

Modelling and Influencing the AI Bidding War: A Research Agenda

The Anh Han
Teesside University
Middlesbrough, UK
t.han@tees.ac.uk

Luís Moniz Pereira
Universidade Nova de Lisboa
Lisbon, Portugal
lmp@fct.unl.pt

Tom Lenaerts
Université Libre de Bruxelles and
Vrije Universiteit Brussel
Brussels, Belgium
tlenaert@ulb.ac.be

ABSTRACT

A race for technological supremacy in AI could lead to serious negative consequences, especially whenever ethical and safety procedures are underestimated or even ignored, leading potentially to the rejection of AI in general. For all to enjoy the benefits provided by safe, ethical and trustworthy AI systems, it is crucial to incentivise participants with appropriate strategies that ensure mutually beneficial normative behaviour and safety-compliance from all parties involved. Little attention has been given to understanding the dynamics and emergent behaviours arising from this AI bidding war, and moreover, how to influence it to achieve certain desirable outcomes (e.g. AI for public good and participant compliance). To bridge this gap, this paper proposes a research agenda to develop theoretical models that capture key factors of the AI race, revealing which strategic behaviours may emerge and hypothetical scenarios therein. Strategies from incentive and agreement modelling are directly applicable to systematically analyse how different types of incentives (namely, positive vs. negative, peer vs. institutional, and their combinations) influence safety-compliant behaviours over time, and how such behaviours should be configured to ensure desired global outcomes, studying at the same time how these mechanisms influence AI development. This agenda will provide actionable policies, showing how they need to be employed and deployed in order to achieve compliance and thereby avoid disasters as well as losing confidence and trust in AI in general.

CCS CONCEPTS

• **Computing methodologies** → **Model development and analysis**; *Cognitive science*; • **Mathematics of computing** → *Stochastic processes*.

KEYWORDS

artificial intelligence race, artificial general intelligence, safety AI, incentives, evolutionary game theory

ACM Reference Format:

The Anh Han, Luís Moniz Pereira, and Tom Lenaerts. 2019. Modelling and Influencing the AI Bidding War: A Research Agenda. In *2019 AAAI/ACM Conference on AI, Ethics, and Society (AIIES'19)*, January 27–28, 2019, Honolulu, HI, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3306618.3314265>

1 INTRODUCTION

Research and development in different areas of fundamental and applied Artificial Intelligence (AI) have been making encouraging progress. Within the research community, there is a growing effort to make progress towards Artificial General Intelligence (AGI). AI brings enormous potential benefits across many sectors, being recognised as a strategic priority by a range of actors and stakeholders, including representatives of various businesses, private research groups, companies, and governments [2]. The media attention as well as the (un)announced business and political ambitions indicate that an AI bidding war has been triggered, competing to be the first to develop and deploy a powerful, transformative AI [3, 6, 7, 10, 11, 37]. These AI systems could be either AGI, able to perform a broad set of intellectual tasks while continually improving itself, sufficiently powerful specialised AIs or even AI specifically developed for espionage and cyberterrorism.

An AI race for technological advantage towards powerful AI systems could lead to serious negative consequences, especially when ethical and safety procedures are underestimated or even ignored [3, 11]. Safety and ethical agreements and regulations can be adopted to ensure that all parties involved in the race will comply with a set of mutually agreed standards and norms [31, 31]. However, as experience with many international treaties, such as climate change, fisheries and timber agreements [5, 13, 27] has shown, the autonomy and sovereignty of the parties involved will make monitoring and compliance enforcement difficult (if not impossible). Therefore, for all to enjoy the benefits provided by a safe, ethical and trustworthy AI, it is crucial to enact appropriate incentivising strategies in order to ensure mutual benefits and safety-compliance from all sides involved.

This position paper sets out a research agenda (i) to develop theoretical models (both analytic and simulated) that capture key factors of an AI race, in order to provide understanding on the dynamics and emergent strategic behaviours for different hypothetical race scenarios; and furthermore, (ii) to examine how incentives can be used to ensure desired outcomes and equilibria in this race. To this end, we combine research on incentives and agreement modelling, with the dynamical approaches that analyse population-wide dynamics, typically employed in Evolutionary Game Theory (EGT)

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

AIIES '19, January 27–28, 2019, Honolulu, HI, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6324-2/19/01...\$15.00

<https://doi.org/10.1145/3306618.3314265>

[23, 32]. Together, they will permit us to systematically explore how different types of incentives (namely, positive vs. negative, peer vs. institutional, and their combinations) can influence safety-compliance behaviours over time, and how such behaviours should be configured to ensure desired global outcomes (i.e. high levels of safety-compliance, possibly including compliant information sharing).

This research agenda is expected to lead to specific outcomes concerning the AI bidding war we are currently observing. On the one hand, it will provide methodologies for the investigation of AI race dynamics, including intervention mechanisms to influence outcomes. On the other, it will provide systematic understanding on how to promote safety compliance and the production of AI systems that benefit the many as opposed to the few. Moreover, conclusions derived from this project have the potential to resolve issues in other multiparty interaction domains, such as business transactions and international environmental agreements, for which commitments have been proven crucial [27, 35].

The rest of the paper is structured as follows. The next section describes relevant factors influencing the AI race and existing approaches for agreements and incentives modelling. It then describes our research agenda, which includes three directions aiming to provide a methodology for AI race modelling. Some preliminary results will be briefly described, elaborated further in the Supplementary Information section.

2 BACKGROUND AND CHALLENGES

Below we review relevant issues and potential factors in AI race modelling, and EGT [32] research on incentives, commitment and agreement modelling.

2.1 AI race and modelling

Potential AI disaster scenarios are many [3, 28, 37]. However the uncertainties in accurately predicting these outcomes are high. In general, the first AGI are likely to be extremely powerful and, if not developed and controlled properly, might not be guaranteed human friendly [3, 8, 10]. Those that first successfully develop more powerful AI will have significant benefits over others, and even ‘the winner takes all’ scenario might come about [2]. It is generally agreed that an AI-related disaster is more likely to occur when safety measures are ignored more often, which might be fostered by the speed and competition of the race. The risk of AI-related disaster increases when teams/developers do not devote sufficient attention and resources to safety in such a powerful system [2] as a result of the on-going race’s pressure.

Regardless of this disaster risk, people and teams can have different (biased) perceptions or beliefs about the level and nature of the risk as well as adopt different risk-taking behaviours (e.g. with different levels of risk-aversion). A high belief in risk might lead to fear mongering concerning both AI and AGI, thereby leading to overregulation and unnecessary obstruction to the development of AI and the realisation of social/economic benefits. On the other hand, a low belief in risk might lead to more risk-taking behaviours, and hence safety measures being omitted more often. Furthermore, envisaging the AI race at diverse temporal scales is likely to bring about distinct aspects that should be focused on [1]. For instance, each team might

anticipate different speeds of reaching the first general AI system. A low belief about fast AGI arrival could result in miscalculating the risks of unsafe AGI deployment by non-compliant rogue teams.

Other important strategic aspects in the AI development race concern different types of openness, including openness about source code, science, data, safety techniques, capabilities, and goals [7]. Some forms of openness, e.g. about safety measures, are plausibly positive both short-term and long-term. Others, such as openness about source code, science, and possible capabilities, could intensify the competitive situation. A central concern is that openness could exacerbate a racing dynamic: competitors trying to be the first to develop true AIs may accept higher levels of (existential) risk in order to accelerate their progress. On the other hand, openness may reduce the probability of AI benefits being monopolized by a small group and facilitate counter measures.

Given the several factors and biases driving the strategic safety-compliant behaviours and the dynamics of interactions among development teams, the outcome of the race (e.g. whether or not disaster occurs) is difficult to predict. Very few efforts have been made in modelling AI disaster and race dynamics among teams. In [3], the authors provide a game-theoretical model of an AI race, under the assumption that the first AI will be very powerful and transformative, so that each team is encouraged to finish first and thus skimp on safety precautions. This work, however, does not consider dynamical aspects of the race (e.g. how teams might adapt their strategies over time), and which strategic behaviours emerge in different race scenarios. Also, this work does not study how positive or negative incentives can be used to enhance teams’s safety compliance.

To the contrary, there has been a significant body of computational modelling research regarding game-theoretical and EGT analyses of other disasters, such as climate change and nuclear war [4, 30]. However, the AI race and its related risks are quite unique, according to an analysis of 12 large global catastrophic risks in [28]. Climate change disaster analyses primarily focus on the unwillingness of participants to take on themselves a personal cost for a jointly desired collective target, and is conjoined to collective risk by all parties involved [30]. In contrast, in AI racing, the winner(s) will derive significant relative advantage over others, risk being more individualized. The AI race is also different from the nuclear arms race, in that the former potentially poses greater effective achievement or otherwise risks to its creator, whereas nuclear powers are generally not at direct risk originating from their own arsenals [3].

Finally, despite a number of proposals and debates on how to prevent, regulate, or solve the AI race, there is significant lack of rigorous modelling studies [6, 11, 14, 31, 36]. Our proposed research agenda aims to bridge this gap by providing basic models to systematically understand the dynamics of the race and what strategic behaviours would likely emerge, in different hypothetical scenarios and conditions of the race concerning disaster risk, risk-perception behaviours, level of information openness, number of racing teams, incentives deployment, etc. These models can then be used and perfected in discussions with stakeholders as well as policy makers.

2.2 Reward and punishment

Punishment and reward are major forms of incentives, widely adopted for enforcing cooperative behaviours among self-interested agents and enacting norm compliance in both social interactions and computerised systems [12, 21, 33]. Various forms of such incentives have been studied, which can be roughly categorised into peer and institutional ones.

To provide a peer incentive, agents pay a personal cost to punish a violator (peer punishment) or reward a cooperator (peer reward), after an interaction. As a result, the punished violator and rewarded cooperator incur a decrease or increase, respectively, in their payoff. In the context of AI race behaviours, examples of peer incentives use are when teams refuse to support, and so do not share development progress and knowledge with non-compliant teams, consequently slowing down their development. They can also arrange, for instance, cyber-attacks against non-compliant teams, or even ascribe and spread their bad reputation, leading to those teams being unable to recruit and retain developers. On the other hand, highly compliant teams can be rewarded, obtaining more support, such as knowledge and experience sharing, and thus move faster along in the development race.

In contrast to peer incentives, institutional incentives assume instead the existence of an institution (with a budget) to take care of the incentivising process [12]. Pool and coordinated group incentives, where agents can put together a fund before an interaction occurs to provide incentives subsequent to the interaction, can be considered as a first step towards institutionalisation of incentives [9, 34]. Examples of institutions providing incentives are modern courts and policing systems, as well as international organisations such as the United Nations (UN), the European Union (EU), the World Trade Organization (WTO) or the Organisation for Economic Co-operation and Development (OECD). Setting-up and maintaining institutions are costly, but the presence of a powerful authority can effectively restrict individuals' strategic options and provide stronger incentives [22]. In the context of an AI race, a centralised access to AI-related knowledge, algorithms and tools, as in the EU platform call (H2020-ICT-2018-2020) for instance, might provide institutional incentives such as strong support and punishment (e.g. allowing high levels of access or exclusion from the centralised pool of AI knowledge).

For both peer and institutional incentives, the critical conditions for cooperation to be achieved and sustained in evolutionary models, as well as observed in lab experiments, require that incentives 'fit the crime' [9, 34]. That is, they increase with the severity of the violation or the merit of the cooperation. Moreover, they need to be cost-effective, that is, the effect of an incentive for its receiver should be sufficiently large compared to the cost to the provider. Interestingly, for institutional incentives, the centralisation by which an institution might observe a global population state enables more efficient approaches. For instance, the 'first carrot, then stick' policy that switches the incentive from reward to punishment whenever the frequency of cooperators in a population exceeds a threshold, is highly efficient at establishing full cooperation, better than either reward or punishment alone [21]. Furthermore, local institutions that provide incentives based on local neighbourhood properties (such as its cooperativeness level) have proved more efficient than

large global governance overseeing the whole population, in the context of climate change games and cooperation dilemmas [16, 38]. It may be that recourse to costly institutional supervision is only enacted when spontaneous peer cooperation becomes insufficient or inadequate.

2.3 Commitment modelling in dynamical systems

Commitment-based formalisms have been widely adopted in Multi Agent Systems (MAS) as modelling and engineering tools [35], with important applications in business protocols, transactions, and software-oriented architectures. Representing commitments that agents have to one another specifies their expected interaction and correct behaviour (e.g. cooperation and norm compliance). The formalisms allow flexible capturing of contractual relationships amongst the agents concerned and, as such, incentivise their correct desirable behaviour. In a commitment-based MAS, it is crucial to understand how an adopted commitment formalism, including its compliance incentivising approach, influences agents' behaviours and the dynamics of their interactions and, as a result, determines whether it can actually help foster and secure substantial commitment compliance and correct behaviour. This is typically studied in the literature on dynamical MAS, using methods from EGT. In this dynamical context, agents interact and adapt their behaviours over time, e.g. via social learning [17, 32, 34].

In fact, commitments have been shown to provide important pathways to reaching mutual cooperation in the context of social dilemma interactions (namely, in the Prisoners Dilemma (PD) [18] and Public Goods Game (PGG) [17] settings), within populations of self-interested agents. This dynamical approach has provided important insights into the design of commitment-based MAS, enabling identification of areas of strengths and weaknesses of commitment-based mechanisms. Consequentially, it is now possible to make specific efforts to improve on concrete issues. This approach also provides novel understanding towards the long quest for the evolution of commitments and their roles in the evolution of cooperation [27]; with important practical ramifications for social and economic interactions, ranging from personal relationships to business contracts and to international agreements [5, 13].

However, these modelling studies focus on standard cooperation dilemmas settings, i.e. the PD and PGG (for both one-shot and repeated interactions). When moving to AI race interactions, several factors need to be taken into account, as described above, such as the risk of disaster, risk perception and levels of openness. And the incentives need to be accordingly designed to account for these new factors. For instance, by requiring that signatories to an AI race treaty or agreement will commit to allow internationally governed inspections concerning a defined minimal transparency or the detection of secret of undertakings. Our study may incorporate specific distinct punishments or rewards for the case of violation or otherwise, or lack of due diligence, in both such cases. Public and employee opinion and consumer practice can also be mustered as another form of incentive.

3 RESEARCH AGENDA: PROBLEMS TO BE STUDIED

The literature reviewed above clearly shows the importance and gaps in modelling research to understand the dynamics of interactions and emergent behaviours in an AI race, as well as how incentive mechanisms might be used effectively to promote desired behaviours, such as safety standards and norms compliance from all parties involved in the race.

Our research agenda below aims to bridge this gap, and the following three directions or parts are envisaged to achieve that goal. The first one will develop baseline models for an AI race, on the basis of which the incentives for influencing the race will be studied in subsequent projects.

3.1 Development of baseline AI race models

The first direction aims to develop new, baseline strategic decision-making models that capture key factors relevant to the competition and cooperation among AI teams, and that investigate how they influence the outcome of the AI race. Namely, we will develop EGT models to answer the following questions:

- Which strategic behaviours emerge in the AI race?
- How does increasing the number of competing teams influence the evolutionary dynamics and outcome of the race, namely in preventing disaster?
- How does the risk-perception probability of AI disaster influence the race outcome?
- How does the heterogeneity of teams' development capacities influence the race outcome and strategic behaviours?

We envisage multi-player game models of multiple stages and/or repeated interactions, where a team might react to the development of strategic behaviours of other teams. The first teams to successfully develop AI will have significant benefits over others, and even the 'winner takes all' scenario might come about. Although we focus on multiplayer games, for there are most likely multiple competing AI teams, as the race evolves the game might end up at a later stage with only a few or even two (strongest) teams. Our analysis will start with pair-wise interactions, modelled as two-player games, which will afterwards be generalised to multi-team interactions, modelled as multiplayer games, in order to provide a more comprehensive understanding.

We start with the most basic scenario where at each round of the race a team or player is faced with two possible choices: to follow the safety precaution (SAFE) or to ignore this safety precaution (UNSAFE). Since it takes more time and effort to comply with the precautionary requirements, playing SAFE is not only costlier, but also implies a slower development speed, compared to playing UNSAFE. As a generalisation of this binary-choice model we will consider continuous games where a player can choose the level of safety-precaution to adopt (i.e. SAFE and UNSAFE correspond to the two extreme cases of complete precaution and none at all, respectively).

Teams can collect benefits from their intermediate AI products. However, differently from the standard repeated games [32] where all players collect benefit at every round, we will need to consider a new time scale, where different teams might collect benefits at

different speeds. That would mean a possible time delay in players' decision-making, during the course of a repeated interaction, because they might want to wait for the outcome of a co-player's decision to see what choice he/she has adopted and/or will adopt in the next development round. Thus, a player has to decide whether to make an immediate move based on just present information (and hence be quicker to collect the next benefit and move faster in the race) but at the risk of making a worse choice, different from one that would have been chosen had the player already known the co-player's decision. Moreover, since noise is a key factor driving the emergent strategic behaviours in the context of repeated games [32]—for instance when a team might (non-deliberately) make a mistake in the safety process, which might intensify the on-going race and trigger long-term retaliation [19, 24, 25]—we will consider conflict resolution mechanisms such as apology and forgiveness [24, 26] for simmering down the noise effects on the race.

There will be a perception probability (by teams) that an AI disaster will occur wherein all teams lose the race, incurring a significant reduction in their payoff, and this risk-probability will follow a certain probability distribution. It is natural to assume that this probability increases with the frequency with which teams violate the safety requirements (i.e. play UNSAFE). Given the disaster probability, teams might have a different perception of the risk. For example, risk-taking (risk-avoiding) teams might underestimate (overestimate) this risk, leading to more (less) violations of the safety requirements. As in the case of climate change games [30], the perception of risk is a key factor driving the evolutionary outcome (whether or not disaster happens). The main difference is that in climate change models the risk (of failure in avoiding disaster) is collective, whereas in the AI race risk is more individualised: should an individual team ignore safety requirements too often, the more likely their AI product will lead to disaster.

Additionally, assuming that teams have distinct development capacities, i.e. that they might move at different speeds in the race, how does that change the strategic behaviours in the race by the different teams? Stronger teams might want to spend more effort with safety to guarantee no disaster occurs and ensure all the benefits from a powerful AI. On the other hand, weaker teams might want to water down or cease the safety efforts in order to catch up with the stronger teams.

3.2 Peer and group incentives for safety-compliance

This part of the agenda strives to investigate how peer incentives, such as peer punishment and reward, can be efficiently used, whether separately or jointly, for enhancing safety-compliance behaviours in the AI race games developed in the first part. Namely, we will address the following questions:

- What is the influence of using different types of peer incentives (separately or jointly), on safety-agreement compliance in the two-team AI race game?
- Generalising to multi-team agreements, how can peer incentives be used efficiently, taking into account the level of participation in the agreements as well as in a coordinated manner (i.e. pool and coordinated incentives)?

- How should incentives be customised to account for the level of risk (that an AI disaster will occur) as well as for teams with different capabilities of development and of commitment?

We shall start by extending the regimented agreement models for pair-wise and multiparty games, in both one-shot and repeated games, introduced in previous works [17, 18, 24], where interactions occur in three (decision-making) stages: (i) Before the interaction, agents choose whether to propose an agreement; (ii) When receiving a proposal, agents decide whether to accept or reject it; (iii) During the course of the (repeated) interaction, agents choose whether to play SAFE or UNSAFE, depending on whether the agreement was formed; and in the case of multiplayer games, also on how many players committed to following the safety precaution, because a minimum number of committing players may be required for an agreement to be formed and put into effect.

The regimentation assumption entails that agents who accept an agreement proposal yet dishonour it by defecting during the interaction (i.e. fake committers [18]), always honour and pay a compensation. Removing this assumption, another stage or decision point will be added: (iv) After each round of the (repeated) interaction, agents decide whether to use any type of peer incentives, depending on the decisions made by co-players and by themselves in the previous three stages.

This new decision-making stage adds extra layers of complexity, not just on how commitment behaviours are influenced by incentives, but also on how different types of incentives interact. Namely, we need to distinguish between incentives when an agreement is in place (i.e. incentives for agreement fulfillers or for violators), and when agreement is absent (i.e. incentives for mere cooperators or defectors). Considering all strategic behaviours concerning how to use incentives in co-presence in a population, we will examine how these two incentive sorts should be treated differently to foster safety-agreement compliance. Moreover, in order to have a clear understanding of the strength and weakness of each type of incentives, we will start from minimal models where only one type of (peer) incentive is present at a time. Increasingly more complex models, which include other types of incentive, will then be constructed and analysed to see how they interact and influence together the outcome of the race. So doing will provide deeper systematic understanding of how different types of incentives, whether separately or jointly, can be used to achieve high levels of safety compliance.

Rogue AI actors and teams might likely exercise antisocial incentives or bullying [1, 21]. We will consider this possibility in our models. Antisocial or bullying teams most likely refuse to join a safety agreement and are detrimental to cooperation. However, setting up a pre-agreement can provide an efficient solution. One can implement measures to restrict access of non-participating (rogue) teams to AI knowledge [17]; while misbehaving participating teams can be appropriately handled through the agreement's terms and conditions, both by coordinated and institutional punishments [9, 34]. This relates to the AI race, when a small number of major monopolizing actors or teams need to be confronted by the collective as a whole. Furthermore, promoting guilt [29] may be envisaged as a further way to simmer down the AI race, viz. the recent developments regarding Facebook policy change promises

and Cambridge Analytica's declared bankruptcy, by appealing to public discomfort and the chastising of data bullies, again, possibly with the help of crowd sourcing or inside employees.

Our analysis will first examine two-player games since they will permit us to focus on the effects of different types of peer incentives. We will then extend the analysis to multiplayer games, which will include strategic behaviours conditional on the level of participation of teams in a safety-agreement; that is, how many teams agree to comply with the safety requirements. Previous works have shown that increasing the number of players in interaction significantly magnifies the complexity of the evolutionary dynamics and outcome. Firstly, since the number of behavioural equilibria might significantly increase, it is important to study under what conditions and how likely is it that desirable equilibria can be reached; and, moreover, how different types of incentives should be used, separately or jointly, to improve the chance of reaching such desirable outcomes. Secondly, it appears that collective decisions are more difficult to be made for larger groups [15]. We will study how increasing group size influences the probability of safety-agreement formation and, once formed, which are the players' compliant behaviours. Indeed, as shown in previous works for regimented commitments within group interactions, it is crucial to closely monitor the minimum level of participation when deciding whether a commitment should be effectively formed in order to achieve an optimal cooperation outcome [17]. Such a minimum membership requirement can be found in the creation of treaties that address international environmental issues and is also important for AI regulation agreements, as full consensus is rarely reached (as are the cases for cyberspace agreements organized by the UN [36]). Removing the regimentation, we will examine how different types and arrangements of incentives influence the participation level and how this level should be monitored to ensure safety-agreement compliance. Last but not least, in multiplayer agreements a new possibility arises where a group of players might coordinate to provide incentives with more substantial effects (e.g. gang up on free-riders). We will start by applying the existing models of coordinated and pool incentives [9, 34], where a group of incentive providers can share the cost of providing incentives and can decide to actually provide the incentives only if there is sufficient interest in sharing the cost.

In both pair-wise and multiplayer games, we will closely examine how incentives should be used differently, taking into account the diverse risk levels or probabilities that an AI disaster will occur. Does a high risk require stronger (or less so) incentives to ensure a high level of safety compliance? Also, considering that teams might have different capacities (for AI development), the question arises of how incentives should be used differently in light of a given capacity. For instance, stronger punishment might be required against teams who frequently violate safety requirements, especially if they are close to the finish line of the race. In contradistinction, stronger support and reward may be enacted for highly compliant teams, to ensure they win the race with a safe and powerful AI product.

3.3 Institutional incentives for safety-compliance

This final part of the agenda aims to examine how institutional incentives can be efficiently used for enhancing safety-compliance behaviours in the AI race games, and how they interact with peer incentives. The following questions will be addressed:

- What is the influence of different institutional incentives, separately or jointly, on safety-agreement compliance?
- How institutional incentives interact with peer incentives and how these two types can be jointly used to provide an efficient hybrid incentive strategy?

We will study institutional incentives in two scenarios: i) in the absence and ii) in the presence, of peer incentives. The former represents a fully centralised approach to regulating the AI development race, while the latter is a hybrid of centralised and decentralised regulations. In this latter case we explore different ways in which the two types of incentives interact. For example, being autonomous entities, the teams can decide by themselves the type of incentives used to enforce the safety-agreement, i.e. a peer or institutional incentive. It is natural to ask which option would emerge as the preferred one in the population and when. This can be answered by analysing EGT models in which all types of incentives are allowed to be adopted by agents in the population. Since forming a regulating institution is usually costly [22], we will identify when peer incentives are sufficiently efficient so that the institutional setting-up and maintenance cost can be avoided. This outcome is particularly important given the lack of attention so far to peer incentives for behavioural regulation in normative MAS, and moreover, to avoid (institutional) overregulation of AI development from the start.

In addition, a distinctive feature of the institutional setting is that the institution can have access to some global information, such as the current population composition. We explore approaches exploiting this distinctiveness. We will start by extending and generalising the ‘first carrot, then stick’ policy. The challenge is that there will be multiple types of strategic behaviours to incentivise, compared to only two types as in the usual standard institutional incentive models. At each time step, the institution needs to decide whether to incentivise (subject to a given budget) one type or even a subset of distinct types of incentive, depending on their current frequencies in the population. This institutional decision-making process is a complex multi-agent resource allocation problem [16, 20], for which appropriate resource allocation optimisation methods (e.g. from AI literature) can potentially be utilised.

4 PRELIMINARY RESULTS

We have obtained preliminary results for a two-team model of the AI race¹. The race is represented by a repeated game, consisting of a number of AI development rounds, where in each round teams can choose either to play SAFE or UNSAFE. The former choice is not just costlier, but also takes longer (i.e. slower speed of development). The team that wins the race will claim a significant benefit, unless AI disaster occurs. Our analysis considers a population of

teams who can either play SAFE or UNSAFE in all the development rounds, or they can choose to adopt a reciprocal strategy (namely, conditionally SAFE). The teams interact and can adapt their strategy through social learning (i.e. copying the strategy of those who are more successful than them).

In general the analysis points to the direction that when the benefit from winning the race is high, teams that always choose UNSAFE dominate the population for a large range of parameters’ values. This result shows that, in the context of the AI race with repeated interactions, the strategic nature and the outcome are different from those of standard repeated games. In the AI race the rogue teams can move faster in the race by ignoring safety precautions, and reciprocal strategies such as tit-for-tat still lose because of being nice initially. This initial finding suggests that, to drive the race in the more beneficial directions it is important to enact measures that influence (prohibit or accelerate) the speed of AI development of teams, since reciprocal behaviours might not be sufficient to promote cooperative or safety behaviours in this context.

5 ACKNOWLEDGEMENTS

All authors of this work are supported by Future of Life Institute (grant RFP2-154). L. M. Pereira also by FCT/MEC NOVA LINCS PEst UID/CEC/04516/2013.

REFERENCES

- [1] AI-Roadmap-Institute. 2017. Avoiding The Precipice: Race Avoidance in the Development of Artificial General Intelligence. (2017).
- [2] AI-Roadmap-Institute. 2017. Report from the AI Race Avoidance Workshop, Tokyo. (2017).
- [3] Stuart Armstrong, Nick Bostrom, and Carl Shulman. 2016. Racing to the precipice: a model of artificial intelligence development. *AI & society* 31, 2 (2016), 201–206.
- [4] Sandeep Baliga and Tomas Sjöström. 2004. Arms races and negotiations. *The Review of Economic Studies* 71, 2 (2004), 351–369.
- [5] Scott Barrett. 2016. Coordination vs. voluntarism and enforcement in sustaining international environmental cooperation. *Proceedings of the National Academy of Sciences* 113, 51 (2016), 14515–14522.
- [6] Seth D Baum. 2017. On the promotion of safe and socially beneficial artificial intelligence. *AI & SOCIETY* 32, 4 (2017), 543–551.
- [7] Nick Bostrom. 2017. Strategic implications of openness in AI development. *Global Policy* 8, 2 (2017), 135–148.
- [8] Nick Bostrom. 2014. Superintelligence: paths, dangers, strategies.
- [9] Robert Boyd, Herbert Gintis, and Samuel Bowles. 2010. Coordinated punishment of defectors sustains cooperation and can proliferate when rare. *Science* 328, 5978 (2010), 617–620.
- [10] Stephen Cave and Kanta Dihal. 2019. Hopes and fears for intelligent machines in fiction and reality. *Nature Machine Intelligence* 1, 2 (2019), 74.
- [11] Stephen Cave and Seán Ó hÉigeartaigh. 2018. An AI Race for Strategic Advantage: Rhetoric and Risks. In *AAAI/ACM Conference on Artificial Intelligence, Ethics and Society*.
- [12] Xiaojie Chen, Tatsuya Sasaki, Åke Brännström, and Ulf Dieckmann. 2015. First carrot, then stick: how the adaptive hybridization of incentives promotes cooperation. *Journal of The Royal Society Interface* 12, 102 (2015), 20140935.
- [13] Todd L Cherry and David M McEvoy. 2013. Enforcing compliance with environmental agreements in the absence of strong institutions: An experimental analysis. *Environmental and Resource Economics* 54, 1 (2013), 63–77.
- [14] Edward Moore Geist. 2016. It’s already too late to stop the AI arms race? We must manage it instead. *Bulletin of the Atomic Scientists* 72, 5 (2016), 318–321.
- [15] Chaitanya S. Gokhale and Arne Traulsen. 2010. Evolutionary games in the multiverse. *Proc. Natl. Acad. Sci. U.S.A.* 107, 12 (March 2010), 5500–5504. <https://doi.org/10.1073/pnas.0912214107>
- [16] The Anh Han, Simon Lynch, Long Tran-Thanh, and Francisco C. Santos. 2018. Fostering Cooperation in Structured Populations Through Local and Global Interference Strategies. In *IJCAI-ECAI 2018*. 289–295.
- [17] The Anh Han, Luís Moniz Pereira, and Tom Lenaerts. 2016. Evolution of commitment and level of participation in public goods games. *Autonomous Agents and Multi-Agent Systems* (2016), 1–23. <https://doi.org/10.1007/s10458-016-9338-4>

¹The preliminary results can be obtained from the following link: tinyurl.com/y2gwhysy

- [18] T. A. Han, L. M. Pereira, F. C. Santos, and T. Lenaerts. 2013. Good Agreements Make Good Friends. *Scientific reports* 3, 2695 (2013).
- [19] T. A. Han, L. M. Pereira, F. C. Santos, and T. Lenaerts. 2013. Why Is It So Hard to Say Sorry: The Evolution of Apology with Commitments in the Iterated Prisoner's Dilemma. In *IJCAI'2013*. AAAI Press, 177–183.
- [20] The Anh Han and Long Tran-Thanh. 2018. Cost-effective external interference for promoting the evolution of cooperation. *Scientific reports* 8 (2018).
- [21] Benedikt Herrmann, Christian Thöni, and Simon Gächter. 2008. Antisocial Punishment Across Societies. *Science* 319, 5868 (March 2008), 1362–1367.
- [22] Christian Hilbe, Arne Traulsen, Torsten Röhl, and Manfred Milinski. 2014. Democratic decisions establish stable authorities that overcome the paradox of second-order punishment. *PNAS* 111, 2 (2014), 752–756.
- [23] J. Hofbauer and K. Sigmund. 1998. *Evolutionary Games and Population Dynamics*. Cambridge University Press.
- [24] Luis A Martinez-Vaquero, The Anh Han, Luis Moniz Pereira, and Tom Lenaerts. 2015. Apology and forgiveness evolve to resolve failures in cooperative agreements. *Scientific reports* 5, 10639 (2015).
- [25] Luis A Martinez-Vaquero, The Anh Han, Luis Moniz Pereira, and Tom Lenaerts. 2017. When agreement-accepting free-riders are a necessary evil for the evolution of cooperation. *Scientific reports* 7, 1 (2017), 2478.
- [26] Michael McCullough. 2008. *Beyond revenge: The evolution of the forgiveness instinct*. John Wiley & Sons.
- [27] R. M. Nesse. 2001. *Evolution and the capacity for commitment*. Russell Sage.
- [28] Dennis Pamlin and Stuart Armstrong. 2015. Global challenges: 12 risks that threaten human civilization. *Global Challenges Foundation, Stockholm* (2015).
- [29] Luis Moniz Pereira, Tom Lenaerts, Luis A Martinez-Vaquero, and The Anh Han. 2017. Social manifestation of guilt leads to stable cooperation in multi-agent systems. In *AAMAS*. 1422–1430.
- [30] Francisco C. Santos and Jorge M. Pacheco. 2011. Risk of collective failure provides an escape from the tragedy of the commons. *PNAS* 108, 26 (2011), 10421–10425.
- [31] Carl Shulman and Stuart Armstrong. 2009. Arms control and intelligence explosions. In *7th European Conference on Computing and Philosophy (ECAP)*, Bellaterra, Spain, July. 2–4.
- [32] Karl Sigmund. 2010. *The Calculus of Selfishness*. Princeton University Press.
- [33] K Sigmund, C Hauert, and M Nowak. 2001. Reward and punishment. *P Natl Acad Sci USA* 98, 19 (2001), 10757–10762.
- [34] K. Sigmund, H. De Silva, A. Traulsen, and C. Hauert. 2010. Social learning promotes institutions for governing the commons. *Nature* 466 (2010), 7308.
- [35] Munindar P Singh. 2013. Norms as a basis for governing sociotechnical systems. *ACM Transactions on Intelligent Systems and Technology (TIST)* 5, 1 (2013), 21.
- [36] Mariarosaria Taddeo and Luciano Floridi. 2018. Regulate artificial intelligence to avert cyber arms race. *Nature* 556, 7701 (2018), 296–298.
- [37] Max Tegmark. 2017. *Life 3.0: Being human in the age of artificial intelligence*. Knopf.
- [38] Vitor V Vasconcelos, Francisco C Santos, and Jorge M Pacheco. 2013. A bottom-up institutional approach to cooperative governance of risky commons. *Nature Climate Change* 3, 9 (2013), 797.