Selection pressures for a theory-of-mind faculty in artificial agents

Jason Noble, Tom Hebbron,
Johannes van der Horst, Rob Mills, Simon T. Powers & Richard A. Watson

Science & Engineering of Natural Systems group
University of Southampton

Theory of mind?

Anticipating the behaviour of others based on what you think they know.

Need to represent the environment, and the idea that others also have such representations.

Adult humans can do it. Children and primates are borderline cases. Many animals can't do it.

Implicated in some big developments, i.e., language, culture, teamwork, self-awareness.

So: how does it evolve? *When would you want it?*

Previous work

Earlier ALife work has looked at ToM in collision-avoidance situations (Takano & Arita, ALife X, and Zanlungo's PhD thesis). Think of that little dance that two pedestrians do when trying to guess the other's intentions.

Interesting work, but it's strictly an anti-coordination game, and we wanted to look at more general conditions for ToM to evolve.

Kinds of minds

A classification scheme borrowed from Dennett.

- Zero-order agent: purely reactive to perceptual inputs.
- 1st-order agent: builds on this by including internal state that maps the environment, e.g., remembering where a predator was last seen.
- 2nd-order agent: has basic ToM via a world-model that includes the internal states of other agents (e.g., "there's a predator behind that tree, but my friend hasn't seen it yet").
- 3rd- and higher-order agents include a recursive aspect, i. e., a model of what I think he thinks I am thinking.

Pragmatic approach

ALife's individual-based simulations allow us to ask: when would this capacity be selected for? In what ecological niches would it be favoured?

Complex cognition doesn't appear for no reason: extra brain functions are going to incur costs, and so there will need to be a benefit if a novel mutation is going to fix in the population.

We ask: in what environments will it be selectively advantageous to climb Dennett's ToM ladder?

The answer should be useful in figuring out what's special about ToM-equipped animals like ourselves.

Simulation environment

Ten agents occupy a ring of 10 abstract locations.

They can perceive their current location and its two neighbours.

Locations payoffs depend on the number of agents present, and are either cooperative (e.g., hunting) or competitive (e.g., splitting a limited food supply).

Thus, high payoffs are achieved by either coordinating or anti-coordinating behaviour.

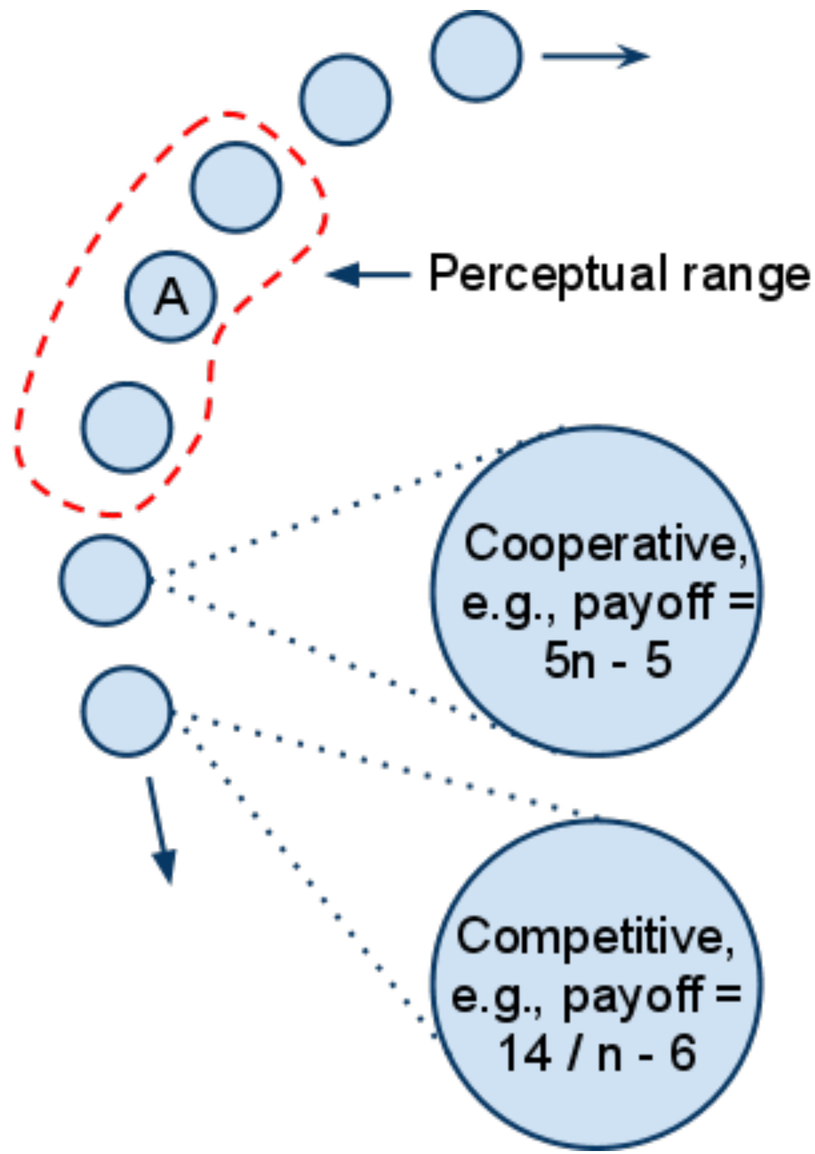

Invasion studies

No genetic algorithm: we looked instead at a population of agents and asked whether a mutant that was one level up the scale could successfully invade.

We also looked at back-sliding: could a mutant from one level *down* invade a population of higher-order agents?

We explored multiple parameters (e.g., percentage of the world visible to one agent) but the major one was the degree of cooperation vs. competition in the environmental payoffs.

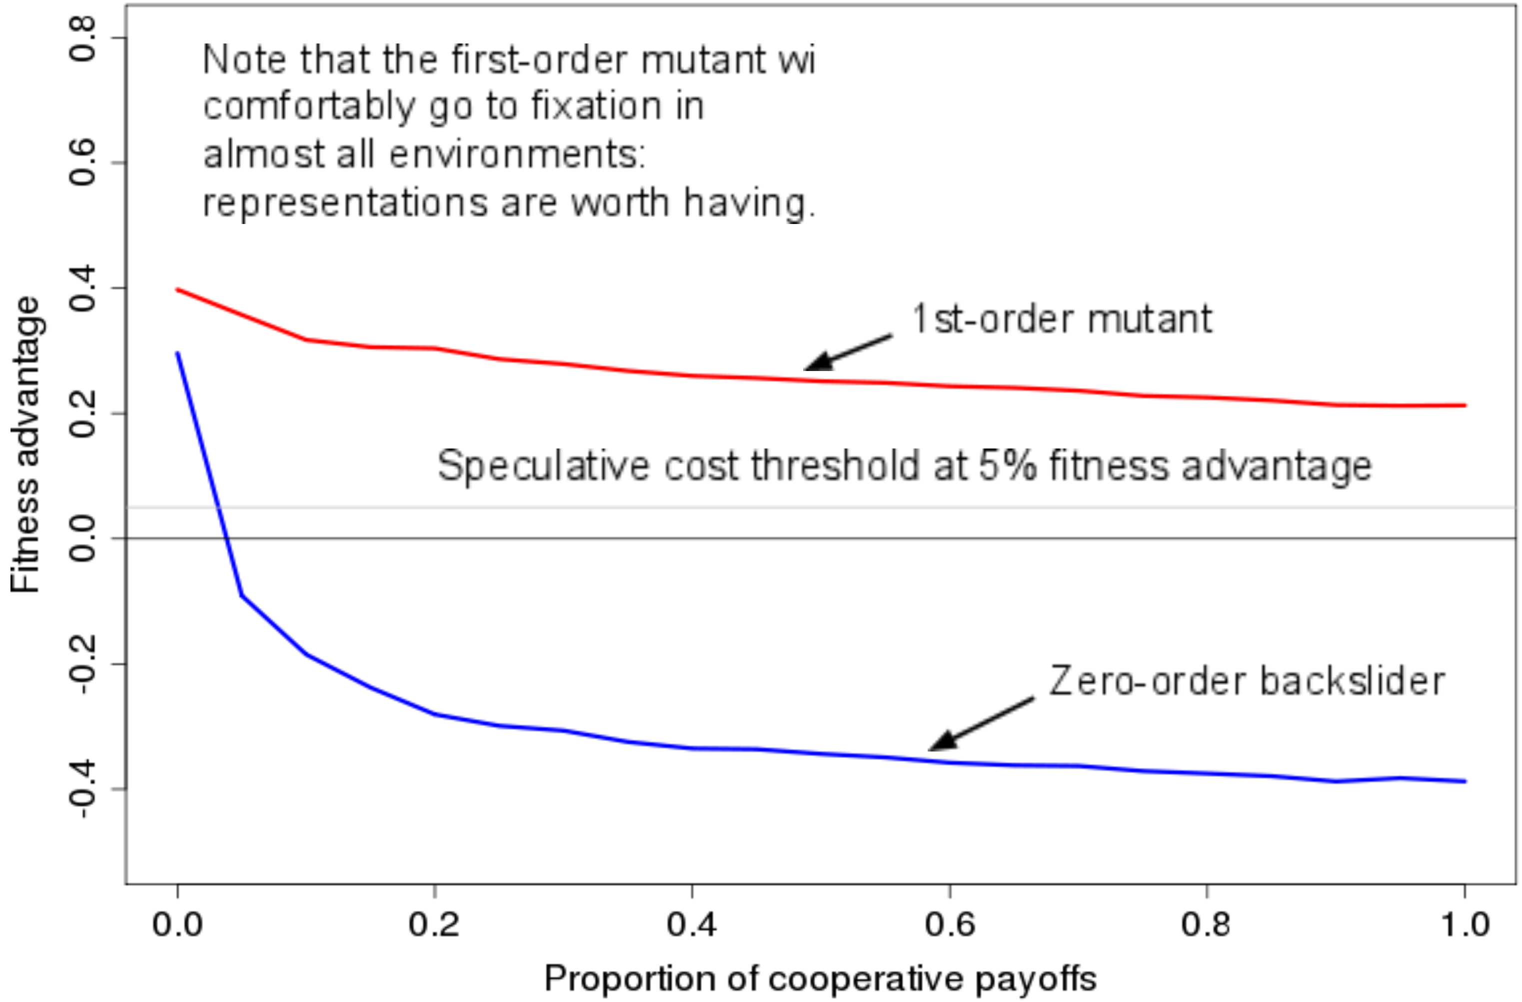

Do 1st-order agents invade a zero-order population?

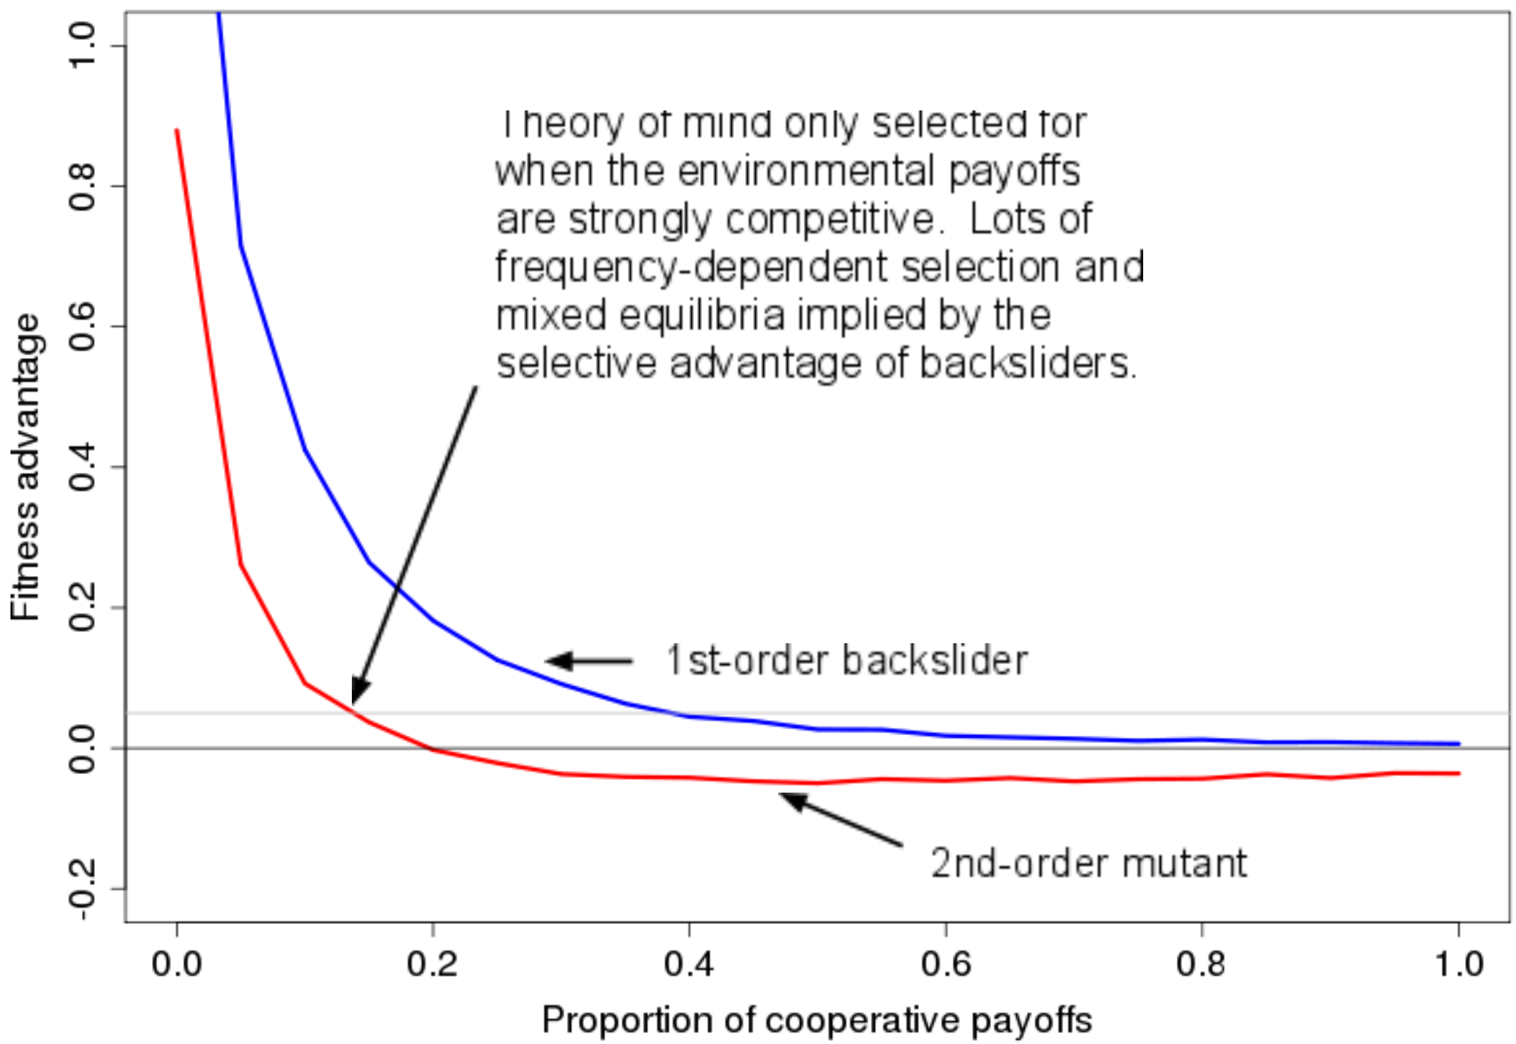

Do 2nd-order agents invade a 1st-order population?

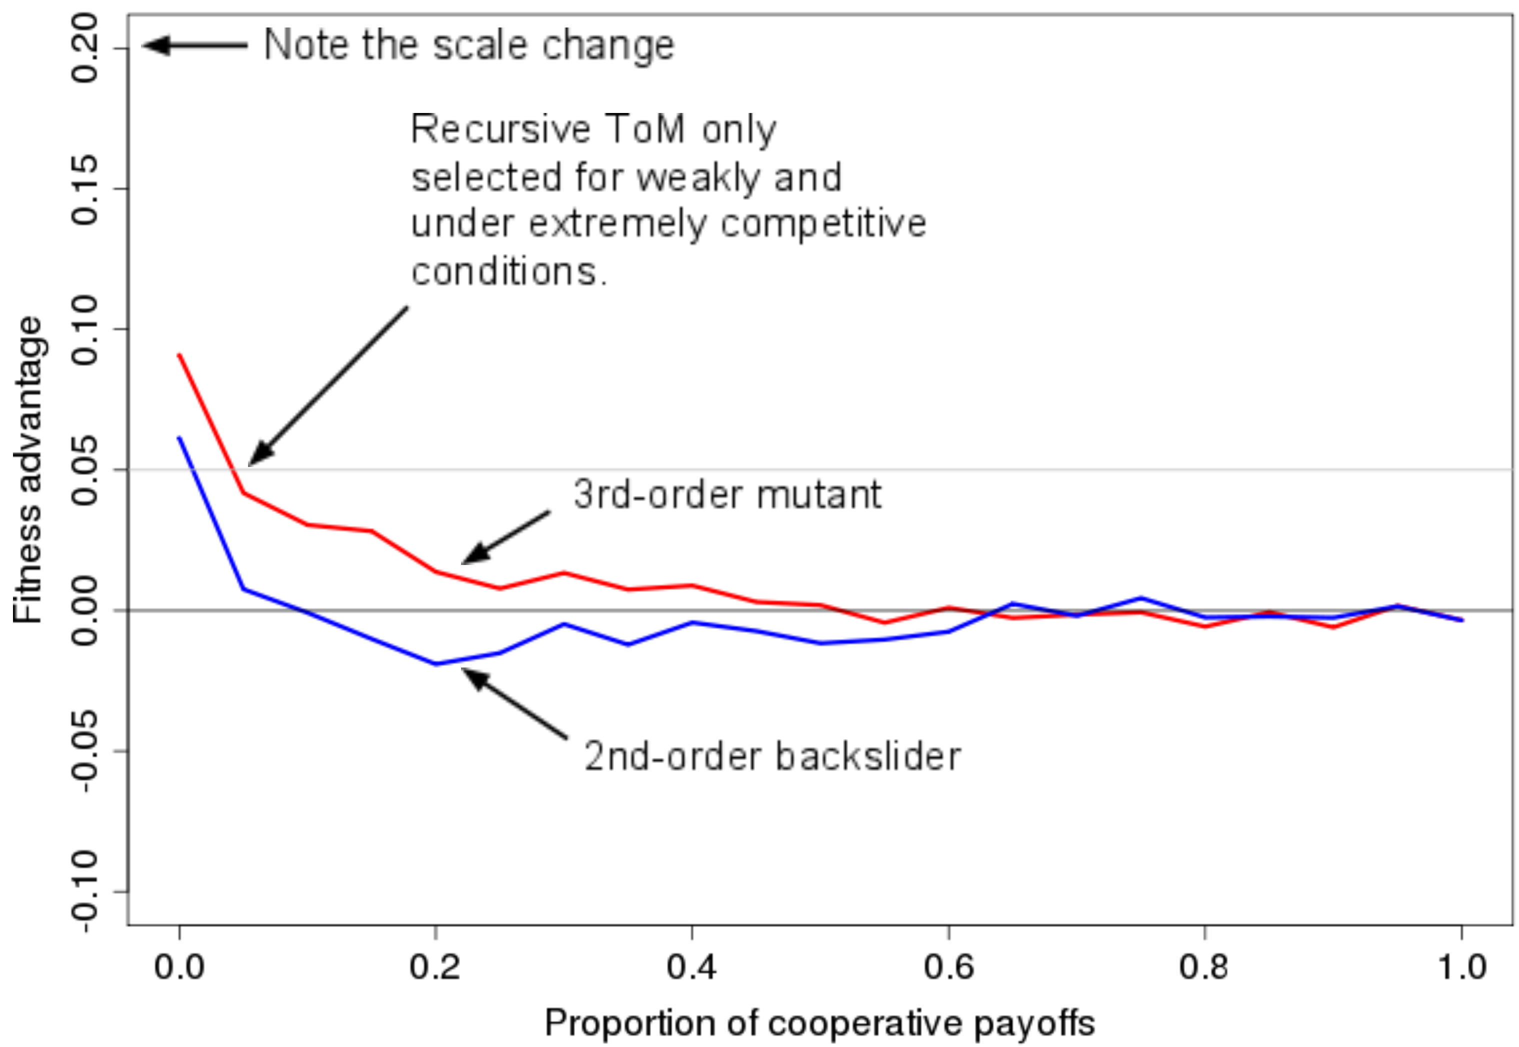

Do 3rd-order agents invade a 2nd-order population?

Conclusions

1st-order (representational) agents are broadly selected for.

2nd-order agents (ToM) only happen if the environment is very competitive, and will be in a mixed-strategy equilibrium with 1st-order agents.

Machiavellian 3rd-order agents are only weakly selected for even in highly competitive contexts.

Just-so stories about ToM having its origins in cooperative hunting look weak. If you have lots of cooperative options already, there is no pressure to model the mental states of others as "selfish" behaviour by all will lead to good outcomes.

What if the world gets bigger?

In other words, what if there are more agents, all looking at distinct parts of a bigger environment?

To get 1st-order systems, the world needs to be bigger than what you can see in one go, otherwise direct perception would do just as well.

Being a 2nd-order system is selected for over a variety of intermediate population sizes / levels of perceptual overlap.

Being a 3rd-order system is not selected for very strongly once the population becomes large.

Some parameter values

NUM_AGENTS 10
NUM_LOCATIONS 10
NUM_REPEATS 200
NUM_TIME_STEPS 2000

VIEWABLE_LOCS 3

PROB_MOVEMENT 1.0
PROB_RANDOM_JUMP 0.05
PROB_CHANGE_PAYOFF 0.02