for each vessel and b) the 'global' return of the entire fleet. Optimising either of them in isolation is known to be sub-optimal, either because the average catch is sub-optimal or because the distribution between individual agents is not balanced. In the previous section we discussed the outcome of optimising the 'private' return. The 'global' return is optimised in what is called a 'team game' in which each agent obtains an equal share of the global catch, independent of its individual performance; it is known that this approach reaches an optimal exploitation distribution only for very small problems [12].

The Coin approach [12,14,15,16,17] can be seen as a sort of compromise between these two. It is based on each vessels aiming to optimise a differently defined 'private' return, which is a function of how each vessel influences the 'global' return. In particular, each vessel tries to maximise its *impact* on the global return, where the impact is defined as the difference between the global catch for the fleet and the catch that the fleet *would have caught* had the vessel not being present.

Let's assume we have a fleet of *N* vessels. We define the impact of vessel *n* on the overall fleet as the difference between the global catch of the fleet and the catch the fleet *would have caught* if vessel *n* had not gone fishing:

$$IMpact_n = TotalCatch - TotalCatch^{-n} \quad , \qquad (5)$$

where the superscript '-*n*' refers to values calculated for the fleet *without* vessel *n*. Notice that, because the catch of each vessel is limited by physical constraint or quota restrictions, we may have: $TotalCatch^{-n} \neq TotalCatch - Catch_n$, which makes Eq (5) non-trivial.

Thankfully, we do not need to calculate Eq (5) exactly (which would be impossible in most cases); the Coin literature shows that a remarkable simplification is effective, thereby the impact can be approximated by simply removing vessel *n* from the fleet and, leaving everything unchanged, approximating $TotalCatch^{-n}$. Let's suppose vessel *n* targeted fishing zone *z*. Clearly it could not have any impact on any of the remaining zones, so we need to concern ourselves only with zone *z*. It thus follows that:

$$\begin{aligned} &IMpact_n = TotalCatch - TotalCatch^{-n} = \\ &Fleet_z * Min(Fish_z / Fleet_z, Quota) - \\ &Fleet_z^{-n} * Min(Fish_z / Fleet_z^{-n}, Quota) \end{aligned} \quad (6)$$

where $Fleet_z$ is the size of the fleet which targeted zone *z* and the possible catch per vessel is constrained by the quota.
Since $Fleet_z^{-n} = Fleet_z - 1$, we have thus have:

$$\begin{aligned} &IMpact_n = Fleet_z * Min(Fish_z / Fleet_z, Quota) - \\ &(Fleet_z - 1) * Min(Fish_z /(Fleet_z - 1), Quota) \end{aligned} \quad (7)$$

This is the equation used in our calculation (see [4] for more details).

Importantly, Eq 7 can be calculated using only local information about the area targeted by the vessel, without any need for global information.

5 Previous Results – Evolutionary Game Theory

In [4] we described a number of virtual experiments mimicking a fishing fleet operating in areas of different fishing capacity, but fully renewable resource. In that scenario, Coin provides optimal catches for the fleet while, at the same time, each individual vessel also maximizes its own profit; this shows that although, in principle, Coin requires fishing vessels not to act greedily, no individual sacrifice is imposed on the vessels by achieving the common goal. We also showed that a fleet following a Coin strategy adapts much faster to change in resource distributions, promising increased benefits over standard approaches in volatile environments.

These results were extended to scenarios in which the resource dynamics were explicitly modelled [6]. We also explored the trade-off between long and short term planning by providing vessels with some knowledge of the time evolution of the resource, thereby allowing them to plan their fishing behaviour in light of predicted long-term resource behaviour [7].

Most relevant to the present work, in [7] vessels were allowed to choose dynamically one of four possible strategies: MG, Coin, team game and random (each vessel decides randomly where to go fishing at each fishing period). The choice of strategy was carried out according to each strategy's past return, in a typical evolutionary economics scenario [10]. The insight gained from these experiments is that the balance of vessels choosing the different strategies depends crucially on the resource availability. We summarise the results via the sketch in Figure 1, while we refer the reader to the original paper [7] for more details and numerical results.

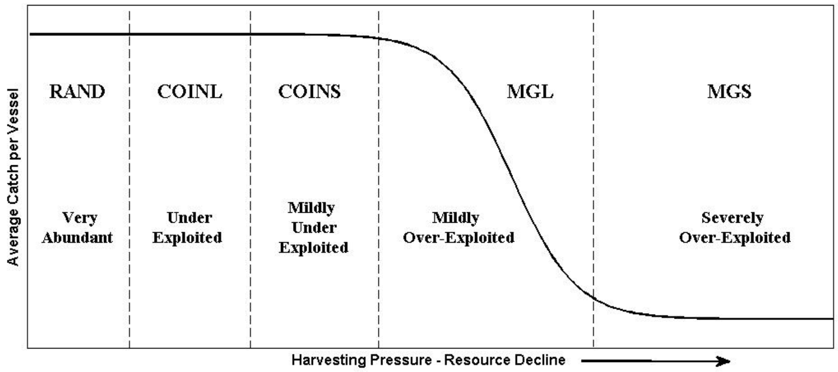

Figure 1. Optimal strategy for different levels of resource abundance. Average catch per vessel (Y) as a function of resource availably (X). The vertical dashed bars indicate which strategy is optimal for different resource abundance; RAND= random choice of fishing zone, COINL= Coin with long term projection of resource dynamics, COINS=Coin without projection of resource dynamics, MGL= greedy strategy with long term projection of resource dynamics, MGS= greedy strategy without projection of resource dynamics

In Figure 1 the X axis gives an indication of resource abundance in relation to the fishing capacity of the overall fleet, ranging from strongly under-exploited (catches have little effect to the resource status) to strongly over-exploited (the fleet can fish more than the resource can sustain). The Y axis represents the modelled average catch per vessel (notice that the catch plateaus for abundant resources because of the limit in vessel fishing capacity due to physical limitations or quotas). For different levels of resource abundance, Figure 1 shows which fishing strategy dominates in the mixed-strategy fleet, that is, which strategy provides the best catches for that specific resource state. Starting from the right-hand side and moving leftward we can mimic the process of resource decline, as typically happens in real world fisheries; Figure 1 then shows how the optimal fishing strategy changes at different stages during the worsening resource state. When the resource is very abundant (basically unlimited in comparison to the fleet capacity) there is no need to put much effort in choosing where and how to fish, and consequently a random strategy performs very well. When the resource is abundant, but not unlimited, a Coin strategy accounting for long term resource dynamics (CoinL) is best; cooperation among vessels and long-term planning allows for optimal resource exploitation. From now on, further resource reduction favours more and more competitive fishing strategies; when the resource is limited but not over-exploited Coin with short-term goal fares best (CoinS), while when the resource gets overexploited fully competitive, greedy behaviours become optimal and thus dominant in the fleet (MGL, greedy behaviour with long term projection and MGS, greedy behaviour with no long term projection).

Figure 1 is interesting for a number of reasons. If we interpret it within a co-evolution framework, in which the resource responds to the level of exploitation and the level of exploitation to the resource status, the figure shows that when a resource is abundant, cooperation and long term planning, as implied by the CoinL is not only beneficial, but actually optimal. However, if the resource starts declining, either by its own internal dynamics or by external pressure, then more and more greedy exploitation patterns become optimal and consequently spread in the population. Cooperation is no longer viable and the community itself becomes more greedy and selfish. Effectively, the system falls into a Tragedy of the Commons situation [11], in which resources are scarce but it is still economically rational to keep on exploiting the limited resource competitively. At this stage, resource recovery is possible only if the fishery becomes uneconomical before it crashes completely. Within this modelled framework, hardship discourages, rather than fosters, collaboration.

Important for our discussion is the transition from Coin to greedy strategies (MG). As shown in Figure 1, this transition coincides with the transition between under-exploited and over-exploited resources. In principle, this suggests that this transition could be used as an indicator of the resource status. Clearly, this information would be of great value in resource management.

Unfortunately, there are several reasons why monitoring the balance between Coin versus MG would be unfeasible in practise. Apart from issues related to the adoption of the approach in the real world, one obvious hurdle would be convincing vessel managers to follow the evolutionary economics principle in order to adopt a strategy (Coin or MG) which may be economically viable only under specific resource conditions (under or over-exploited) for the purpose of detecting those conditions themselves, trusting that this approach will provide best returns in the long-term.

An option, however, may be available: rather than asking vessels to choose dynamically whether to follow Coin *or* MG according to past returns, we may let them use information from both Coin *and* MG at each period, and choose dynamically which one is more likely to offer more accurate information. In order to do so, we need to understand the reasons behind the different performances of the strategies under different resource conditions.

6 Tuning the fishing strategy to the resource status

In this section we describe how fishing vessels could use information from both MG and Coin dynamically. In order to do so we first need to explain how information from past catches is processed.

In the model used in the previous experiments, the fishing strategy for vessel *n* can be summarised as follows:

1) At fishing period *t*, store the *profit* from that period. Here we define $profit_n = catch_n$ from Eq (1) for MG and $profit_n = IMpact_n$ from Eq (7) for Coin.
2) Clearly, vessel *n* obtained a profit (and thus information) at *t* only from the zone *z* it targeted. Define $W^t_{n,z} = \delta^t_z profit^t_n$, where $\delta^t_z = 1$ if the vessel fished in zone *z* at time *t* and $\delta^t_z = 0$ otherwise.
3) Store the value of $W^t_{n,z}$ for the last *T* periods. The value of *T* can be seen as a characteristic memory time for the planning process. $W^t_{n,z}$ now contains

information about all zones z which have been visited in the last T fishing periods.

4) In order to account for non-stationarity in the predictions, discount the past catches. In this work we use a linear discount rate $W^t_{n,z} \to W^t_{n,z} \frac{T-t+1}{T}$.
5) For each zone z, sum the discounted weights $W^t_{n,z}$ over the T fishing periods $P^z_n = \sum_{t=1}^{T} W^t_{n,z}$ and normalise over the Z fishing zones $P^z_n \to P^z_n / \sum_z P^z_n$. P^z_n is now interpreted as the probability that zone z provides the best profit for vessel n at the next fishing period and is used, stochastically, by vessel n to choose the next zone to target. Should a zone have $P^z_n = 0$ we assign $P^z_n = \varepsilon$ and renormalize, where ε is a parameter chosen to prevent certain fishing zones to be 'forgotten'. Alternatively, ε can be seen as a factor which allows exploitation of yet unexplored zones.

In order to appreciate the difference between Coin and MG it is important to understand the difference between $IMpact_n$ and $Catch_n$. If vessel n chooses a zone which is under-exploited, that is a zone where vessels can fish at their maximum capacity (this is clearly a profitable action to take), $IMpact_n$ in Equation (7) will be large, since vessel n will be able to catch its share without affecting the other vessels' catch. In this case $IMpact_n = Catch_n$. If vessel n choose a zone which is over-exploited, that is a zone where vessels cannot fish at their maximum capacity (this is clearly an action to avoid), $IMpact_n < Catch_n$ since all vessels need to share the limited amount of fish available and no extra catch is possible. Had the vessel not gone fishing, its catch $Catch_n$ would have been collected by other vessels which did not succeed in fishing at maximum capacity.

In general, we thus have $IMpact_n \le Catch_n$. Since $Catch_n$ is the contribution of vessels n to the overall catch, we see that the impact a vessel can exert on the community is always less or equal to its contribution. The aim of Coin is thus to maximise impact, that is to make the impact as close as possible to the contribution. The key to its success is its directing vessels where their contribution really 'makes a difference' on the overall catch.

Let us consider a mildly under-exploited scenario, under which we know Coin outperforms MG. We can expect that whichever zone vessel n chooses it will be able to obtain a certain catch, that is $Catch_{n,z} \neq 0, \vee z, z \in Z$. However, it is reasonable to expect that there may be zones for which the impact is zero $\exists z, z \in Z, IMpact_{n,z} = 0$, or at least for which $\exists z, z \in Z, IMpact_{n,z} < Catch_{n,z}$.

Let us now consider an over-exploited scenario, under which we know MG outperforms Coin. Here as well we can expect that whichever zone vessel n chooses it will be able to obtain a certain catch, $Catch_{n,z} \neq 0, \vee z, z \in Z$. However, it may now happen that vessel n is not able to exert any impact in any zone. This may happen in the case in which the fleet would be able to catch all available resource in each zone, without the contribution of vessel n. In this case we may have $IMpact_{n,z} = 0, \vee z, z \in Z$.

When this happens Coin is not provided with any information on the profit from its fishing period and thus has to rely only on information from past catches in order to direct future fishing effort. Crucially, should this happen for *T* iterations, Coin would be left with no information at all and its behaviour would become random.

From the above analysis it is reasonable to suggest that the strategy which performs better is the one which provides most *unequivocal* information about the likelihood of future catches. This is not surprising of course, and suggests an alternative approach:

1) At each fishing period *t*, compute both $Catch_n$ and $IMpact_n$
2) Store both values in two separate tables ${}^{MG}W^t_{n,z}$ and ${}^{C}W^t_{n,z}$, where *MG* stays for the greedy strategy MG and *C* for Coin.
3) Discount and normalise both ${}^{MG}W^t_{n,z}$ and ${}^{C}W^t_{n,z}$ and obtain the probabilities ${}^{MG}P^z_n$ and ${}^{C}P^z_n$ for both MG and Coin as at point 5 above.
4) Define the measure of the information content of ${}^{MG}P^z_n$ and ${}^{C}P^z_n$ as their entropy ${}^{MG}E_n = -\sum_z {}^{MG}P^z_n \log({}^{MG}P^z_n)$ and ${}^{C}E_n = -\sum_z {}^{C}P^z_n \log({}^{C}P^z_n)$; the choice of the entropy as a measure of information content is a cornerstone of information theory [18].
5) In order to decide where to fish next, choose either ${}^{MG}P^z_n$ or ${}^{C}P^z_n$ depending on which one has lowest entropy. In the rest of the paper we call this strategy *MaxInfo*.

In Figure 2a we show the performances of the Coin, MG and *MaxInfo* strategies as a function of relative resource abundance. This is calculated as the ratio between available resource and the maximum fishing capacity of the fleet, ranging from *.5* (resource equal to half the fishing capacity) to *2* (resource equal twice the fishing capacity). We modelled 50 fishing vessels targeting 2 fishing zones with uneven resource distribution, $Fish_2 / Fish_1 = 1/3$.

The results are given for runs of 100 fishing periods, and include the initial transient during which the strategies train themselves starting from initially uniform $W^t_{n,z}$. On the Y axis we plot the average fishing efficiency, that is the ratio between the average and the maximum possible catch per vessel. All results are averaged over 200 runs. In the figure, red lines show the Coin results, blue lines show the results of the MG and black line shows the results of the *MaxInfo*.

The comparison between Coin and MG confirms the result in Figures 1. For abundant resources (*X>1*) Coin performs better while for scarce resource (*X<1*) MG gives better catches. The transition from MG to Coin best performance occurs for values of the relative resource abundance slightly less than 1.

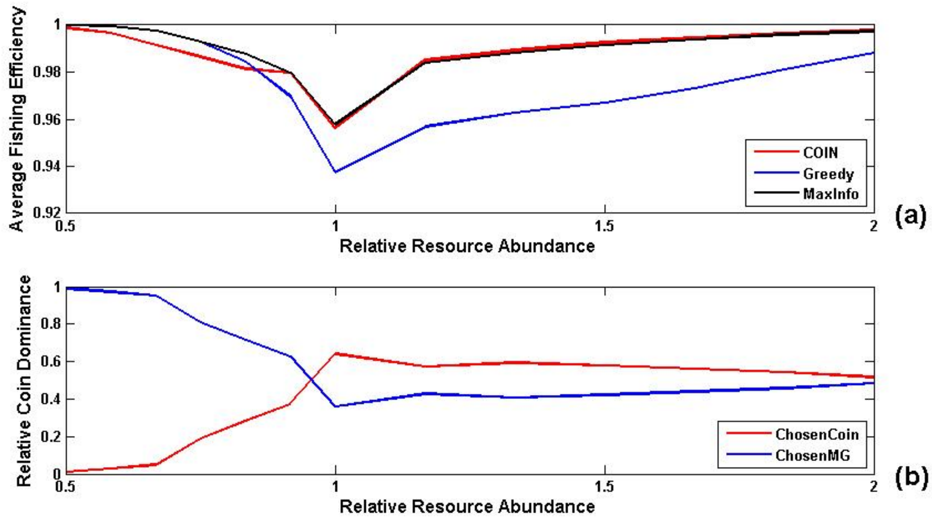

Figure 2. (a) Average fishing efficiency as a function of relative resource abundance for the MG (blue), Coin (red) and *MaxInfo* (black). *MaxInfo* matches MG for scarce resource and Coin for abundant resource and outperforms both for intermediate resource status. (b) Ratio of individual vessels for which Coin provides a more informative option than MG within the *MaxInfo* approach. For very scarce resource almost all agents choose the MG; for increasing resource more and more agents adopts Coin. Coin choice dominates exactly at the transition between over-exploited and under-exploited resources.

In Figure 2a *MaxInfo* performance matches MG for low resource and Coin for ample resource levels, thus following the optimal strategy at both extremes of the plot. This suggests that *MaxInfo* adapts itself to the resource status. Importantly, *MaxInfo* behaves better than both Coin and MG close to the transition from under to over-exploited resource. This is the resource distribution for which the average fishing efficiency is at its minimum [8] since in this configuration the catches are most sensitive to the fleet effort distribution. This also suggests that an adaptive strategy is particularly beneficial in these difficult scenarios.

In Figure 2b, for the *MaxInfo* strategy, we show the average ratio of individual vessels for which ${}^{C}W^{t}_{n1,z}$ at point 3 above provided the most unequivocal forecast (red line) and the ratio of individual vessels which used ${}^{MG}W^{t}_{n,z}$ (blue line). Basically, Figure 2b shows how many individual vessels chose Coin versus MG for their dynamical decision making, as a function of resource abundance. As expected, for very scarce resource almost all agents choose the MG; for increasing resource more and more agents adopts Coin. Two further features can be noticed. First, the transition from Coin to MG dominance occurs for a resource abundance slightly larger than the one for which Coin outperforms MG in Figure 2a. Second, the ratio of vessels choosing Coin reaches a maximum in the vicinity of the transition from scarce to abundant recourse.

Once again, this result confirms the conjecture we already suggested from Figure 1: The balance between vessels dynamically adopting Coin versus MG gives an indication of resource abundance. However, unlike the previous results, this approach suggests a more direct way in which this information could be used in order to reduce the fishing pressure on the resource and we discuss this in the next section.

7 Centralised versus decentralised strategies.

In the previous section we modelled a fleet of 50 independent vessels. As in a typical Tragedy of the Commons, should a vessel realise via the curve in Figure 2b that the resource is over-exploited, it would have little incentive not to go fishing, because this would result in lost income. However, if the vessel is part of a larger fleet, there could be an incentive for the fleet manager to prevent the vessel from going fishing, since his fleet would obtain the same overall catch with reduced costs.

In the following we thus apply the algorithm described above to a setting in which the fleet is subdivided into smaller sub-fleets, each directed by its own manager. While the scenario in the previous section describes a totally decentralised fleet, in which each vessel takes its own independent decision, here we model various levels of centralisation, depending on the size of the sub-fleets. In the limit in which the entire fleet is directed by a single manager we have a fully centralised decision process.

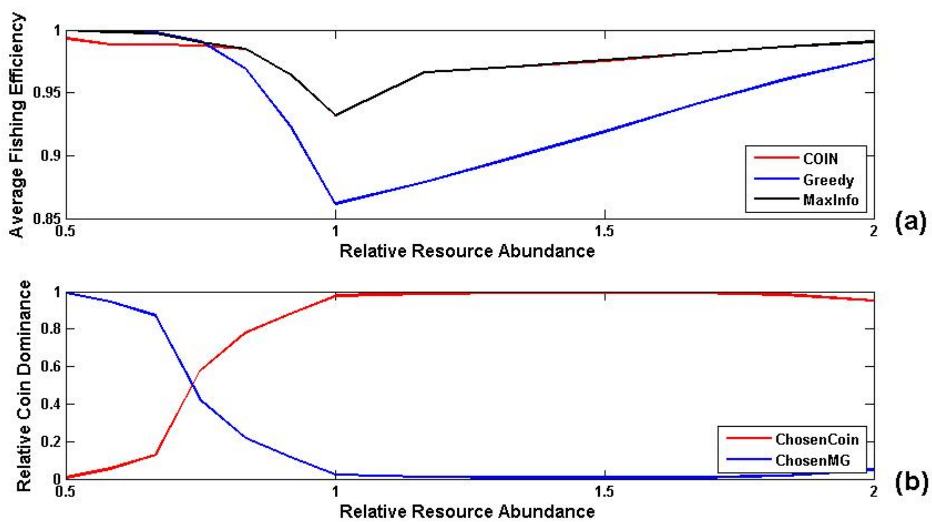

Figure 3. 50 vessels divided into 10 sub-fleets of 5 vessels each. Decision making is carried by the 10 sub-fleet managers, not by the 50 individual vessels skippers. (a) Average fishing efficiency as a function of relative resource abundance for the MG (blue), Coin (red) and *MaxInfo* (black). (b) Ratio of individual vessels for which Coin provides a more informative option than MG within the *MaxInfo* approach.

In Figure 3, we show the results for $N=50$ vessels divided into 5 sub-fleets of 10 vessels each. Two main differences can be noted between Figure 2 and 3. First, the performances of all strategies worsen for this mildly decentralised scenario. However, the MG worsens considerably more than Coin or *MixInfo*. In Figure 3a, as in 2a, *MaxInfo* combines the benefits of the MG for scarce resources and of Coin for abundant resources, giving the best performance in both resource states. Second, while the transition from MG to Coin dominance in Figure 3a occurs for a resource balance very close to the one in Figure 2a, the transition from MG to Coin dominance in Figure 3b occurs for scarcer resources and thus there is a closer correspondence between the transition between Figures 3a and b than between Figures 2a and 2b. Finally, while the maximum of Coin use in *MaxInfo* is still reached for a relative resource abundance very close to 1, now the difference between the number of vessels adopting Coin versus MG in *MaxInfo* is much stronger.

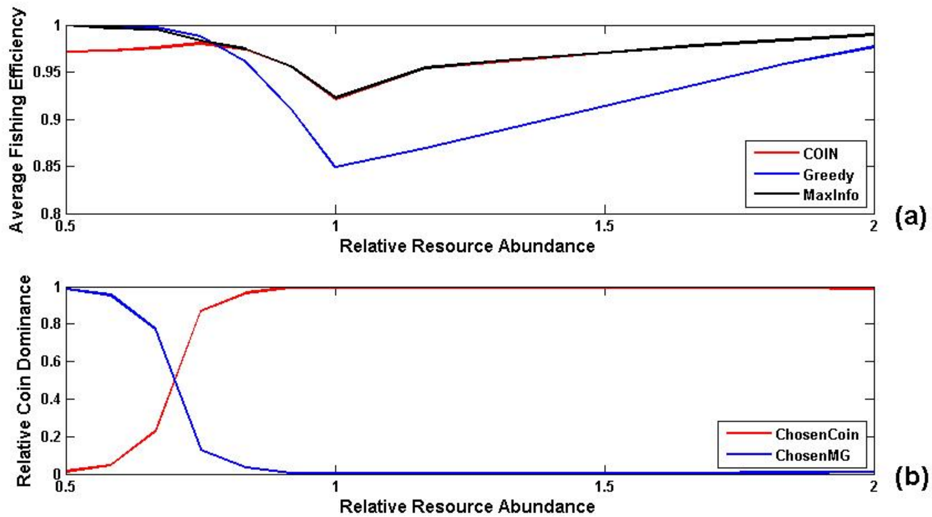

Figure 4. 50 vessels behaving as a single fleets under centralised decision making. (a) Average fishing efficiency as a function of relative resource abundance for the MG (blue), Coin (red) and *MaxInfo* (black). (b) Ratio of individual vessels for which Coin provides a more informative option than MG within the *MaxInfo* approach.

Finally in Figure 4 we show the results for a fully centralised fleet, in which decision making for all vessels is carried our by a single manager. The main difference, compared to the previous figures, is that the dominance of Coin choices in Figure 4b occurs for even scarcer resources, suggesting that increasing centralisation favours the adoption of Coin behaviour.

8 Discussion

The core of the approach we presented in this paper lies in choosing dynamically what strategy to employ to carry out a complex task in a competitive environment. Unlike previous work in [7], the dynamical choice is not carried out via standard evolutionary criteria, that is by evaluating a strategy according to past performances, rather via an information theoretical criterion, which is used to evaluate and compare the information contents of possible strategies. The first results seem promising, but a number of issues need to be addressed before firm conclusions on the potential of this approach can be made. We discuss some of these in this section.

In the view of a potential adoption by human agents, the information based approach seems to have some advantages against the one based on performance. The agents need to carry out more book-keeping, since records of both catches and impacts need to be stored and calculated, however the processing of this records requires few simple operations for which in principle no computation devise is strictly necessary. More important, the agents do not need to commit in advance to a specific strategy in order to evaluate its returns, nor do they need to trust reports from colleague/competitors in order to judge a strategy they have not employed; all they need to do is to check which of their own record (Coin or MG) provides a more informative instruction. This may make it easier to convince agents in the real world to adopt this approach.

However, because of the nature of the problem we address, any decision taken by an agent will affect the future behaviour of the community and thus of future resource status as well the future behaviour of the agent itself. For example, previous results shows that, in the case of slightly over-abundant resource, if most agents follow a Coin strategy, and consequently the fleet spreads its effort according to the resource distribution, the catch improves. It thus follows that, by following Coin, the difference between impact and contribution is reduced. We encountered this very result in the discussion of Figure 2b, in which the difference between information content of Coin and MG for a fully decentralised approach was unexpectedly low. It is reasonable to expect that such an information gap would have been higher, had less agents chosen a Coin strategy. We should thus expect that the behaviour of the *MaxInfo* approach will be affected by the history of the simulation and more experiments to evaluate it are needed.

The second thread of the paper is the potential use of the *MaxInfo* approach to monitor the resource status. This also shows promising results, but leads to the question of what to do with such information should it suggest that the resource is badly over-exploited. Decision making in this scenario requires a proper model of exploitation cost, as well as alternative employment options for fishers and managers should a reduced fishing effort be necessary. We plan to analyse this in future work by casting it into an evolution dynamics scenario as in [13].

In this paper we modelled a fully renewable resource, and perfectly rational agents, which not only do not make mistakes, but also never cheat. Both assumptions are unrealistic, and have been chosen because of a ‘reductionist’ intent of separating the factors affecting a problem and evaluating the *MaxInfo* approach in an easy to understand test case. Clearly the approach will need to be assessed under more realistic scenarios as in [6,7], which we also endeavour to do in the near future.

Finally, it is crucial to assess the receptivity of the approach for real agents. It is equally important to evaluate what modifications real agents may impose on the method, should they adopt it. We initially tested this in a role playing experiment with human actors and witnessed the unexpected solutions human subjects may provide to given problems. These are very hard to model and forecast in detail. More extensive experiments in this setting have already been planned and we believe are crucial to properly evaluate this approach and the best adoption path.

9 Conclusions

We presented a method which allows agents to choose dynamically between purely competitive and collaborative strategies in a mixed-strategy approach. In a minority game-like problem, this approach shows improved performance for both centralised and decentralised decision making. The balance between agents choosing the competitive and collaborative strategy also gives an indication of the resource state and may be used for both monitoring and management.

References

[1] Blackmore, S., (2001), Evolution and Memes: The human brain as a selective imitation device, Cybernetics and Systems, Vol 32:1, 225-255.
[2] Lee, S., (2004), Scientific alternatives to the anthropic principle. eprint arXiv:hep-th/0407213
[3] Arthur, W.B., (1994), Inductive behaviour and bounded rationality, The American Economic Review, 84, 406-411.
[4] Boschetti, F., (2007), Improving resource exploitation via Collective Intelligence by assessing agents' impact on the community outcome, Ecological Economics, 63, 553-562.
[5] Batten, D., (2007), Are some human ecosystems self-defeating?, Environmental Modelling and Software, 22(5), 649-655
[6] Brede, M., F. Boschetti & D. McDonald, (2007), Strategies for Resource Exploitation, Ecological Complexity, accepted.
[7] Brede, M., & De Vries, (2007), Harvesting Heterogeneous Renewable Resources: Uncoordinated, Selfish, Team-, and Community-Oriented Strategies, Environmental Modelling & Software, in review.
[8] Brede, M., Boschetti, F., & McDonald, D., (2007), Managing renewable resources via Collective Intelligence, ModSim07, ChristChurch, NZ, December 2007.
[9] Challet, D. and Y-C, Zhang, (1998), On the minority game: analytical and numerical studies. Physica A, 256, 514.
[10] Gintis, H., (2000), Game Theory Evolving: A Problem-Centered Introduction to Modeling Strategic Interaction, Princeton University Press, Princeton, NY
[11] Hardin, G., (1968), The tragedy of the commons, Science, 162, 1243-8.
[12] Wolpert, D. and K. Tumer, (2001), Optimal Payoff Functions for Members of Collectives, Advances in Complex Systems, 4(2/3):265–279.
[13] Brede, M., and Boschetti, F., 2007, Commons and Anticommons in a simple Renewable Resource Harvest Model, submitted.
[14] Wolpert, D., K. Wheeler, and K. Tumer, (2000), Collective Intelligence for Control of Distributed Dynamical Systems. Europhys. Lett., 49 (6), p. 708.
[15] Wolpert, D. and K. Tumer, (1999), An Introduction To Collective Intelligence, Tech Report: NASA-ARC-IC-99-63, http://ic.arc.nasa.gov/ic/people/kagan/coin pubs.html.
[16] Wolpert, D., K. Tumer, and D. Bandari, (2004), Improving Search Algorithms by Using Intelligent Coordinates, Physical Review E, 69, 2004.
[17] Bieniawski, S., D. Wolpert, & I. Kroo, (2004), Discrete, Continuous, and Constrained Optimization Using Collectives, AIAA Paper 2004-4580, Presented at the 10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Albany, NY, August 30-September 1, 2004.

[18] Shannon, C.E. (1948), "A Mathematical Theory of Communication", Bell System Technical Journal, 27, pp. 379–423 & 623–656, July & October, 1948.