



Full Length Article



Trust AI regulation? Discerning users are vital to build trust and effective AI regulation

Zainab Alalawi ^a, Paolo Bova ^{a, }, Theodor Cimpeanu ^{b, }, Alessandro Di Stefano ^{a, },
 Manh Hong Duong ^{c, }, Elias Fernández Domingos ^{d, e, }, The Anh Han ^{a, }, ^{*},
 Marcus Krellner ^{b, }, Ndidi Bianca Ogbo ^{a, }, Simon T. Powers ^{f, },
 Filippo Zimmaro ^{g, h, }

^a School Computing, Engineering and Digital Technologies, Teesside University, United Kingdom

^b School of Mathematics and Statistics, University of St Andrews, United Kingdom

^c School of Mathematics, University of Birmingham, United Kingdom

^d Machine Learning Group, Université libre de Bruxelles, Belgium

^e AI Lab, Vrije Universiteit Brussel, Belgium

^f Division of Computing Science and Mathematics, University of Stirling, United Kingdom

^g Department of Mathematics, University of Bologna, Italy

^h Department of Computer Science, University of Pisa, Italy

A B S T R A C T

There is general agreement that some form of regulation is necessary both for AI creators to be incentivised to develop trustworthy systems, and for users to actually trust those systems. But there is much debate about what form these regulations should take and how they should be implemented. Most work in this area has been qualitative, and has not been able to make formal predictions. Here, we propose that evolutionary game theory can be used to quantitatively model the dilemmas faced by users, AI creators, and regulators, and provide insights into the possible effects of different regulatory regimes. We show that achieving safe AI and user trust requires regulators to be incentivised to regulate effectively. We demonstrate two effective mechanisms. In the first, governments can recognise and reward regulators that do a good job. In that case, if the AI technology is not too risky, some level of safe development and user trust evolves. In the second mechanism, users can condition their trust decision on the effectiveness of the regulators. This leads to effective regulation, and consequently the development of trustworthy AI and user trust, provided that the cost of implementing regulations is not too high. Our findings highlight the importance of considering the effect of different regulatory regimes from an evolutionary game theoretic perspective.

1. Introduction

Debates are taking place across the world about what kind of regulation should apply to the development of artificial intelligence (AI) systems [1–5]. Governments want *trustworthy* and safe AI systems to be developed, and for users to actually *trust* these systems – this is exemplified by the aims of the EU AI act [6]. Regulation is typically assumed to be the way to achieve this [7]. Economic history supports this assumption; safety critical systems such as automotive vehicles and medical devices are indeed regulated in this way. But it is less clear what forms the regulations should take and who should implement them. See, for example, proposals in the EU compared to the USA on the amount of restrictions that should be placed on AI creators [2,8]. Moreover, there is the key question

* Corresponding author.

E-mail address: T.Han@tees.ac.uk (T.A. Han).

<https://doi.org/10.1016/j.amc.2025.129627>

Received 13 June 2025; Received in revised form 1 July 2025; Accepted 5 July 2025

Available online 15 July 2025

0096-3003/© 2025 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).




Fig. 1. Core features. The figure schematically illustrates the core features of the three-population model of AI governance. Users (left) can either trust (T) or not trust (N) the AI system, in which case they do not adopt the system and get zero benefit. Creators (centre) can either defect by creating unsafe AI products (D) or cooperate by creating safe ones (C), which entails additional costs. The government (right) can choose to employ regulators (C) at some costs. Those will punish defecting creators with further costs to themselves.

of *who* should create and enforce the regulations? Should the regulations be *public* or *private* [9]? That is, should they be created by government bodies, or by industrial organisations and accrediting bodies such as the IEEE? And once they have been created, who will enforce them? Enforcement could involve actions such as auditing source code, training runs and training data, and monitoring for consumer complaints. This could be done by government-sector organisations, but it could also be done by third-party auditors, and there are ongoing debates about the benefits and drawbacks of each approach [10,11].

To help governments choose the most suitable kind of regulatory framework, we need to be able to predict the effects of different regulatory systems. Most of the discourse around this at the moment is qualitative and does not lead to formal predictions [11–15]. This limits the ability of governments, technology companies, and citizens to foresee what the effects of different regulatory systems might be [16]. There is a small amount of literature on AI race and AI governance modelling [17–19], including from the perspective of (evolutionary) game theory [20–25], but this has not considered how different regulatory mechanisms influence both user trust and the compliance of companies with AI ethical and safety considerations. To address this, we propose that evolutionary game theory [26] can be used to formally model the effects of different regulatory systems in terms of their incentives on tech companies, end users, and regulators.

We present a framework for formalising the strategic interaction between users, AI system creators, and regulators as a game (Fig. 1). This game captures three key decisions that these actors face:

1. **Users:** do they trust and hence use an AI system or not?
2. **AI creators:** do they follow a safety optimal development path in compliance with regulations, or do they pursue a competitive development path that violates regulations in a race to the market?
3. **Regulators:** do they invest in monitoring AI system creators and enforcing the regulations effectively, or do they cut corners to save the costs of doing this?

This highlights the dilemmas facing users, creators, and regulators. Users can benefit from using an AI system, but also run the risk that the system may not act in their best interest, i.e. may not be trustworthy [27]. This follows from the fact that AI creators are themselves in competition with each other, as highlighted by the current “AI race” to develop artificial general intelligence [11,17,28–30]. Consequently, we cannot assume that creators will always act in the best interests of their users by complying with regulations and developing systems worthy of user trust. Finally, regulators themselves may be self-interested. This may occur when governments delegate the enforcement of regulations to other actors, such as private audit firms [10,13,31]. This kind of delegation may reduce government costs, but it also introduces a principal agent problem [32]: regulators are themselves agents with their own profit-maximising goals.

To analyse the model, we use the methods of evolutionary game theory [26,33]. Evolutionary game theory is based on the idea that agents can learn behaviours that benefit them by social learning, i.e. by copying the behaviour of other agents in their population

that are doing better than themselves. This avoids the need to assume that the agents are fully rational and have complete information. In our model, we consider three populations corresponding to the three actors: users, creators, and regulators.

Our analysis demonstrates how incentives for regulators are important. Governments desire that all creators produce safe and trustworthy AI systems, and all users trust these systems. Such a state cannot be reached if regulators that do their job properly cannot be distinguished from regulators that cut corners. This holds regardless of the severity of punishment for creators that develop unsafe systems.¹

We consider two possible institutional solutions to this problem. First, we show that if governments can provide rewards to regulators that do a good job, and use of the AI system is not too risky for users, then some level of safe development and trust by users occurs. We then consider an alternative solution, where users may condition their trust decision on the effectiveness of the regulators, for example, where information about the past performance of regulators is available. This leads to effective regulation, and consequently the development of safe AI and user trust, provided that the cost of implementing regulations is not too high.

In the next section, we describe the models and methods, including a three-population model of AI governance and the evolutionary methods for analysing the model from both finite and infinite population perspectives. Results for each type of analysis and Discussion sections will follow.

2. Models and methods

2.1. Three population model of AI governance

We start by constructing a **baseline model** of an AI development ecosystem (following Figure 1 of [7]). This is similar to current models of regulations for information systems, e.g. the General Data Protection Regulations (GDPR) and the AI Act of the European Union. The model involves three populations representing the three actors in the ecosystem: AI users, AI system creators, and regulators. In each population, individuals can choose different options (also called strategies). A user can decide to trust (T) or not (N) an AI system: This is a combination of trust in the regulator and the creator. The creator either complies (C) with the rules set out by the regulators or not (D). The regulator chooses whether to create rules and enforce compliance (C) or not (D).

The individual payoff earned in any one encounter (also called a game) depends on the strategy of the participating individuals. In each game, one user, one creator, and one regulator participate. If the user trusts and adopts an AI system when both the creator and the regulator cooperate (by complying and enforcing, respectively), the user benefits significantly from AI adoption, denoted by b_U . However, if AI adoption occurs when the creator defects (does not comply with regulations), the user is affected by unsafe AI, gaining a reduced or even negative benefit, denoted by $\epsilon \times b_U$, where $\epsilon \in [-\infty, 1]$. This parameter, ϵ , also represents a *risk factor* that users take when trusting and adopting the AI system. For the AI creator, since it takes more time and effort to comply with the precautionary requirements, we assume that playing C requires an extra cost c_P , compared to playing D (the cost for this is normalised to 0). Regulators earn a benefit b_R when the user trusts and adopts AI. This corresponds to regulation being funded by taxes on the sales of AI products, or by governments investing more in regulation when there is more uptake of AI. Since it is costly to create rules and technologies to prevent and detect unsafe development, we assume that playing C requires an extra cost c_R , compared to playing D (the cost for playing D is normalised to 0). If a creator defects (D) and does not comply with regulations and they interact with a cooperative regulator (playing C), then the creator suffers institutional punishment from the regulator at a cost of u . In this case, the cooperative regulator pays an additional cost v to administer this punishment. The model strategies and parameters are summarised in Table 1.

In the general case, this leads to a frequency-dependent dynamics in each of the three populations, where the dynamical process of strategy adoption assumes that those individuals that are more fit are more often imitated by their peers, an adaptive scheme akin to social learning (see Methods below). As shown in our analysis below, in this system, regulators are always better off not cooperating, and trust by users rarely develops. Thus, below we consider two extended models with institutional arrangements [32] that might enable investment in high quality regulation, safe AI development, and users' trust in these safe systems.

1. First, we consider that cooperative regulators are rewarded an amount of b_{fo} if they catch defecting creators (when users trust and adopt AI). Fig. 2 illustrates this conceptually, while the corresponding payoff matrix is shown in Table 2.
2. Second, we assume that regulators' reputation is publicly available and users can act conditionally on whether the regulators' reputation is good or not; see Fig. 3 and the payoff matrix in Table 3. This allows us to consider a CT strategy: these are discerning users who do not trust when the regulator's reputation is bad.

2.2. Methods

2.2.1. Stochastic dynamics for finite populations

The baseline model. We consider three different well-mixed populations of Users (U), Creators (C) and Regulators (R) of sizes, respectively N_U , N_C and N_R . Let x be the fraction of users that trust the AI system. Let y and z be respectively the fraction of

¹ We note that trustworthiness and safety are two distinct concepts, with safety focusing more on technical issues, while trustworthiness covers a broader range of concerns including ethics, fairness and accountability [81]. For the remainder of this paper, we refer to safety for ease of exposition, while noting that because our models operate at a high level of abstraction, the same concepts would apply to regulations governing trustworthiness more generally.

Table 1

AI Governance modelling: The top table summaries the key parameters in the models. The bottom table presents the baseline model. (User *Us.*, Creator *Cr* and Regulator *Re.*)

Parameter description	Symbol
Users trust (T) or do not (N) trust an AI system — this is a combination of trust in regulators and trust in creators	T, N
Creators comply with the rules (C) or not (D) as set out by the Regulators	C, D
Regulators create rules and ensure compliance (C) or not (D)	C, D
Benefit users get from trusting and adopting the AI system	b_U
Fraction of user benefit when creators play D , where ε in $[-\infty, 1]$, also referred to as the (inverse) risk factor users take when adopting an AI system	ε
Benefit the creator gets from selling the product	b_P
Additional cost of creating safe AI (the cost of creating unsafe AI is normalised to 0)	c_P
Funding for regulators (which is only generated upon AI adoption)	b_R
The cost of developing rules and technologies to capture unsafe AI (the cost of not doing this is normalised to 0)	c_R
The cost of institutional punishment	v
The impact of institutional punishment	u
Reward granted to cooperative regulators	b_{f0}

Strategies			Payoffs		
Us.	Cr	Re	User	Creator	Regulator
T	C	C	b_U	$b_P - c_P$	$b_R - c_R$
T	C	D	b_U	$b_P - c_P$	b_R
T	D	C	εb_U	$b_P - u$	$b_R - c_R - v$
T	D	D	εb_U	b_P	b_R
N	C	C	0	$-c_P$	$-c_R$
N	C	D	0	$-c_P$	0
N	D	C	0	0	$-c_R$
N	D	D	0	0	0




Fig. 2. Reward for regulators. The figure shows the changes to the original model. Regulators gain extra reward for capturing unsafe development.

Creators and Regulators that cooperate. Each game involves an individual randomly drawn from each population. The fitness that a user, creator and regulator obtains in each game is respectively given by

$$f_{X \in \{T, N\}}^U = yzP_{XC}^U + (1-y)zP_{XD}^U + y(1-z)P_{XC}^U + (1-y)(1-z)P_{XD}^U, \quad (1)$$

$$f_{Y \in \{C, D\}}^C = xzP_{YC}^C + (1-x)zP_{YC}^C + x(1-z)P_{TD}^C + (1-x)(1-z)P_{ND}^C, \quad (2)$$

$$f_{Z \in \{C, D\}}^R = xyP_{TCZ}^R + (1-x)zP_{NCZ}^R + x(1-y)P_{TDZ}^R + (1-x)(1-y)P_{NDZ}^R. \quad (3)$$

Table 2

AI Governance extended model with incentives for regulators (User $Us.$, Creator Cr and Regulator Re), where regulators are rewarded for capturing unsafe creators (when users trust and adopt).

Strategies			Payoffs		
Us.	Cr	Re	User	Creator	Regulator
T	C	C	b_U	$b_P - c_P$	$b_R - c_R$
T	C	D	b_U	$b_P - c_P$	b_R
T	D	C	ϵb_U	$b_P - u$	$b_R - c_R - v + b_{fo}$
T	D	D	ϵb_U	b_P	b_R
N	C	C	0	$-c_P$	$-c_R$
N	C	D	0	$-c_P$	0
N	D	C	0	0	$-c_R$
N	D	D	0	0	0




Fig. 3. Conditional trust. The figure shows how the state of regulation changes the behaviour of users, which impacts the whole system.

We now derive the fitness/average payoffs for each model. First, for the baseline model, from (1) and the payoff Table 1, we calculate explicitly the difference of the fitness between two strategies in users:

$$\begin{aligned}
 f_T^U - f_N^U &= yz(P_{TCC}^U - P_{NCC}^U) + (1-y)z(P_{TDC}^U - P_{NDC}^U) + y(1-z)(P_{TCD}^U - P_{NCD}^U) + (1-y)(1-z)(P_{TDD}^U - P_{NDD}^U) \\
 &= yzb_U + (1-y)z(\epsilon b_U) + y(1-z)b_U + (1-y)(1-z)(\epsilon b_U) \\
 &= b_U \left(yz + \epsilon(1-y)z + y(1-z) + \epsilon(1-y)(1-z) \right) \\
 &= b_U(y + \epsilon(1-y)).
 \end{aligned} \tag{4}$$

Similarly, the difference of the fitness between two strategies in creators is

$$\begin{aligned}
 f_C^C - f_D^C &= xz(P_{TCC}^C - P_{TDC}^C) + (1-x)z(P_{NCC}^C - P_{NDC}^C) + x(1-z)(P_{TCD}^C - P_{TDD}^C) + (1-x)(1-z)(P_{NCD}^C - P_{NDD}^C) \\
 &= xz(b_P - c_P - (b_P - u)) + (1-x)z(-c_P) + x(1-z)(b_P - c_P - b_P) + (1-x)(1-z)(-c_P) \\
 &= xz(u - c_P) + (1-x)z(-c_P) + x(1-z)(-c_P) + (1-x)(1-z)(-c_P) \\
 &= -c_P + uxz.
 \end{aligned} \tag{5}$$

Finally, the difference of the fitness between two strategies in regulators is

Table 3

AI Governance extended model with conditional trust (User *Us.*, Creator *Cr* and Regulator *Re*). In this model, regulators' reputation are assumed to be publicly available before the game, and users can act conditionally on it: CT does not trust if the regulator has a bad reputation.

Strategies			Payoffs		
Us.	Cr	Re	User	Creator	Regulator
CT	C	C	b_U	$b_P - c_P$	$b_R - c_R$
CT	C	D	0	$-c_P$	0
CT	D	C	ϵb_U	$b_P - u$	$b_R - c_R - v + b_{fo}$
CT	D	D	0	0	0
N	C	C	0	$-c_P$	$-c_R$
N	C	D	0	$-c_P$	0
N	D	C	0	0	$-c_R$
N	D	D	0	0	0

$$\begin{aligned}
 f_C^R - f_D^R &= xy(P_{TCC}^R - P_{TCD}^R) + (1-x)y(P_{NCC}^R - P_{NCD}^R) + x(1-y)(P_{TDC}^R - P_{TDD}^R) + (1-x)(1-y)(P_{NDC}^R - P_{NDD}^R) \\
 &= xy(b_R - c_R - b_R) + (1-x)y(-c_R) + x(1-y)(b_R - c_R - v + b_R) + (1-x)(1-y)(-c_R) \\
 &= xy(-c_R) + (1-x)y(-c_R) + x(1-y)(-c_R - v) + (1-x)(1-y)(-c_R) \\
 &= -c_R - x(1-y)v.
 \end{aligned} \tag{6}$$

We notice that $f_C^R - f_D^R$ is always negative. Thus, it is always better off for the regulator to defect.

The extended model with regulators' incentive. Now, considering that a reward b_{fo} is provided to regulators for capturing unsafe creators and if they catch defective creators (when users adopt only), see Fig. 2. In this case, only the payoff difference $f_C^R - f_D^R$ is changed to

$$f_C^R - f_D^R = -c_R + x(1-y)(b_{fo} - v). \tag{7}$$

The extended model with conditional trust. Finally, considering that users can make a conditional decision based on the regulator's reputation (see Fig. 3), we have

$$\begin{aligned}
 f_T^U - f_N^U &= b_U z(y + \epsilon(1-y)), \\
 f_C^C - f_D^C &= -c_P + uxz, \\
 f_C^R - f_D^R &= -c_R + b_{RX} + (b_{fo} - v)x(1-y).
 \end{aligned} \tag{8}$$

For a finite population setting, at each time step, a randomly selected individual A, with fitness f_A , may adopt a different strategy by imitating a randomly chosen individual B from the same population (with fitness f_B) with probability given by the Fermi distribution [34,35]

$$p = [1 + e^{-\beta(f_B - f_A)}]^{-1},$$

where $\beta \geq 0$ is the strength of selection. $\beta = 0$ corresponds to neutral drift where imitation decisions are random, while for large $\beta \rightarrow \infty$, the imitation decision becomes increasingly deterministic.

In the absence of mutations or exploration, the end states of evolution are inevitably monomorphic: once such a state is reached, it cannot be escaped through imitation. We thus further assume that with a certain mutation probability, an agent switches randomly to a different strategy without imitating another agent. In the limit of small mutation rates, the dynamics will proceed with, at most, two strategies in the population, such that the behavioural dynamics can be conveniently described by a Markov chain, where each state represents a monomorphic population, whereas the transition probabilities are given by the fixation probability of a single mutant [36–38]. The resulting Markov chain has a stationary distribution, which characterises the average time the population spends in each of these monomorphic end states.

Now, the probability to change the number k of agents using strategy A by $\pm$ one in each time step can be written as (Z is the population size) [39]

$$T^\pm(k) = \frac{Z-k}{Z} \frac{k}{Z} [1 + e^{\mp\beta[f_A(k) - f_B(k)]}]^{-1}. \tag{9}$$

The fixation probability of a single mutant with a strategy A in a population of $(Z-1)$ agents using B is given by [37,39]

$$\rho_{B,A} = \left(1 + \sum_{i=1}^{Z-1} \prod_{j=1}^i \frac{T^-(j)}{T^+(j)} \right)^{-1}. \tag{10}$$

The transition matrix Λ corresponding to the set of $\{1, \dots, s\}$ strategies is given by:

$$\Lambda_{ij,j \neq i} = \frac{\rho_{ji}}{3} \quad \text{and} \quad \Lambda_{ii} = 1 - \sum_{j=1, j \neq i}^s \Lambda_{ij}. \quad (11)$$

Fixation probability ρ_{ij} denotes the likelihood that a population transitions from a state i to a different state j when a mutant of one of the populations adopts an alternate strategy s . The fixation probability is divided by the number of populations (3) representing the interaction of three players at a time [40,41].

2.2.2. Population dynamics for infinite populations: the multi-population replicator dynamics

In this section, we recall the framework of the replicator dynamics for multi-populations [42,43]. To describe the dynamics, we consider a set of m different populations (m is some positive integer), which are infinitely large and well-mixed. Each population i , $i = 1, \dots, m$, consists of n_i (n_i is some positive integer) different strategies (types). Let x_{ij} , $1 \leq i \leq m$, $1 \leq j \leq n_i$, be the frequency of the strategy j in the population i . We denote by $x_i = (x_{ij})_{j=1}^{n_i}$, which is the collection of all strategies in the population i , and $x = (x_1, \dots, x_m)$, which is the collection of all strategies in all populations.

For each $i \in \{1, \dots, m\}$ and $j \in \{1, \dots, n_i\}$, let $f_{ij}(x)$ be the fitness (reproductive rate) of the strategy j in the population i . This fitness is obtained when the strategy j interacts with all other strategies in all populations; thus, it depends on all the strategies in the populations. The average fitness of the population i is defined by

$$\bar{f}_i(x) = \sum_{j=1}^{n_i} x_{ij} f_{ij}(x).$$

The multi-population replicator dynamics is then given by

$$\dot{x}_{ij} = x_{ij}(f_{ij}(x) - \bar{f}_i(x)), \quad 1 \leq i \leq m, \quad 1 \leq j \leq n_i. \quad (12)$$

This is in general an ODE system of $\sum_{i=1}^m n_i$ equations. Noting, however that since $\sum_{j=1}^{n_i} x_{ij} = 1$ for all $i = 1, \dots, m$, we can reduce the above system to a system of $\sum_{i=1}^m n_i - m$ equations.

Now we focus on the case when there are two strategies in each population (which is the case for our models of AI governance and trust in the present paper), that is $n_i = 2$ for all $i = 1, \dots, m$. Let η_i be the frequency of the first strategy in the population i , $i = 1, \dots, m$ (thus $1 - \eta_i$ will be the frequency of the second strategy in the population i), let $\eta = (\eta_1, \dots, \eta_m)$. Let $f_{1i}(\eta)$ and $f_{2i}(\eta)$ be the fitness of the first and second strategy in the population i . Since

$$\bar{f}_i(\eta) = \eta_i f_{1i}(\eta) + (1 - \eta_i) f_{2i}(\eta),$$

we have

$$f_{1i}(\eta) - \bar{f}_i(\eta) = f_{1i}(\eta) - (\eta_i f_{1i}(\eta) + (1 - \eta_i) f_{2i}(\eta)) = (1 - \eta_i)(f_{1i}(\eta) - f_{2i}(\eta)).$$

Thus we obtain the following system of equations

$$\dot{\eta}_i = \eta_i(1 - \eta_i)(f_{1i}(\eta) - f_{2i}(\eta)), \quad i = 1, \dots, m. \quad (13)$$

This is a system of m coupled nonlinear ordinary differential equations for m variables.

In the subsequent sections, we employ (13) to our models of AI governance trust, where the fitness is computed from the payoff matrix constructed in the models, see Tables 1-2-3, where we assume that the payoffs are directly translated to the biological fitnesses (infinite strength of selection).

3. Equilibrium analysis in infinite populations

3.1. The baseline model

We consider three different well-mixed populations of Users (U), Creators (C) and Regulators (R). Let x , y and z be the frequencies of users, creators and regulators that play T, C and C, respectively. The replicator dynamics can be written as follows (see again Table 1)

$$\dot{x} = x(1 - x)(f_T^U - f_N^U) = x(1 - x)b_U(y + \varepsilon(1 - y)), \quad x(0) = x_0 \quad (14a)$$

$$\dot{y} = y(1 - y)(f_C^C - f_D^C) = y(1 - y)(-c_P + uxz), \quad y(0) = y_0 \quad (14b)$$

$$\dot{z} = z(1 - z)(f_C^R - f_D^R) = z(1 - z)(-c_R - x(1 - y)v), \quad z(0) = z_0 \quad (14c)$$

where $(x_0, y_0, z_0) \in [0, 1]^3$ is the initial data. This is a complex, non-linear system of three coupled ordinary differential equations (ODEs) for three variables. In general, the dynamical solutions $x(t), y(t), z(t)$ may show different behaviour depending on the values of the parameters and the initial data.

The equilibrium points in our models (including other models below) represent different long-term states of the AI governance ecosystem. For example, an equilibrium where users consistently trust AI, creators consistently comply with regulations, and regulators consistently enforce them represents a desirable, well-functioning system. Conversely, an equilibrium where users distrust AI, creators do not comply with regulations, and/or regulators do not enforce rules indicates a problematic scenario. The stability of these equilibria is crucial: a stable equilibrium implies that if the system is slightly perturbed from that state, it will tend to return to it, indicating a resilient outcome. An unstable equilibrium, on the other hand, means that even small changes can push the system towards a different, potentially undesirable state. Hence, understanding the stability of different equilibria is important for designing efficient interventions that can drive the AI ecosystem towards beneficial outcomes.

3.1.1. Stationary solutions (equilibrium points)

Stationary solutions of (14) are points $(x^*, y^*, z^*) \in [0, 1]^3$ that make the right hand side of (14) vanish. It is clear that the vertices of the unit cube $(x, y, z) \in \{0, 1\}^3$ are stationary solutions (we call them vertex equilibria). There may exist other, non-vertex, stationary solutions depending on the sign of ε and the relation between c_P and u .

If $\varepsilon > 0$, then $y + \varepsilon(1 - y) > 0$ for all $y \in [0, 1]$. It implies that there are no other stationary solutions (except in the very specific case where $u = c_P$, in which $(x^*, y^*, z^*) \in \{0, 1\} \times [0, 1] \times \{0, 1\}$ is a stationary solution).

If $\varepsilon < 0$, and $0 < c_P \leq u$, then other, non-vertex, stationary solutions are given by

$$y^* = \frac{\varepsilon}{\varepsilon - 1}, \quad -c_P + ux^*z^* = 0, \quad z^* = 1,$$

which is

$$x^* = \frac{c_P}{u}, \quad y^* = \frac{\varepsilon}{\varepsilon - 1}, \quad z^* = 1. \quad (15)$$

3.1.2. Stability analysis

Next, we study the stability of the equilibria. To this end, we define

$$F(x, y, z) = (F_1(x, y, z), F_2(x, y, z), F_3(x, y, z)), \quad F_1(x, y, z) = x(1 - x)b_U(y + \varepsilon(1 - y)),$$

$$F_2(x, y, z) = y(1 - y)(-c_P + uxz), \quad F_3(x, y, z) = z(1 - z)(-c_R - x(1 - y)v).$$

Then (14) can be reformulated as

$$\dot{x} = F_1(x, y, z), \quad \dot{y} = F_2(x, y, z), \quad \dot{z} = F_3(x, y, z).$$

The Jacobian matrix is

$$DF(x, y, z) = \begin{pmatrix} \frac{\partial F_1}{\partial x} & \frac{\partial F_1}{\partial y} & \frac{\partial F_1}{\partial z} \\ \frac{\partial F_2}{\partial x} & \frac{\partial F_2}{\partial y} & \frac{\partial F_2}{\partial z} \\ \frac{\partial F_3}{\partial x} & \frac{\partial F_3}{\partial y} & \frac{\partial F_3}{\partial z} \end{pmatrix} = \begin{pmatrix} b_U(1 - 2x)(y + \varepsilon(1 - y)) & b_U x(1 - x)(1 - \varepsilon) & 0 \\ uy(1 - y)z & (1 - 2y)(-c_P + uxz) & uy(1 - y)x \\ -z(1 - z)(1 - y)v & x vz(1 - z) & (1 - 2z)(-c_R - x(1 - y)v) \end{pmatrix}.$$

We recall that an equilibrium (x^*, y^*, z^*) is asymptotically stable if all eigenvalues of $DF(x^*, y^*, z^*)$ have negative real parts.

3.1.3. Stability analysis for vertex equilibria

We now determine the stability for vertex-equilibria, $(x^*, y^*, z^*) \in \{0, 1\}^3$. For these equilibria, $DF(x^*, y^*, z^*)$ is always a diagonal matrix and is explicitly given by

$$DF(x^*, y^*, z^*) = \begin{pmatrix} b_U(1 - 2x^*)(y^* + \varepsilon(1 - y^*)) & 0 & 0 \\ 0 & (1 - 2y^*)(-c_P + ux^*z^*) & 0 \\ 0 & 0 & (1 - 2z^*)(-c_R - x^*(1 - y^*)v) \end{pmatrix}.$$

For instance at the origin $(x^*, y^*, z^*) = (0, 0, 0)$ we have

$$DF(x, y, z) = \begin{pmatrix} \varepsilon b_U & 0 & 0 \\ 0 & -c_P & 0 \\ 0 & 0 & -c_R \end{pmatrix}.$$

For another example at $(x^*, y^*, z^*) = (1, 0, 0)$ then

$$DF(x, y, z) = \begin{pmatrix} -\varepsilon b_U & 0 & 0 \\ 0 & -c_P & 0 \\ 0 & 0 & -c_R \end{pmatrix}.$$

Since for vertex equilibria, $DF(x^*, y^*, z^*)$ is diagonal and its eigenvalues read

$$\lambda_1 = b_U(1 - 2x^*)(y^* + \varepsilon(1 - y^*)), \quad \lambda_2 = (1 - 2y^*)(-c_P + ux^*z^*), \quad \lambda_3 = (1 - 2z^*)(-c_R - x^*(1 - y^*)v).$$

Thus (x^*, y^*, z^*) is stable only if

$$\lambda_1 < 0, \quad \lambda_2 < 0, \quad \lambda_3 < 0.$$

Since $(-c_R - x^*(1 - y^*)v) < 0$, it follows that $\lambda_3 < 0$ iff $z^* = 0$. For $z^* = 0$, $\lambda_2 = -c_P(1 - 2y^*) < 0$ iff $y^* = 0$. Then $\lambda_1 = \varepsilon b_U(1 - 2x^*)$. For $\varepsilon > 0$, then $\lambda_1 < 0$ iff $x^* = 1$. On the other hand, if $\varepsilon < 0$, then $\lambda_1 < 0$ iff $x^* = 0$. In conclusion, for vertex equilibria, $(x^*, y^*, z^*) \in \{0, 1\}^3$:

- when $\varepsilon > 0$, the equilibrium $(x^*, y^*, z^*) = (1, 0, 0)$ is stable, thus the users will fully trust (even if the creators and the regulators defect).
- when $\varepsilon < 0$, the origin $(x^*, y^*, z^*) = (0, 0, 0)$ is stable, both regulators and creators defect and there is no trust from users.

3.1.4. Stability analysis for the non-vertex equilibrium

For the non-vertex equilibrium (15), the Jacobian matrix is of the form

$$DF = \begin{pmatrix} 0 & b_U \frac{c_P}{u} (1 - \frac{c_P}{u})(1 - \varepsilon) & 0 \\ -u \frac{\varepsilon}{(\varepsilon - 1)^2} & 0 & -u \frac{\varepsilon}{(\varepsilon - 1)^2} \frac{c_P}{u} \\ 0 & 0 & c_R + \frac{c_P}{u} \frac{1}{1 - \varepsilon} v \end{pmatrix}.$$

Thus it has one positive (recalling that $\varepsilon < 0$ in this case) real root, which is the last diagonal entry. Therefore, the non-vertex equilibrium is not stable.

Conclusion: our analysis of the baseline model (without incentives for regulators) showed that in all case, both creators and regulators will defect. Moreover, when adopting unsafe AI has a low risk (i.e. positive ε), users will choose to trust and adopt AI. Otherwise, they will not trust AI systems.

3.2. Extended model: incentive for regulators to cooperate

For the extended model, using $f_C^R - f_D^R$ calculated in (7) in (14), the replicator dynamics become (see again Table 2)

$$\dot{x} = x(1 - x)(f_T^U - f_N^U) = x(1 - x)b_U(y + \varepsilon(1 - y)), \quad x(0) = x_0 \quad (16a)$$

$$\dot{y} = y(1 - y)(f_C^C - f_D^C) = y(1 - y)(-c_P + uxz), \quad y(0) = y_0 \quad (16b)$$

$$\dot{z} = z(1 - z)(f_C^R - f_D^R) = z(1 - z)(-c_R + x(1 - y)(b_{fo} - v)), \quad z(0) = z_0 \quad (16c)$$

where $(x_0, y_0, z_0) \in [0, 1]^3$ is the initial data. Practically, compared to the baseline model, the only difference is that we allowed v to be negative. This is again a complex nonlinear coupled system of three ordinary differential equations. Its solutions $x(t), y(t), z(t)$ may exhibit different behaviour depending on the values of the parameters and the initial data, see Fig. 4.

The vertex equilibria in the baseline model are obviously still equilibria for this extended model. In addition, due to the presence of the incentive b_{fo} there may be another internal equilibrium

$$x^* = \frac{c_R(1 - \varepsilon)}{b_{fo} - v}, \quad y^* = \frac{\varepsilon}{\varepsilon - 1}, \quad z^* = \frac{c_P(b_{fo} - v)}{u(1 - \varepsilon)}, \quad (17)$$

if $\varepsilon < 0$ and the parameters $c_P, c_R, (b_{fo} - v)$ are such that $0 < x^*, y^*, z^* < 1$ (in particular, $b_{fo} < v$). The Jacobian matrix is given by

$$DF(x, y, z) = \begin{pmatrix} b_U(1 - 2x)(y + \varepsilon(1 - y)) & b_U x(1 - x)(1 - \varepsilon) & 0 \\ uy(1 - y)z & (1 - 2y)(-c_P + uxz) & uy(1 - y)x \\ -z(1 - z)(1 - y)(v - b_{fo}) & x(v - b_{fo})z(1 - z) & (1 - 2z)(-c_R - x(1 - y)(v - b_{fo})) \end{pmatrix}.$$

We are now interested in determining the stability for vertex equilibria, $(z^*, y^*, x^*) \in \{0, 1\}^3$. For these equilibria, $DF(x^*, y^*, z^*)$ is always a diagonal matrix and is explicitly given by

$$DF(x^*, y^*, z^*) = \begin{pmatrix} b_U(1 - 2x^*)(y^* + \varepsilon(1 - y^*)) & 0 & 0 \\ 0 & (1 - 2y^*)(-c_P + ux^*z^*) & 0 \\ 0 & 0 & (1 - 2z^*)(-c_R - x^*(1 - y^*)(v - b_{fo})) \end{pmatrix}.$$

$DF(x^*, y^*, z^*)$ still has three real eigenvalues, corresponding to the diagonal entries:

$$\lambda_1 = b_U(1 - 2x^*)(y^* + \varepsilon(1 - y^*)), \quad \lambda_2 = (1 - 2y^*)(-c_P + ux^*z^*), \quad \lambda_3 = (1 - 2z^*)(-c_R - x^*(1 - y^*)(v - b_{fo})).$$

Thus (x^*, y^*, z^*) is stable only if

$$\lambda_1 < 0, \quad \lambda_2 < 0, \quad \lambda_3 < 0.$$







(a) $x_0 = y_0 = z_0 = 0.1$, $\epsilon = 0.01$ (b) $x_0 = y_0 = z_0 = 0.5$, $\epsilon = 0.01$ (c) $x_0 = y_0 = z_0 = 0.1$, $\epsilon = -0.5$ (d) $x_0 = y_0 = z_0 = 0.5$, $\epsilon = -0.5$ (e) $x_0 = y_0 = z_0 = 0.1$, $\epsilon = -1$ (f) $x_0 = y_0 = z_0 = 0.5$, $\epsilon = -1$

Fig. 4. Time evolution of the density of trusting users $x(t)$, cooperative creators $y(t)$ and cooperative regulators $z(t)$ for the extended model with regulator' incentive (without conditional trust), as described in Equation (16). In all these simulations: $b_U = 4$, $c_P = 0.5$, $u = 1.5$, $c_R = 0.5$, $b_{fo} - v = 1.5$. For the initial conditions and the other parameters, see the captions.

If $v > b_{fo}$ then $-c_R - x^*(1 - y^*)(v - b_{fo}) < 0$, the stability for vertex equilibria is the same as in the case where $b_{fo} = 0$. We consider now $v < b_{fo}$:

When $\epsilon > 0$ then $\lambda_1 < 0$ iff $x^* = 1$, in which case λ_2 and λ_3 reduce to

$$\lambda_2 = (1 - 2y^*)(-c_P + uz^*), \lambda_3 = (1 - 2z^*)(-c_R - (1 - y^*)(v - b_{fo})).$$

- $(y^*, z^*) = (0, 0)$. Then $\lambda_2 = -c_P < 0$, $\lambda_3 = -c_R - (v - b_{fo})$. Thus the equilibrium $(1, 0, 0)$ is stable if $-c_R < v - b_{fo} < 0$ and unstable otherwise.
- $(y^*, z^*) = (0, 1)$. Then $\lambda_2 = -c_P + u$, $\lambda_3 = c_R + (v - b_{fo})$. Thus $(1, 0, 1)$ is stable only if $u < c_P$ and $c_R + (v - b_{fo}) < 0$ and is unstable otherwise.
- $(y^*, z^*) = (1, 0)$. Then $\lambda_2 = c_P > 0$. Thus the equilibrium $(1, 1, 0)$ is unstable.

- $(y^*, z^*) = (1, 1)$. Then $\lambda_2 = c_P - u$, $\lambda_3 = c_R > 0$. Thus $(1, 1, 1)$ is unstable.

When $\varepsilon < 0$

- $y^* = 0$. Then $\lambda_1 = b_U(1 - 2x^*)\varepsilon$. Hence $\lambda_1 < 0$ for $x^* = 0$. Then $\lambda_2 = -c_P < 0$ and $\lambda_3 = -c_R(1 - 2z^*)$. Hence $\lambda_3 < 0$ for $z^* = 0$. In this case, $(0, 0, 0)$ is a stable equilibrium.
- $y^* = 1$. Then $\lambda_1 = b_U(1 - 2x^*)$. Hence $\lambda_1 < 0$ for $x^* = 1$. Then $\lambda_2 = c_P - uz^*$. Thus $\lambda_2 < 0$ for $z^* = 0$ and $c_P < u$. Thus $\lambda_3 = -c_R < 0$. If $c_P < u$ then $(1, 1, 0)$ is a stable equilibrium.

Conclusion: When cooperative regulators are rewarded, a richer set of stable equilibrium points are possible, including $(1, 0, 1)$, but it is not sufficient to ensure the desirable one $(1, 1, 1)$.

3.3. Extended model with conditional trust

For the model with conditional trust, by calculating the differences in fitness through the payoffs in the matrix of Table 3, as we have done in the two previous models, we find that the replicator dynamics reads

$$\dot{x} = x(1 - x)b_U z(y + \varepsilon(1 - y)), \quad x(0) = x_0 \quad (18a)$$

$$\dot{y} = y(1 - y)[-c_P + uxz], \quad y(0) = y_0 \quad (18b)$$

$$\dot{z} = z(1 - z)[-c_R + b_R x + (b_{fo} - v)x(1 - y)], \quad z(0) = z_0 \quad (18c)$$

The stability analysis for this model is similar to the extended model, however the presence of a factor z in the evolution of x makes a continuum of fixed points emerge. Indeed, as well as the vertex equilibria

$$(x^*, y^*, z^*) \in \{0, 1\}^3 \quad (19)$$

there are

$$(x^*, y^*, z^*) \in [0, 1] \times \{0, 1\} \times \{0\} \quad (20)$$

The Jacobian reads

$$DF(x, y, z) = \begin{pmatrix} b_U z(1 - 2x)(y + \varepsilon(1 - y)) & b_U zx(1 - x)(1 - \varepsilon) & x(1 - x)b_U(y + \varepsilon(1 - y)) \\ uy(1 - y)z & (1 - 2y)(-c_P + uxz) & uy(1 - y)x \\ -z(1 - z)[b_R - (1 - y)(v - b_{fo})] & x(v - b_{fo})z(1 - z) & (1 - 2z)(-c_R + xb_R - x(1 - y)(v - b_{fo})) \end{pmatrix},$$

and it is still diagonal at the vertex equilibria $(x^*, y^*, z^*) \in \{0, 1\}^3$

$$DF(\{0, 1\}^3) = \begin{pmatrix} b_U z^*(1 - 2x^*)(y^* + \varepsilon(1 - y^*)) & 0 & 0 \\ 0 & (1 - 2y^*)(-c_P + ux^*z^*) & 0 \\ 0 & 0 & (1 - 2z^*)(-c_R + x^*b_R - x^*(1 - y^*)(v - b_{fo})) \end{pmatrix}.$$

Thus, one finds similar conditions for the stability of the vertex equilibria as the ones found for the previous model, with the remarkable exception that for $z^* = 0$ the Jacobian has a null eigenvalue. Specifically, with $z^* = 1$ the eigenvalues for the corresponding vertex equilibria read

$$\lambda_1 = b_U(1 - 2x^*)(y^* + \varepsilon(1 - y^*))$$

$$\lambda_2 = (1 - 2y^*)(-c_P + ux^*)$$

$$\lambda_3 = (-1)(-c_R + x^*b_R - x^*(1 - y^*)(v - b_{fo}))$$

thus one can conclude that

- $(1, 1, 1)$ is stable if $u > c_P$ and $b_R > c_R$
- $(0, 1, 1)$ is never stable
- $(1, 0, 1)$ is stable if $\varepsilon > 0$, $u < c_P$ and $b_R + b_{fo} > c_R + v$
- $(0, 0, 1)$ is never stable

Conclusion: A key general finding from our analysis of the three models is that, the desirable equilibrium point $(1, 1, 1)$ can be stable only in the model with conditional trust. Moreover, the cost of regulation needs to be sufficiently small ($c_R < b_R$) and the impact of punishment on unsafe creators need to be sufficiently high ($u > c_P$).




Fig. 5. Time evolution of the density of trusting users $x(t)$, cooperative creators $y(t)$ and cooperative regulators $z(t)$ for the extended model with conditional trust, as described in Equation (18). In all these simulations: $c_p = 0.5$, $u = 1.5$, $c_R = 0.5$, $b_{fo} - v = 5$. For the initial conditions and the other parameters, see the captions.

3.4. Numerical validation for replicator dynamics

We present here numerical results for the infinite population setting, validating the analytical observations shown above. First, Fig. 4 shows different outcomes in the evolutionary dynamics when varying initial conditions (i.e. initial fractions of x , y , z in the population), for the extended model where cooperative regulators are provided with a reward (without conditional trust). In general, as shown in the theoretical analysis, we observe that it is not possible to reach the stable equilibrium $(1, 1, 1)$ in this case, when users always trust and creators and regulators cooperate. This desirable equilibrium can become stable when users are enabled to make conditional decision based on regulators' cooperative behaviour, see Fig. 5 (panels d and f).

We consider evolutionary dynamics in more detail. In Fig. 4a), the system first reaches trust and regulation (TCC), which leads to profitability for regulators to stop enforcing safety standards (TCD). Interestingly, in panels 4b) and d) we observe a stable limit cycle between regulation and unsafe development. Users fully trust, but regulators mirror unsafe AI development. Safe AI reduces




Fig. 6. Low regulation cost ($c_R = 0.5$). Conditional trust can lead to full trust, cooperative regulation and safe development. Parameters set to: $b_U = b_R = b_P = 4$, $u = 1.5$, $v = 0.5$, $c_P = 0.5$, $\beta = 0.1$, $N_U = N_P = N_R = 100$.

regulation, whereas unsafe AI prompts them to reinstate incentives. Thus, we could simplify the equations (16), taking $x(t) = 1$ constantly and explaining mathematically the existence of the cycle. Qualitatively, when defecting creators emerge, regulators enforce standards (because $b_{fo} > v + c_R$), thus it is profitable to comply when meeting a defecting creator. Conversely, when cooperating regulators emerge, for creators it is better to cooperate as well (as $u > c_P$); but, if creators cooperate, the enforcing standards becomes unprofitable and regulators defect. In the absence of regulation, unsafe creators re-emerge and the limit cycle restarts. In 4c), e) and f), trust is never established in the users' population, there is no incentive to regulate as there is no adoption, and creators ignore safety standards - NDD is reached.

Fig. 5 shows that $(1, 1, 1)$ can be a stable equilibrium for the condition theoretically derived above (i.e. $u > c_P$ and $b_R > c_R$). But it only occurs in non-favourable initial conditions (see left column).

4. Stochastic evolutionary dynamics

We now study evolutionary game dynamics in finite populations (see Methods, Section 2.2.1). Compared to traditional concepts of evolutionary stability and dynamics of infinite populations, stochastic effects in finite population dynamics, including errors in social learning, can have dramatic effects on evolutionary outcomes [44–46]. Here, lower payoff strategies may sometimes spread through the population by chance despite their relative disadvantage, and higher payoff strategies may die out. This stochastic approach has been shown to be powerful in explaining empirical observations in human behavioural experiments [45,46].

Fig. 6 shows the frequencies of the eight strategies when varying the reward for cooperative regulators, b_{fo} , for both proposed extended models when the cost of regulation, c_R , is sufficiently small. We identify two regimes, based on the sign of the risk factor ε (see again Table 1). We assume that defection in the creator population (i.e. strategies of the form xDz) is the most deleterious






Fig. 7. High regulation cost ($c_R = 5$). Conditional trust is not sufficient to enable safe development and AI adoption when the cost of regulation is high, as this discourages regulators to cooperate. Parameters set to: $b_U = b_R = b_P = 4$, $u = 1.5$, $v = 0.5$, $c_P = 0.5$, $\beta = 0.1$, $N_U = N_P = N_R = 100$.

outcome for users, but there still exists a positive, partial benefit when the sign of ϵ is positive. In such cases, we observe the evolution of trust regardless of the behaviour of regulators. Given low incentives for capturing creators of unsafe AI, regulators fail to adequately enforce safety. Increasing incentives compels regulators to act, thus giving rise to a mixed outcome between TCC , TCD and TDC . For a clear view of the dominance cycles that govern these states, in Fig. 8, we show the fixation probabilities and stationary distributions. They were computed using the Methods in Section 2.2.1. Whether the population evolves to TDD (for low b_{f0}) or the mixed state, we note that for this regime, it is always beneficial to adopt even unsafe AI (see discussion). If $\epsilon < 0$, no trust can evolve, and AI is never adopted.

Importantly, when regulation is cheap to implement (low c_R), conditional trust solves these issues for all incentives and in both risk regimes.

Fig. 7 shows results for varying b_{f0} for different models, when the cost of regulation c_R is large compared to the benefit, e.g. when AI is highly sophisticated and it is difficult to develop technology to capture unsafe behaviour. In the regime of high risk $\epsilon < 0$, we observe defection because users do not trust the technology and the regulators. For very low values of b_{f0} , there is no incentive to deviate from defection, and this remains the main regime even for increasing b_{f0} values. With conditional trust, illustrated in Fig. 6 (b), even in the absence of trust in regulations, other strategies start emerging, due to cyclic dominance between $TCC \rightarrow TCD \rightarrow TDD \rightarrow TDC$ (see Fig. 8). This cycle only emerges if regulators are given a high enough incentive to enforce safety standards (b_{f0}). With a high risk factor $\epsilon < 0$, trust, regulation and safety never emerge. Interestingly, extending this model to conditional trust, we see a decrease in trust and adoption. Previously, low risk ($\epsilon > 0$) provided enough benefit for users to adopt AI even in the case of rapid, risky development, yet regulators prevent this from occurring. Regulation is too costly, thus regulators restrict access to AI, deeming it unsafe. If we consider conditional trust in such cases, users cannot elect to adopt technology unless regulators approve of creators, which hinges on the cost of regulation.

(I) Without CT

(A) $b_{f0} = 0, \varepsilon = 0.2, c_R = 0.5$ (B) $b_{f0}=5, \varepsilon = 0.2, c_R = 0.5$

(II) With CT



(C) $b_{f0}=5, \varepsilon = -0.1, c_R = 5$ (D) $b_{f0}=5, \varepsilon = -0.1, c_R = 0.5$

Fig. 8. Stationary distribution and transitions among the strategies (with conditional trust and without conditional trust). Only arrows where the transition probability is larger than the opposite one, are shown. Dashed lines represent neutral transitions. Parameters set to: $b_U = b_P = b_R = 4$, $u = 1.5$, $v = 0.5$, $c_P = 0.5$, $\beta = 0.1$. Population sizes: $N_U = N_P = N_R = 100$.

Overall, similarly to the replicator dynamics analysis, we observe that the desirable outcome of a high frequency of TCC (i.e. the equilibrium point $(1, 1, 1)$ is stable) can be achieved when conditional trust is enabled, the cost of regulation is sufficiently small ($c_R < b_R$) and the impact of punishment on unsafe creators is sufficiently high.

5. Discussion

5.1. Key takeaways

Our conceptual model of an AI governance ecosystem highlights several key factors relevant to both current and future regulatory landscapes for AI:

- I. AI companies are motivated to implement safety protocols only when regulatory bodies have the capacity to enforce compliance;
- II. In scenarios where the potential negative consequences of unsafe AI development are significant, user trust in the system hinges on the dependability of regulatory authorities;
- III. For regulatory bodies to consistently enforce safety measures, the cost associated with regulation must be manageable, and their capacity to identify non-compliant AI creators should be high;
- IV. Discerning users, namely those whose trust in the regulatory system is conditional on its perceived effectiveness, are essential to enable a stable and long-lasting trust in the system, where creators comply with regulations and regulators continue to enforce them.

We found that trust can only emerge when regulators have an incentive to regulate. Moreover, the presence of discerning users, that is, users that place their trust conditionally, is able to break the cyclic dynamics that we observe in Fig. 4a. Figs. 8b and d show how the presence of these users who trust conditionally breaks the cycle in the dimension of the regulator and creator populations, by making TCC dominant over TCD. That is, when users trust conditionally, enforcing regulation risk-dominates not enforcing it.

5.2. Implications for AI governance

We focus on the policy implications of our main findings from the analyses in Sections 3 and 4: high costs in providing effective regulation pose an obstacle to building trust in future AI systems, and that it is vital to help users discern which systems have undergone effective regulation.

Scenarios where regulators face large costs in assessing the risks of AI systems are plausible. Existing AI evaluators already experience relatively high costs in providing effective regulation for frontier AI systems. Foundation models like Open AI's GPT-4 require large amounts of computational hardware and access to the top AI talent [29]. Evaluations of these systems are likewise expensive and demand teams with expertise in AI. Academia and public institutions who wish to assess the risks of new AI systems therefore compete with AI creators in acquiring these resources (although see recent progress in government-led evaluations of new language models [47,48]).

To overcome these barriers to entry, policy makers should consider (i) investing further in government capacity to assess the risks of large AI systems [49], and (ii) subsidising organisations that show promise in providing services that improve the effectiveness of AI regulation [50]. Governments may therefore wish to take heed of the UK's investment in their AI Safety Institute [51].

Governments and regulators can benefit from communicating information that allows users to discern which methods that regulators use are reliable (so that, in turn, they can discern which AI systems are safe and trustworthy). Doing so requires investing in the expertise needed to follow the recommendations on how to assess the risks presented by AI systems [12,52]. This could mean conducting machine learning model evaluations or other tests as part of an audit or risk assessment of new frontier AI systems [53,54]. The role of academia and the commentariat should also not be underestimated here [7]. Academia and commentators can play a key role in communicating to users information about the effectiveness of regulations. Our model can be extended to capture this by including academics and commentators as an additional evolving population of agents.

In the case where the regulator is a public institution with greater powers, they may wish to monitor the size of training runs for frontier AI systems or to impose fines on companies whose AI systems cause significant economic harm [55]. The governance of computing hardware (shortened to compute) has also been noted as a point of potentially high leverage, given that large amounts of physical infrastructure are much easier to monitor, track, and limit in principle than software, data, or model weights [56].

One more avenue regulators can pursue to help demonstrate to users that trustworthy AI systems will be developed and deployed safely is to encourage international coordination in the governance of AI systems, perhaps by proposing designs for international institutions to govern new AI breakthroughs or AI safety projects [57–60]. World powers have expressed an awareness that AI could pose enormous risks when signing the Bletchley Declaration [61], highlighting that international coordination of some form is tractable and desired.

Finally, public institutions will be necessary to ensure that companies and users, particularly large enterprises that serve a wide number of additional users, can form reliable impressions of the reputation of different regulators [51]. Model registration is an essential prerequisite for such a reputation system; Users need to know which regulatory evaluations match which AI systems [62]. Moreover, governments may wish those providing regulatory services to preregister a set of experiments they plan to run to test the model. Regulators could later add details about how they dealt with unforeseen findings in their experiments. This would help them to demonstrate vigilance (i.e. that they have a high likelihood of finding early signs of dangerous capabilities when they in fact exist), while ensuring that their insights can be shared with other organisations quickly and safely [49,63]. Keep in mind that, as our results show, this information is useful whenever we have multiple organisations evaluating frontier AI systems, including at the international level [58].

5.3. Limitations and areas for future research

Above, we discussed how decision makers can act on the insights from our results. Now, we turn to the limitations of the present model. An important simplification we made was to randomly match users, creators and regulators. Future research could address these limitations by incorporating partner selection between populations into these models. In a more realistic setting, users would choose which AI system they want to use. Companies too may choose to relocate to avoid especially burdensome regulation.

Another important area is to model more explicitly competition between different regulatory blocs. One appropriate formalism for achieving this would be to model each regulatory bloc as an island network that is connected loosely to other islands. These islands of regulation would be in competition with each other through a process of cultural group selection [64–66]. For example, creators may move to regulatory regimes more favourable to themselves, while regulators may imitate the regulatory regimes of other groups that are more successful than themselves.

Moreover, it might be valuable to move beyond the assumption of uniform actors, studying the impact of heterogeneity and influence that are prevalent in an AI ecosystem. For instance, some users, AI creators, or regulators may possess disproportionate influence due to their market share, computing capacities, or political power [3,67]. Therefore, models that incorporate non-uniform interactions, potentially drawing upon techniques from contagion or network spreading [68], could provide an alternative representation of the AI governance ecosystem dynamics. Such models could capture how the behaviours and decisions of influential actors cascade

through the system. This would allow us to examine the effectiveness of regulatory interventions targeting specific influential actors [69].

Our model considers reward as a form of incentive for encouraging regulators to capture unsafe AI creators. Future research could also benefit from drawing on insights from efficient incentive mechanisms in other settings, such as social dilemmas, bargaining games and climate change mitigation [70–73]. For instance, recent evolutionary game models of institutional incentives have shown that hybrid and adaptive combinations of different forms of positive and negative incentives can lead to cost-efficient approaches for promoting various prosocial behaviours such as cooperation, fairness and trust [69,74–78]. Furthermore, it would be beneficial to consider providing incentives for other actors within the model. For example, users can be encouraged to actively participate by reporting instances of misuse or harm caused by AI systems, such as through media channels [79].

Our model also does not fully capture the race dynamics between AI companies. OpenAI's initial commercial breakthrough with ChatGPT has sparked interest in building more capable foundation models among the world's leading tech companies [29]. As predicted by other game-theoretic models of tech race dynamics, remaining at the frontier of AI capabilities appears to take precedence over pursuing AI safety research agendas [17,20]. It is noteworthy that to date, companies have had relatively limited success making generative AI systems fundamentally safer.

Scaling AI systems quickly encounters a brisk trade-off with safety research. It allows less time to investigate and understand the impacts of smaller models, particularly when new algorithms are used to pre-train or fine tune those models. More talent and computational resources must also be allocated to these large experiments, allowing fewer resources to pursue research agendas on AI safety. Therefore, we cannot expect even relatively safety-minded AI labs to choose the optimal trade-off, given the huge competitive dynamics at play. For a further discussion of what can be done to directly curb race dynamics, consult the following references [28,55,63,80].

While we recognise the limitations of our simpler scenario, we find our approach useful as a tool for thinking through what assumptions policymakers must make to assume that different regulatory regimes are likely to achieve their policy objectives.

Further avenues of research should test the predictions of our model, providing real-world evidence of strategic interactions between users, creators, and regulators in the AI domain. Exploring additional incentive structures for regulators could further enhance our understanding of effective AI governance. A useful incentive structure to consider is that AI companies could be mandated to pay for regulatory services that governments consider essential to demonstrate the safety of the model [10,13].

Our findings offer valuable insights into the dynamics at play in the AI regulation landscape and highlight the indispensable role of reliable regulators, incentivised by either governmental rewards or the maintenance of a prestigious reputation. They highlight the value of taking a game-theoretic approach to formally model the effects of different kinds of regulatory system on the behaviour of both AI creators and users.

CRediT authorship contribution statement

Zainab Alalawi: Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **Paolo Bova:** Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **Theodor Cimpanu:** Investigation, Methodology, Validation, Visualization, Writing – original draft, Writing – review & editing. **Alessandro Di Stefano:** Investigation, Methodology, Supervision, Validation, Writing – original draft, Writing – review & editing. **Manh Hong Duong:** Formal analysis, Investigation, Methodology, Resources, Supervision, Validation, Writing – original draft, Writing – review & editing. **Elias Fernández Domingos:** Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **The Anh Han:** Conceptualization, Formal analysis, Investigation, Methodology, Project administration, Resources, Supervision, Validation, Writing – original draft, Writing – review & editing. **Marcus Krellner:** Investigation, Methodology, Validation, Visualization, Writing – original draft, Writing – review & editing. **Ndidi Bianca Ogbo:** Investigation, Methodology, Validation, Visualization, Writing – original draft, Writing – review & editing. **Simon T. Powers:** Conceptualization, Investigation, Methodology, Supervision, Validation, Visualization, Writing – original draft, Writing – review & editing. **Filippo Zimmaro:** Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing.

Acknowledgement

An earlier version of this work was produced during the workshop “AI Governance Modelling”, funded through the generous support from the Future of Life Institute (TAH). T.A.H. is supported by EPSRC (grant EP/Y00857X/1). M.H.D is supported by EPSRC (grant EP/Y008561/1) and a Royal International Exchange Grant IES-R3-223047.

Data availability

No data was used for the research described in the article.

References

- [1] Y. Bengio, G. Hinton, A. Yao, D. Song, P. Abbeel, et al., Managing extreme AI risks amid rapid progress, *Science* 384 (6698) (May 2024) 842–845 [Online]. Available: <https://www.science.org/doi/abs/10.1126/science.adn0117>.

- [2] T. Baker, The Executive Order on Safe, Secure, and Trustworthy AI: Decoding Biden's AI Policy Roadmap, Nov. 2023.
- [3] Y. Bengio, S. Mindermann, D. Privitera, T. Besiroglu, R. Bommasani, S. Casper, Y. Choi, P. Fox, B. Garfinkel, D. Goldfarb, et al., International ai safety report, arXiv preprint, arXiv:2501.17805, 2025.
- [4] G. Finocchiaro, The regulation of artificial intelligence, *AI Soc.* 39 (4) (2024) 1961–1968.
- [5] L. Floridi, J. Cowsls, M. Beltrametti, R. Chatila, P. Chazerand, V. Dignum, C. Luetge, R. Madelin, U. Pagallo, F. Rossi, et al., Ai4people—an ethical framework for a good ai society: opportunities, risks, principles, and recommendations, *Minds Mach.* 28 (2018) 689–707.
- [6] J. Laux, S. Wachter, et al., Trustworthy artificial intelligence and the European Union AI act: on the conflation of trustworthiness and acceptability of risk, *Regul. Gov.* 18 (1) (2024) 3–32.
- [7] S.T. Powers, O. Linnyk, et al., The stuff we swim in: regulation alone will not lead to justifiable trust in AI, *IEEE Technol. Soc. Mag.* 42 (4) (2023) 95–106.
- [8] C. Siegmann, M. Anderljung, The Brussels Effect and Artificial Intelligence, Oct. 2022.
- [9] J. Tallberg, E. Erman, et al., The global governance of artificial intelligence: next steps for empirical and normative research, *Int. Stud. Rev.* 25 (3) (2023) viad040, private vs public regulation.
- [10] J. Clark, G.K. Hadfield, Regulatory Markets for AI Safety, arXiv, Dec. 2019.
- [11] M. Anderljung, J. Barnhart, et al., Frontier AI Regulation: Managing Emerging Risks to Public Safety, Jul. 2023.
- [12] A. Dafoe, AI governance: overview and theoretical lenses, in: J.B. Bullock, Y.-C. Chen, J. Himmelreich, V.M. Hudson, A. Korinek, M.M. Young, B. Zhang (Eds.), *The Oxford Handbook of AI Governance*, Oxford University Press, 2023, pp. 21–44.
- [13] G.K. Hadfield, J. Clark, Regulatory Markets: the Future of AI Governance, Apr. 2023.
- [14] J. Alaga, J. Schuett, Coordinated Pausing: an Evaluation-Based Coordination Scheme for Frontier AI Developers, Sep. 2023.
- [15] E. Zaidan, I.A. Ibrahim, Ai governance in a complex and rapidly changing regulatory landscape: a global perspective, *Humanit. Soc. Sci. Commun.* 11 (1) (2024) 1–18.
- [16] J. Pitt, Chatsh*t and other conversations (that we should be having, but mostly are not), *IEEE Technol. Soc. Mag.* 42 (3) (2023) 7–13.
- [17] S. Armstrong, N. Bostrom, et al., Racing to the precipice: a model of artificial intelligence development, *AI Soc.* 31 (2) (May 2016) 201–206.
- [18] T. LaCroix, A. Mohseni, The tragedy of the AI commons, *Synthese* 200 (4) (2022) 289.
- [19] M. Jensen, N. Emery-Xu, et al., Industrial Policy for Advanced AI: Compute Pricing and the Safety Tax, 2023.
- [20] T.A. Han, L.M. Pereira, et al., To regulate or not: a social dynamics analysis of an idealised AI race, *J. Artif. Intell. Res.* 69 (Nov. 2020) 881–921.
- [21] T.A. Han, Institutional incentives for the evolution of committed cooperation: ensuring participation is as important as enhancing compliance, *J. R. Soc. Interface* 19 (188) (2022) 20220036.
- [22] T. Cîmpeanu, F. Santos, et al., Artificial intelligence development races in heterogeneous settings, *Sci. Rep.* 12 (1) (2022) 1723.
- [23] P. Bova, A. Di Stefano, et al., A tale of two regulatory markets: the role of institutional incentives in supporting sustainable regulatory markets for future AI systems, in: *ALIFE 2023: Ghost in the Machine: Proceedings of the 2023 Artificial Life Conference*, MIT Press, Jul. 2023.
- [24] K.J.D. Chan, G. Papyshev, M. Yarime, Balancing the tradeoff between regulation and innovation for artificial intelligence: an analysis of top-down command and control and bottom-up self-regulatory approaches, *Technol. Soc.* 79 (2024) 102747.
- [25] T.A. Han, L.M. Pereira, et al., Mediating artificial intelligence developments through negative and positive incentives, *PLoS ONE* 16 (1) (Jan. 2021).
- [26] J. Hofbauer, K. Sigmund, *Evolutionary Games and Population Dynamics*, Cambridge University Press, 1998.
- [27] T.A. Han, C. Perret, S.T. Powers, When to (or not to) trust intelligent machines: insights from an evolutionary game theory analysis of trust in repeated games, *Cogn. Syst. Res.* 68 (2021) 111–124.
- [28] A. Askell, M. Brundage, et al., The Role of Cooperation in Responsible AI Development, arXiv, Jul. 2019.
- [29] B. Cottier, T. Besiroglu, et al., Who Is Leading in AI? An Analysis of Industry AI Research, 2024.
- [30] D. Tom, G. Charley, How Microsoft catapulted to \$3 trillion on the back of AI, <https://www.wsj.com/tech/microsoft-closes-above-3-trillion-as-its-big-ai-play-generates-excitement-e37ea2a4>, 2024.
- [31] P. Cihon, M.J. Kleinaltenkamp, et al., AI certification: advancing ethical practice by reducing information asymmetries, *IEEE Trans. Technol. Soc.* 2 (4) (Dec. 2021) 200–209.
- [32] D.C. North, *Institutions, Institutional Change and Economic Performance*, Cambridge University Press, 1990.
- [33] K. Sigmund, The calculus of selfishness, in: *The Calculus of Selfishness*, Princeton University Press, 2010.
- [34] L.E. Blume, The statistical mechanics of strategic interaction, *Games Econ. Behav.* 5 (3) (1993) 387–424.
- [35] G. Szabó, C. Töke, Evolutionary prisoner's dilemma game on a square lattice, *Phys. Rev. E* 58 (1) (1998) 69.
- [36] L.A. Imhof, D. Fudenberg, M.A. Nowak, Evolutionary cycles of cooperation and defection, *Proc. Natl. Acad. Sci. USA* 102 (2005) 10 797–10 800.
- [37] M.A. Nowak, A. Sasaki, C. Taylor, D. Fudenberg, Emergence of cooperation and evolutionary stability in finite populations, *Nature* 428 (2004) 646–650.
- [38] E.F. Domingos, F.C. Santos, T. Lenaerts, Egttools: evolutionary game dynamics in python, *iScience* 26 (4) (2023).
- [39] A. Traulsen, M.A. Nowak, J.M. Pacheco, Stochastic dynamics of invasion and fixation, *Phys. Rev. E* 74 (2006) 11909.
- [40] S. Encarnação, F.P. Santos, F.C. Santos, V. Blass, J.M. Pacheco, J. Portugali, Paradigm shifts and the interplay between state, business and civil sectors, *R. Soc. Open Sci.* 3 (12) (2016) 160753.
- [41] Z. Alalawi, T.A. Han, Y. Zeng, A. Elragig, Pathways to good healthcare services and patient satisfaction: an evolutionary game theoretical approach, in: *Artificial Life Conference Proceedings*, MIT Press One Rogers Street, Cambridge, MA 02142-1209, USA, 2019, pp. 135–142, journals-info
- [42] P.D. Taylor, Evolutionarily stable strategies with two types of player, *J. Appl. Probab.* 16 (1) (1979) 76–83.
- [43] J. Bauer, M. Broom, E. Alonso, The stabilization of equilibria in evolutionary game dynamics through mutation: mutation limits in evolutionary games, *Proc. Royal Soc. A* 475 (2231) (2019) 20190355.
- [44] M.A. Nowak, A. Sasaki, C. Taylor, D. Fudenberg, Emergence of cooperation and evolutionary stability in finite populations, *Nature* 428 (6983) (2004) 646–650.
- [45] D.G. Rand, C.E. Tarnita, H. Ohtsuki, M.A. Nowak, Evolution of fairness in the one-shot anonymous ultimatum game, *Proc. Natl. Acad. Sci. USA* 110 (7) (2013) 2581–2586.
- [46] I. Zisis, S. Di Guida, T.A. Han, G. Kirchsteiger, T. Lenaerts, Generosity motivated by acceptance-evolutionary analysis of an anticipation game, *Sci. Rep.* 5 (1) (2015) 18076.
- [47] UK AI Safety Institute, Pre-deployment evaluation of anthropic's upgraded Claude 3.5 Sonnet [Online]. Available: <https://www.aisi.gov.uk/work/pre-deployment-evaluation-of-anthropics-upgraded-claude-3-5-sonnet>, 2024.
- [48] UK AI Safety Institute, Pre-deployment evaluation of open AI's o1 model [Online]. Available: <https://www.aisi.gov.uk/work/pre-deployment-evaluation-of-openais-o1-model>, 2024.
- [49] J. Whittlestone, J. Clark, Why and How Governments Should Monitor AI Development, arXiv, Aug. 2021.
- [50] M. Kinniment, L.J.K. Sato, H. Du, B. Goodrich, M. Hasin, et al., Evaluating language-model agents on realistic autonomous tasks [Online]. Available: <http://arxiv.org/abs/2312.11671>, Jan. 2024.
- [51] GOV.UK, Introducing the AI safety institute, <https://www.gov.uk/government/publications/ai-safety-institute-overview/introducing-the-ai-safety-institute>, 2023.
- [52] UK AI Safety Institute, Early lessons from evaluating frontier ai systems [Online]. Available: <https://www.aisi.gov.uk/work/early-lessons-from-evaluating-frontier-ai-systems>, 2024.
- [53] T. Shevlane, S. Farquhar, et al., Model Evaluation for Extreme Risks, May 2023.

- [54] L. Koessler, J. Schuett, Risk Assessment at AGI Companies: a Review of Popular Risk Assessment Techniques from Other Safety-Critical Industries, Jul. 2023.
- [55] Y. Shavit, What Does It Take to Catch a Chinchilla? Verifying Rules on Large-Scale Neural Network Training via Compute Monitoring, Mar. 2023.
- [56] G. Sastry, L. Heim, H. Belfield, M. Anderljung, M. Brundage, et al., Computing Power and the Governance of Artificial Intelligence, 2024.
- [57] R.F. Trager, B. Harack, et al., International governance of civilian AI: a jurisdictional certification approach [Online]. Available: <https://dx.doi.org/10.2139/ssrn.4579899>, 2023.
- [58] L. Ho, J. Barnhart, et al., International Institutions for Advanced AI, Jul. 2023.
- [59] R. Gruetzemacher, A. Chan, et al., An international consortium for AI risk evaluations, in: *Socially Responsible Language Modelling Research*, Nov. 2023.
- [60] L. Hammond, A. Chan, J. Clifton, J. Hoelscher-Obermaier, A. Khan, E. McLean, C. Smith, W. Barfuss, J. Foerster, T. Gavenčiak, T.A. Han, E. Hughes, V. Kovafik, J. Kulveit, J.Z. Leibo, C. Oesterheld, C.S. de Witt, N. Shah, M. Wellman, P. Bova, T. Cîmpeanu, C. Ezell, Q. Feuillade-Montixi, M. Franklin, E. Kran, I. Krawczuk, M. Lamparth, N. Lauffer, A. Meinke, S. Motwani, A. Reuel, V. Conitzer, M. Dennis, I. Gabriel, A. Gleave, G. Hadfield, N. Haghtalab, A. Kasirzadeh, S. Krier, K. Larson, J. Lehman, D.C. Parkes, G. Piliouras, I. Rahwan, Multi-agent risks from advanced ai [Online]. Available: <https://arxiv.org/abs/2502.14143>, 2025.
- [61] Gov uk, The bletchley declaration by countries attending the AI safety summit, <https://www.gov.uk/government/publications/ai-safety-summit-2023-the-bletchley-declaration/the-bletchley-declaration-by-countries-attending-the-ai-safety-summit-1-2-november-2023>, 2023.
- [62] G. Hadfield, M.-F. Cuéllar, T. O'Reilly, It's time to create a national registry for large AI models, <https://carnegieendowment.org/2023/07/12/it-s-time-to-create-national-registry-for-large-ai-models-pub-90180>, 2023.
- [63] P. Bova, A. Di Stefano, et al., Both eyes open: vigilant incentives help auditors improve ai safety, *J. Phys. Complex.* 5 (2) (2024) 025009.
- [64] K. Binmore, *Natural Justice*, Oxford University Press, 2005.
- [65] J.C. van den Bergh, J.M. Gowdy, A group selection perspective on economic behavior, institutions and organizations, *J. Econ. Behav. Organ.* 72 (1) (2009) 1–20.
- [66] P. Richerson, R. Baldini, A.V. Bell, K. Demps, K. Frost, V. Hillis, S. Mathew, E.K. Newton, N. Naar, L. Newson, et al., Cultural group selection plays an essential role in explaining human cooperation: a sketch of the evidence, *Behav. Brain Sci.* 39 (2016) e30.
- [67] P. Verdegem, Dismantling ai capitalism: the commons as an alternative to the power concentration of big tech, *AI Soc.* 39 (2) (2024) 727–737.
- [68] D. Brockmann, D. Helbing, The hidden geometry of complex, network-driven contagion phenomena, *Science* 342 (6164) (2013) 1337–1342.
- [69] T. Cîmpeanu, C. Perret, T.A. Han, Cost-efficient interventions for promoting fairness in the ultimatum game, *Knowl.-Based Syst.* 233 (2021) 107545 [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0950705121008078>.
- [70] D. Balliet, L.B. Mulder, P.A. Van Lange, Reward, punishment, and cooperation: a meta-analysis, *Psychol. Bull.* 137 (4) (2011) 594.
- [71] K. Sigmund, C. Hauert, M.A. Nowak, Reward and punishment, *Proc. Natl. Acad. Sci. USA* 98 (19) (2001) 10 757–10 762.
- [72] A.R. Góis, F.P. Santos, J.M. Pacheco, F.C. Santos, Reward and punishment in climate change dilemmas, *Sci. Rep.* 9 (1) (2019) 1–9.
- [73] T.A. Han, Institutional incentives for the evolution of committed cooperation: ensuring participation is as important as enhancing compliance, *J. R. Soc. Interface* 19 (188) (2022) 20220036.
- [74] W. Sun, L. Liu, X. Chen, A. Szolnoki, V.V. Vasconcelos, Combination of institutional incentives for cooperative governance of risky commons, *iScience* 24 (8) (2021).
- [75] M.H. Duong, T.A. Han, Cost efficiency of institutional incentives for promoting cooperation in finite populations, *Proc. Royal Soc. A* 477 (2254) (2021) 20210568.
- [76] S. Wang, X. Chen, Z. Xiao, A. Szolnoki, V.V. Vasconcelos, Optimization of institutional incentives for cooperation in structured populations, *J. R. Soc. Interface* 20 (199) (2023) 20220653.
- [77] Y. Liu, L. Wang, R. Guo, S. Hua, L. Liu, L. Zhang, T.A. Han, Evolution of trust in the n-player trust game with transformation incentive mechanism, *J. R. Soc. Interface* 22 (224) (2025) 20240726.
- [78] T.A. Han, M.H. Duong, M. Perc, Evolutionary mechanisms that promote cooperation may not promote social welfare, *J. R. Soc. Interface* 21 (220) (2024) 20240547.
- [79] N. Balabanova, A. Bashir, P. Bova, A. Buscemi, T. Cîmpeanu, H.C. da Fonseca, A. Di Stefano, M.H. Duong, E.F. Domingos, A. Fernandes, et al., Media and responsible ai governance: a game-theoretic and llm analysis, *arXiv preprint*, arXiv:2503.09858, 2025.
- [80] T.A. Han, T. Lenaerts, F.C. Santos, L.M. Pereira, Voluntary safety commitments provide an escape from over-regulation in AI development, *Technol. Soc.* 68 (2022) 101843.
- [81] Donghee Shin, Yong Jin Park, Role of fairness, accountability, and transparency in algorithmic affordance, *Comput. Hum. Behav.* (ISSN 0747-5632) 98 (2019) 277–284, <https://doi.org/10.1016/j.chb.2019.04.019>, <https://www.sciencedirect.com/science/article/pii/S0747563219301591>.