

# THE TRAGEDY OF THE AI COMMONS

TRAVIS LACROIX<sup>1</sup>

AYDIN MOHSENI<sup>2</sup>

**ABSTRACT.** Policy and guideline proposals for ethical artificial-intelligence research have proliferated in recent years. These are supposed to guide the socially-responsible development of AI for the common good. However, there typically exist incentives for non-cooperation (i.e., non-adherence to such policies and guidelines); and, these proposals often lack effective mechanisms to enforce their own normative claims. The situation just described constitutes a social dilemma—namely, a situation where no one has an individual incentive to cooperate, though mutual cooperation would lead to the best outcome for all involved. In this paper, we use stochastic evolutionary game dynamics to model this social dilemma in the context of the ethical development of artificial intelligence. This formalism allows us to isolate variables that may be intervened upon, thus providing actionable suggestions for increased cooperation amongst numerous stakeholders in AI. Our results show how stochastic effects can help make cooperation viable in such a scenario. They suggest that coordination for a common good should be attempted in smaller groups in which the cost for cooperation is low, and the perceived risk of failure is high. This provides insight into the conditions under which we should expect such ethics proposals to be successful with regard to their scope, scale, and content.

**Keywords** – Stochastic Game Theory, Artificial Intelligence, Social Dilemmas, AI Ethics, Cooperation, Ethics Guidelines, Soft Law, Safe and Responsible AI Research

## 1. INTRODUCTION

Artificial intelligence promises to fundamentally change nearly every facet of our lives, for better or worse (Harari, 2017; Helbing, 2019; Makridakis, 2017). In response to this reality, there has been a proliferation of policy and guideline proposals for ethical artificial-intelligence and machine-learning (AI/ML) research. Jobin et al. (2019) survey several global initiatives for AI/ML and find no fewer than 84 documents containing ethics principles for AI research, with 88% of these having been released since 2016. More broadly, the World Economic Forum has identified almost three hundred separate efforts to develop ethical principles for AI (Russell, 2019). Policy documents of this sort are meant to specify ‘best practices’ to which engineers, developers, researchers, etc. ought to adhere.

---

<sup>1</sup>DEPARTMENT OF PHILOSOPHY, DALHOUSIE UNIVERSITY

<sup>2</sup>CENTER FOR PHILOSOPHY OF SCIENCE, UNIVERSITY OF PITTSBURGH

*E-mail address:* tlacroix@dal.ca (Corresponding); aydin.mohseni@gmail.com.

*Date:* Draft of May 23, 2022.

For example, the *Montréal Declaration for the Responsible Development of AI* identifies a set of abstract normative principles and values intended to promote the fundamental interests of stakeholders; signatories are invited to commit to ‘the development of AI at the service of the individual and the common good’ (Université de Montréal, 2017). Proposals of this sort typically highlight issues concerning transparency, justice and fairness, non-maleficence, responsibility, and privacy, among others (Jobin et al., 2019). These initiatives generally take one of two approaches to foster the ethical practice of AI research: proposing principles to guide the socially-responsible development of AI or examining the societal impacts of AI (Luccioni and Bengio, 2019).

However, policies-and-procedures documents, like the *Montréal Declaration*, are examples of ‘non-legislative policy instruments’ or ‘soft law’ (Sossin and Smith, 2003), which are instruments for cooperation that are not legally binding. This stands in contrast to ‘hard law’, which consists in legally-binding regulations passed by legislatures. By definition, then, these reports are not intended to produce enforceable rules but are meant merely as *guides* for ethical practice. Therefore, the *proliferation* of such guiding principles raises pressing questions about their efficacy. It has been suggested that such declarations have little to no practical effect (Hagendorff, 2019); at least one study finds that the effectiveness of ethical codes for influencing the behaviour of professionals in the technology community is virtually nonexistent (McNamara et al., 2018).

Part of the problem is that, in the AI research community, as in many other social contexts, there are typically costs associated with cooperation—e.g., the additional time and financial resources that are required to ensure research and business practices adhere to the proposed normative standards—in addition to incentives for non-cooperation or non-adherence to ethics agreements—e.g., a lucrative defence contract to develop autonomous weapons. Furthermore, non-legislative policy instruments, like guiding principles, lack effective mechanisms to *enforce* (or *reinforce*) their own normative claims, because they depend on voluntary, non-binding cooperation between the relevant parties—e.g., individual researchers, labs, academic bodies, corporations, governments, international bodies, etc.

This failure should come as no surprise to game theorists and economists: in the situation just described, no one has an individual incentive to cooperate, although mutual cooperation would lead to the best outcome for all those involved. This constitutes a *social dilemma* (Allison and Kerr, 1994; Dawes, 1980; Hobbes, 1651; Hume, 1739; Ross, 2019; Serrano and Feldman, 2013).

In this paper, we use stochastic evolutionary game dynamics to model this social dilemma in the context of the ethical development of AI. This model allows us to isolate variables that may be intervened upon; this, in turn, helps to formulate

actionable suggestions for increased cooperation amongst numerous stakeholders in AI. Our results show how stochastic effects can help make cooperation viable in such a scenario, and they suggest, among other things, that coordination for a common good should be attempted in smaller groups in which the cost for cooperation is low, and the perceived risk of failure is high. This analysis provides insight into the conditions under which we may expect such policy and guideline proposals for safe and ethical AI to be successful in terms of their scope, scale, and content. This, in turn, yields plausible solutions to the social dilemmas that the AI/ML community face for acting in a socially responsible and ethical way that is aligned with a common good.<sup>1</sup>

In Section 2, we discuss related work. In Section 3, we present an evolutionary game which models cooperation (with respect to proposed ethical norms and principles in AI) using a *social dynamics* framework. Select results are presented in Section 4. Section 5 clarifies and discusses a number of morals we derive from our model. Here, we outline a concrete situation under which these morals may be applied and discuss the practical applications that the insights from our model proffer. Section 6 concludes.

## 2. BACKGROUND AND RELATED WORK

In this section, we discuss related work and highlight how our model-based analysis of adherence to AI ethics guidelines, when understood as a social dilemma, provides substantive and actionable suggestions for fostering cooperation.

**Artificial Intelligence and Normative Guides.** As mentioned, non-legislative policy instruments, in the form of ethics guidelines, have proliferated in recent years. These guidelines, codes, and principles for the responsible creation and use of new technologies come from a wide array of sources, including academia, professional associations, and non-profit organisations;<sup>2</sup> governments;<sup>3</sup> and industry, including for-profit corporations.<sup>4</sup> Several researchers have noted that the very fact that a diverse set of stakeholders would exert such an effort to issue AI principles

<sup>1</sup>Of course, it is nontrivial to determine precisely what a ‘common good’ is; see discussion in [Green \(2019\)](#).

<sup>2</sup>See, for example [Future of Life Institute \(2017\)](#); [Gotterbarn et al. \(2018\)](#); [HAIP Initiative \(2018\)](#); [Information Technology Industry Council \(2017\)](#); [Partnership on AI \(2016\)](#); [Royal Statistical Society and the Institute and Faculty of Actuaries \(2019\)](#); [Stanford University \(2018\)](#); [The Future Society \(2017\)](#); [The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems \(2017\)](#); [The Japanese Society for Artificial Intelligence \(2017\)](#); [The Public Voice \(2018\)](#); [UNI Global Union \(2017\)](#); [Université de Montréal \(2017\)](#); [US Public Policy Council, Association for Computing Machinery \(2017\)](#).

<sup>3</sup>e.g., [European Group on Ethics in Science and New Technologies \(2018\)](#); [Government of Japan \(2017, 2018\)](#); [House of Lords, UK \(2018\)](#).

<sup>4</sup>e.g., [DeepMind \(2017\)](#); [Google \(2018\)](#); [IBM \(2017, 2018\)](#); [Microsoft \(2018\)](#); [OpenAI \(2018\)](#); [Sage \(2017\)](#); [SAP \(2018\)](#); [Sony \(2018\)](#).

and policies is strongly indicative that these stakeholders have a vested interest in shaping policies on AI ethics to fit their own priorities (Wagner, 2018; Benkler, 2019; Greene et al., 2019; Jobin et al., 2019).

Perhaps unsurprisingly, it has also been widely noted that ethics policies for AI research are typically or relatively ineffective: they do not have an actual impact on human decision-making in the field of AI/ML (Hagendorff, 2019); they are often too broad or high-level to guide ethics in practice (Whittlestone et al., 2019); and they are, by their very nature, voluntary and nonbinding (Campolo et al., 2017; Whittaker et al., 2018).

The content of ethics guidelines can be varied to promote adherence. For example, substantial changes in abstraction can help the application, impact, and influence of ethics principles for AI by way of *targeted* suggestions that can be implemented easily and concretely in practice (Gebru et al., 2020; Morley et al., 2019). However, while several authors have noted the inefficacy of guiding principles, and have offered proposals to surmount this inefficacy, such proposals are, themselves, often couched in inherently *normative* language. Therefore, they too fail to instantiate any mechanisms for their own reinforcement.

For example, ethics policies and guidelines can be understood *deontically*, insofar as they provide static, universal principles to be adhered to (Ananny, 2016; Mittelstadt, 2019). Hagendorff (2019) takes the inefficacy of such policies to be a failure of deontic ethics, proposing that AI ethics should instead be couched in a virtue-ethics framework.<sup>5</sup> But these types of solutions are themselves ineffective insofar as they are merely *meta-level normative principles* which recommend adherence to object-level principles. Some reflection shows that we can then ask the same question of the meta-level principles as we did of the object-level ones: Should we expect *these* to be effective? By dint of what? Since there is no mechanism by which these meta-level principles can be enforced, they provide no solution to the inefficacy created by the proliferation of principles and guides for ethical AI. Instead, the problem is merely reproduced a level up.

At the same time, many of the problems posed by the implementation of AI systems are *institutional* or *social* rather than explicitly technological.<sup>6</sup> This suggests the need for a structural analysis of the causes of social dilemmas in this particular context. Instead of asking ‘how can policies and procedures be applied or fulfilled to be more effective?’, we ask ‘how can we promote cooperation in *signing on* to such policies in the first place and adhering to these agreements in the second?’

<sup>5</sup>Virtue ethics is a moral theory which emphasises the role of an individual’s character and virtues in evaluating the rightness of actions (Anscombe, 1958; Aristotle, 1995; Crisp and Slote, 1997; Foot, 1978).

<sup>6</sup>See discussion in, e.g., Gabriel and Ghazavi (2021); Miceli et al. (2022); Falbo and LaCroix (2022).

The model that we present in Section 3 is quite general insofar as the dynamics would emerge in any interactive system of agents characterised by the incentive structure of the model. As a result, this model applies to social dilemmas much wider than our target of analysis (cooperative agreement in the context of AI research). However, we show how the results presented in Section 4 apply in the particular context of cooperation for safe and responsible AI research, and in Section 5 we draw some targeted recommendations for researchers in this field.

In asking how the application and fulfilment of AI ethics guidelines can be improved, we shift the focus from the *content* of the principles and related meta-ethical considerations to the *cooperative, social aspects* of the dilemma. Namely, we examine the circumstances under which we should expect cooperative success to be possible or likely in a socio-dynamical context. This analysis has significant practical implications for regulatory principles in AI/ML, which are made all the more pressing as these technologies are increasingly integrated into society.

### 3. A SOCIO-DYNAMIC MODEL OF COOPERATION FOR ETHICAL AI RESEARCH

In this section, we describe the model that we employ using standard techniques from evolutionary game theory. (See Appendix A for some additional background details.) This formal framework has been applied to social dilemmas like climate change (Finus, 2008; Pacheco et al., 2014; Wagner, 2001) nuclear proliferation (Brams and Kilgour, 1987a,b; Kraig, 1999; Zagare, 1987), conflicts over water resources (Madani, 2010), etc. However, it has yet to be applied to the unique cooperative challenges that are faced in the context of AI safety and ethics. The model we present is a straightforward adaptation of Pacheco et al. (2014), who are concerned primarily with cooperative dynamics in the context of climate change agreements.

There are, of course, analogies between the challenges of global cooperation in the context of climate change and AI research. This implies that at least some of the insights that can be gleaned from an analysis of the former challenge(s) may be relevant to the latter. For example, climate change and existential risk from AI both appear to be global challenges that require high degrees of collective cooperation with significant risk for the entire community. However, it is hardly obvious that these situations are similar enough to the AI context to warrant wholesale transfer of the conclusions of those models to the particular situation we discuss here. For example, contemporary public perception of the magnitude of risk and consequences of cooperative failure in the domain of AI research plausibly admits greater variance and uncertainty than it does for climate change. Hence, we highlight the relevance to cooperative AI research as we present the model.

We begin by discussing the mathematical structure of the model, including the payoff structure and dynamics. In Section 4, we investigate the effect of several variables on the possibility or likelihood of cooperation in the context of non-legislative policy instruments for AI research.

**Model and Parameters.** A population of size  $Z$  is organised (divided) into groups of size  $N$ . Each individual in the population may be interpreted as an AI researcher who can realise or ignore some proposed norm(s) of research and practice. The population can be thought of as the global AI community at large, and each group of size  $N$  in the population can be thought of as an organisational unit—e.g., a research laboratory, a university department, a collaborative research group, etc.—which can sustain shared norms. Individuals are classified according to the strategy that they choose with respect to the proposed norms for ethical behaviour within the community: they can choose either to cooperate ( $C$ ), by adhering to the proposed norm, or to defect ( $D$ ), by flouting the proposed norm.<sup>7</sup>

Each individual has an initial endowment,  $b \in \mathbb{R}^+$ . Contributing to a collective, cooperative effort imposes a cost,  $c \in [0, 1]$ , to the cooperator, consisting in a fraction of their endowment; defecting requires no such contribution. The cost may be construed as an additional effort relative to the non-cooperative baseline—e.g., taking extra precautions to ensure research and business practices adhere to the proposed normative standards—or in terms of opportunity costs—e.g., by refusing lucrative projects that would violate the proposed normative standards.

A normative agreement is successful if the fraction of contributors exceeds some threshold,  $p^* \in (0, 1)$ .<sup>8</sup> If this cooperative threshold is not achieved, then with some probability,  $r$ , everyone in the group loses a fraction,  $m$ , of their endowment.  $r$  can be interpreted as the *perceived* risk of negative consequences in the event of a failure to cooperate, and  $m$  can be understood as the *perceived* magnitude of those consequences. That is, when  $m = 1$  (the failure to cooperate is perceived to be maximally impactful), individuals risk losing all of their endowment;  $m = 0$  represents the situation in which failure to reach an agreement has no perceived negative consequence.<sup>9</sup>

This provides a five-dimensional space to analyse the possibility of cooperation in our social dilemma: the size of the groups in which cooperation is attempted,  $N$ ; the perceived risk of negative consequences when cooperation fails,  $r$ ; the perceived

<sup>7</sup>We assume here that individuals who choose to cooperate, or who say they will cooperate, in fact do so.

<sup>8</sup>A social threshold is common to many public endeavours—for example, international agreements that require a ratification threshold to take effect (Chen et al., 2012; Gokhale and Traulsen, 2010; Pacheco et al., 2009, 2014; Santos and Pacheco, 2011; Souza et al., 2009; Wang et al., 2009).

<sup>9</sup>Note that perceived risk of collective failure has proved important for successful collective action in dilemmas of this sort (Milinski et al., 2008; Pacheco et al., 2014; Santos and Pacheco, 2011).

magnitude of said consequences,  $m$ ; the cost of cooperation,  $c$ ; and the critical threshold of cooperators required to avoid negative consequences,  $p^*$ .

**Payoffs.** The payoffs for each strategy determine the game. An individual's payoff depends upon what action the individual chooses and what everyone else in the community is doing. The payoffs to cooperators in a group of size  $N$  when  $n_C$  individuals are cooperating is given in Equation 1,

$$(1) \quad \pi_C(n_C) = b \cdot \Theta(k) + b \cdot (1 - rm) \cdot (1 - \Theta(k)) - cb, \quad k = n_C - n^*.$$

$n^* = \lceil p^* \cdot N \rceil$  is the critical number of cooperators in a group, and  $\Theta$  is the Heaviside step function.<sup>10</sup>

Equation 1 has two key terms. The left-hand summand ( $b \cdot \Theta(k)$ ) provides payoff  $b$  to the cooperator just in case  $k = n_C - n^* \geq 0$ ; this is the payoff when cooperation succeeds. When cooperation fails ( $n_C - n^* \leq 0$ ), the right-hand summand ( $b \cdot (1 - rm) \cdot \Theta(k)$ ) comes into play. This payoff is weighted by the cost of failure as a function of the risk and magnitude of some negative consequence,  $(1 - rm)$ . So, when  $r$  and  $m$  are maximal ( $rm = 1$ ), the entire endowment is lost if cooperation fails; when the risk or magnitude are nonexistent ( $rm = 0$ ), none of the endowment is lost. Finally, the cost of cooperation,  $cb$ , is subtracted from the payoff for successful cooperation or failure to cooperate. The payoff to defectors is defined in terms of the payoff to cooperators, as in Equation 2:

$$(2) \quad \pi_D(n_C) = \pi_C(n_C) + cb.$$

We calculate the average payoff to each type,  $C, D$ , as a function of the group size,  $N$ , and the fraction of each type in the broader population,  $x_C$  and  $x_D = 1 - x_C$ , respectively. First, we find the fraction of cooperators in the population,  $x_C = n_C^Z/Z$ . We then calculate the vector of (binomially-distributed) probabilities for each possible group, composed of  $k$  cooperators and  $N - k$  defectors. We then compute the vector of conditional probabilities ( $k/N$ ) of being a cooperator in each combination. We compute the average payoff to the cooperators by weighing the payoff for cooperation by the probability of being a cooperator, as described. And, *mutatis mutandis* for the average payoff to defectors. The mean payoffs capture the expectation if the entire population,  $Z$ , were randomly paired into groups of size  $N$ . See Appendix B for formal details.

**Dynamics.** We consider the dynamics of small, finite populations. The dynamics determine how the population changes based on assumptions about how the payoffs affect the *fitness* of strategies, given what others are doing. Our model uses a dynamics called the *Fermi process* (Traulsen and Hauert, 2009). This involves a

<sup>10</sup>That is,  $\Theta(k) = 1$  when  $k \geq 0$ , and  $\Theta(k) = 0$  otherwise.

pairwise comparison where two individuals—a *focal individual* and a *role model*—are sampled from the population at random. The focal individual copies the strategy of the role model with probability  $p$ , depending on a comparison of the payoffs of those strategies. If both individuals have the same payoff, the focal individual randomises between the two strategies.<sup>11</sup> The probability is a nonlinear function of the payoff difference for  $p$ , called the *Fermi function* (Dirac, 1926; Fermi, 1926),

$$(3) \quad p = \left[ 1 + e^{\lambda(\pi_f - \pi_r)} \right]^{-1},$$

where  $\lambda \in \mathbb{R}^+$  is the intensity of selection, which specifies the importance of neutral drift compared to the selection dynamics.<sup>12</sup>  $\pi_f, \pi_r$  are the payoffs of the focal individual and the role model, respectively. In addition to the parameters described above, we further assume some small rate of mutation,  $\mu$ , which can be interpreted as mutation, experimentation, error, or noise. For all the results we present in Section 4, we set  $\lambda = 5, \mu = 0.10$ .

In our model, we specify the transition probabilities of the population changing from a state containing  $n_C^Z$  cooperators to one with  $(n_C^Z + 1)$  or  $(n_C^Z - 1)$  cooperators (i.e., the transition from the dynamics yielding one more or one fewer cooperator in the population at large). A transition matrix (for all possible states) is populated by recursive application of the transition probabilities. The transition matrix is used to compute the gradient of selection for our dynamics. This is defined by the difference between the probability that the number of cooperators increases and the probability that the number of cooperators decreases. We compute the stationary distribution of the process using standard techniques (Sandholm, 2007).<sup>13</sup> Again, see Appendix B for formal details.

<sup>11</sup>Note, then, that the focal individual does not always switch to a better strategy; the individual may switch to one that is strictly worse.

<sup>12</sup>Note that when  $\lambda = 0$ , selection is random; when selection is *weak* ( $w \ll 1$ ),  $p$  reduces to a linear function of the payoff difference; when  $\lambda = 1$ , our model gives us back the *replicator dynamic*; and, when  $\lambda \rightarrow \infty$ , we get the *best-response* dynamic (Fudenberg and Tirole, 1991).

<sup>13</sup>Note that we already know that the stationary distribution *exists*, because the addition of mutation makes the Markov process ergodic. That is to say, the Fermi process was already finite and aperiodic; with mutation, it is also irreducible (i.e., has only one recursive class). This is because there is a positive probability path between any two states, and in the limit every state will be visited an infinite number of times. In the absence of mutation, there are two (singleton) recursive classes corresponding to the two absorbing states where the population is composed entirely of defectors or entirely of cooperators. From being ergodic, it follows that the limit distribution is independent of any initial distributions.



#### 4. RESULTS

In this section, we examine the results of our model under simulation.<sup>14</sup> This is useful since human behaviour involves so many degrees of freedom that meaningful analytic results are often unlikely to be obtained. As is common practice in game-theoretic models of social behaviour, we are primarily interested in high-level qualitative results—such as whether cooperation will tend to succeed—rather than exact quantities.

**Simulation Parameter Values.** In our simulations, we fix the population,  $Z = 100$ , and the endowment,  $b = 1$ . All other parameters are variable: this includes the size of the groups,  $N \in \{1, 2, \dots, 20\}$ ; the perceived risk of negative consequences occurring if cooperation fails,  $r \in [0, 1]$ ; the cost of those consequences, understood in terms of the proportion of endowments lost, if cooperation fails,  $m \in [0, 1]$ ; the proportion of the endowment that is contributed by cooperators,  $c \in [0, 1]$ ; and the critical threshold of cooperators required in order for cooperation to succeed,  $p^* \in [0, 1]$ , which determines the number of cooperators required,  $n^* = \lceil p^* \cdot N \rceil$ .

**Qualitative Dynamics.** We compute the mean payoff to each strategy, the gradient of selection, and the stationary distribution for the dynamics as described in Section 3. The gradient of selection gives the probable direction of evolution in the short run, whereas the stationary distribution captures the long-run proportion of time that the process spends at any given state. This is the appropriate method of analysis for such finite-population dynamics (Mohseni, 2019; Traulsen and Hauert, 2009); see Figure 1 for examples.

The gradient may change from selection for an equilibrium where ALL DEFECT is the unique stable state (top of 1a), displaying the qualitative dynamics of a *prisoner's dilemma* (Serrano and Feldman, 2013); to a bi-stable dynamics with both an ALL DEFECT equilibrium and a POLYMORPHIC equilibrium, where both strategies can coexist in the population (top of 1b, 1c—the latter corresponds to an *anti-coordination* game (Rapoport and Chammah, 1966)); to (in rare circumstances) a unique equilibrium of ALL COOPERATE (top of 1d), where the qualitative dynamics yield a *prisoner's delight* (Binmore, 2004).

When there is a unique stable state, the stationary distribution reflects that the process spends the majority of its time near this state (bottom of 1a, 1d). However, when there are multiple stable states (1b, 1c), the dynamics may or may not guarantee population-level success relative to the demands of the cooperative challenge—the stationary distribution shows us the proportion of time spent at each

<sup>14</sup>Our code outputs visual graphics of the selection gradient, average payoffs to each strategy, and the stationary distribution. All of our simulation code is available online at <https://amohseni.shinyapps.io/tragedies-of-the-commons/>.

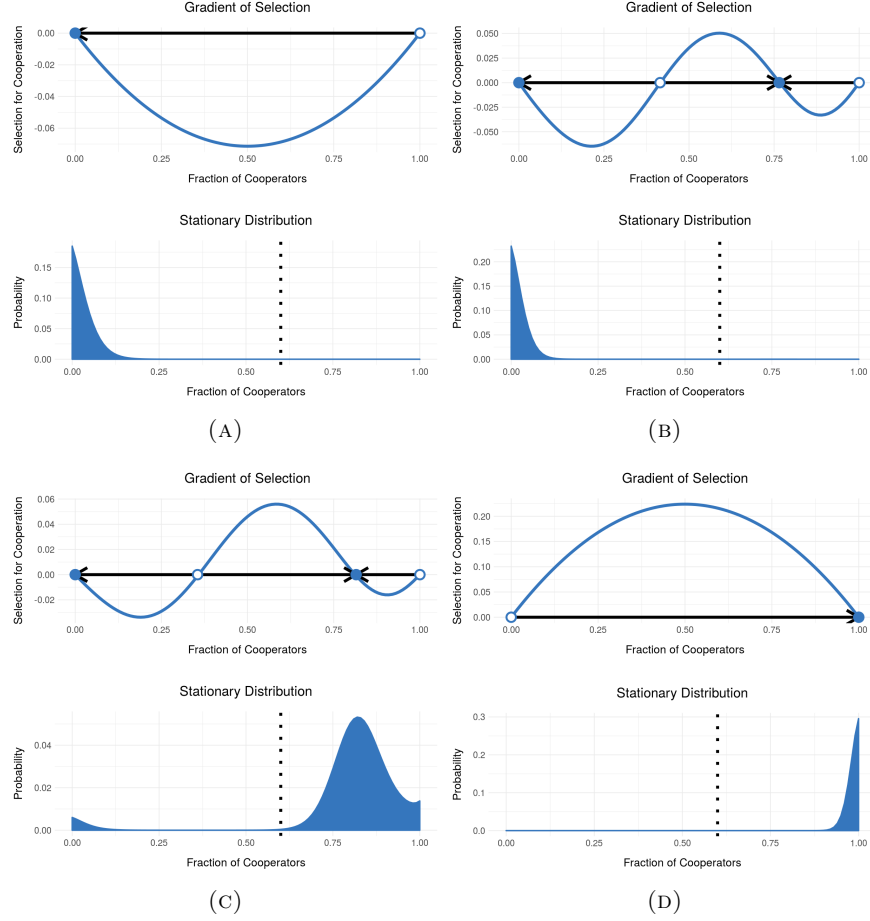


FIGURE 1. Examples of each qualitative dynamic for the gradient of selection (top) and stationary distribution (bottom) of the mean-field dynamics, resulting in (a) a prisoner's dilemma favouring ALL DEFECT; (b) a bi-stable dynamics favouring the ALL DEFECT equilibrium; (c) a bi-stable dynamics favouring the POLYMORPHIC equilibrium; and (d) a prisoner's delight favouring ALL COOPERATE.

state. Hence, although a mixed-population of cooperators and defectors is *stable* in Figure 1b, the stationary distribution indicates that the stable configuration of ALL DEFECT is most probable.

In sum, there are four possible qualitative outcomes for our population under the dynamics described: (1) ALL DEFECT is the unique stable state, near which the process spends all its time; (2) both ALL DEFECT and a POLYMORPHIC state in which both cooperators and defectors co-exists are stable, yet the process spends a majority of its time near the ALL DEFECT equilibrium; (3) both ALL DEFECT and a POLYMORPHIC state are stable, but the process spends the majority of its time

near the polymorphic mixture; and (4) ALL COOPERATE is uniquely stable and the process tends toward this state.

**Simulation Results.** We want to know how the parameters of interest interact with the achievement of cooperative success (the situations shown in Figures 1c or 1d), and which of those will have the greatest effect. Our dynamics encode a process by which strategies that are more successful proliferate via imitation; so, for cooperation to succeed it must be that cooperators do better on average than defectors. Each of the parameters will either work for or against cooperative success. We will examine their respective effects by first considering their extremal values.

It should be obvious that when the cost of cooperation is maximal ( $c = 1$ ), *defect* ( $D$ ) will dominate since cooperators have nothing to gain from cooperation in this scenario. Therefore, the lower the cost of cooperation, the more likely individuals will be to cooperate. However, we will see that even a low cost of cooperation does not ensure cooperative success.

At the other extreme, when the cost of cooperation is minimal ( $c = 0$ ), *cooperate* ( $C$ ) is the strictly dominant strategy in almost all cases—the exception is when there is no benefit to cooperation ( $r \cdot m = 0$ ), and so the payoffs for  $C$  and  $D$  are equivalent. On the other hand, if the cost of cooperation is nonzero ( $c > 0$ ), and there is no payoff for cooperative success ( $r \cdot m = 0$ ), then  $D$  is strictly dominant. That is, improbable but highly consequential failures to cooperate and highly probable but inconsequential failures to cooperate will evoke the same cooperative (or non-cooperative) response. Transitions between qualitative dynamics as a function of the risk,  $r$ , and magnitude,  $m$ , of the consequences of cooperative failure are given in Figure 2.

Further, Figure 3 makes clear that cooperative success requires the ratio of the benefit of avoiding the failure to cooperate ( $r \cdot m$ ) and the cost of cooperation ( $c$ ) to be sufficiently favourable.

Figures 2 and 3 also illustrate the critical role that group size plays in the prospects of cooperative success. In particular, all else equal, smaller group sizes are more favourable for cooperation. This is perhaps counter-intuitive; however, the result can be explained as follows. Recall that cooperators must obtain higher average payoffs for cooperation to be favoured. Moreover, agents contribute to the prevalence of cooperation in their respective groups by choosing to cooperate (or not).

When group size is minimal ( $N = 1$ ), an agent who chooses to cooperate ensures cooperative success by dint of her action alone. Thus, all groups composed of cooperators succeed, and all groups composed of defectors fail. (This holds for

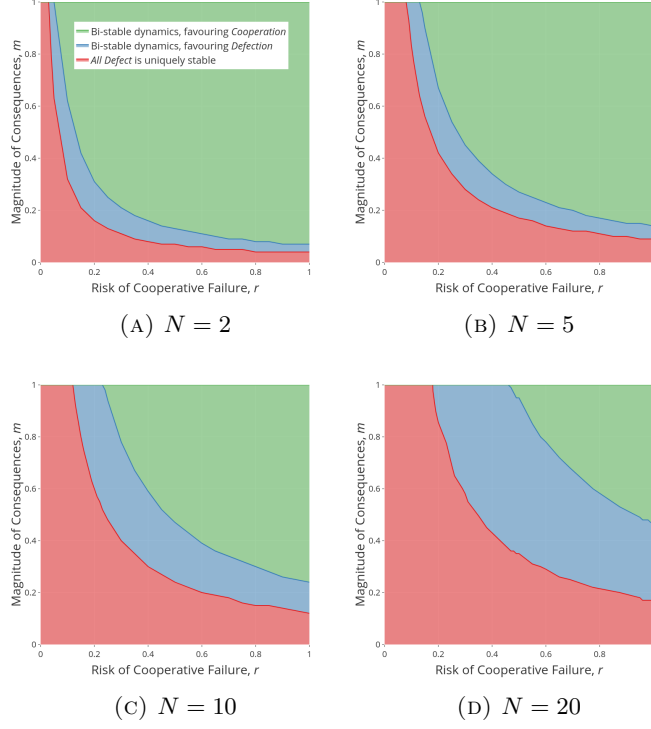


FIGURE 2. Qualitative dynamics as a function of the risk,  $r$ , and magnitude,  $m$ , of the consequences of a failure to cooperate. In each case,  $p^* = 0.5$ ,  $c = 0.01$ .

non-extremal values; i.e.,  $c > 0$ ,  $m \cdot r > 0$ , and  $p^* > 0$ .) This is sufficient to ensure selection for cooperation.

At the other extreme, when the cooperative group grows arbitrarily large ( $N \rightarrow \infty$ ), an agent's action makes a vanishing contribution to the prospect of cooperative success.<sup>15</sup> Furthermore, by the law of large numbers (Bernoulli, 2005; Poisson, 1837), the proportion of cooperators and defectors in large groups will approach the precise proportion of cooperators and defectors in the whole population. Hence, any group will tend to be equally likely to succeed or fail. Yet, in such groups, defectors will obtain higher payoffs by dint of not having paid the cost to cooperate. This is sufficient to ensure selection for defection.

In between these extremes, smaller group sizes will tend to favour cooperation and disfavour defection. Thus, we find a special instance of the general pattern that correlation between strategic actions conduces to pro-social behaviour (Hamilton,

<sup>15</sup>The reader may identify that this strategic structure is analogous to that of the Paradox of Voting (de Caritat Condorcet, 1793).

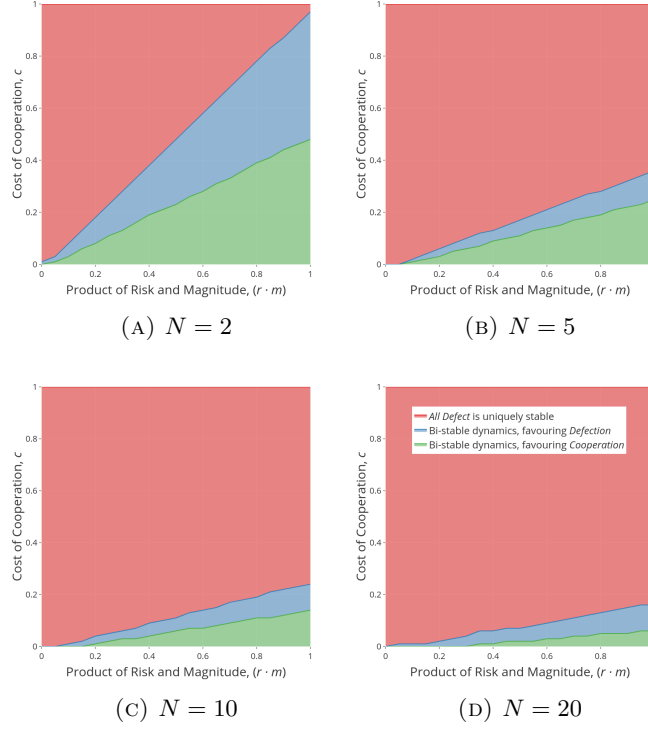


FIGURE 3. Qualitative dynamics as a function of the product of the risk,  $r$ , and magnitude,  $m$ , of the consequences to fail to cooperate ( $r \cdot m$ ) and the cost of cooperation,  $c$ . In each case,  $p^* = 0.50$ .

1964a,b; Skyrms, 1994), where here the relevant form of correlation is realised by the contribution of the individual's action to the composition of the group's actions.<sup>16</sup>

Finally, we can consider the effect of the critical threshold for cooperation,  $p^*$ , on the prospects for cooperative success; see Figure 4.

Logically, at the extreme where the critical threshold is minimal ( $p^* = 0$ ), cooperation 'succeeds' without need of cooperators, and so defection is dominant. However, when the critical threshold is nonzero ( $p^* > 0$ ), the general pattern is that the smaller the fraction of cooperators required for cooperation, the more likely cooperation is to succeed. Indeed, sufficiently demanding cooperative endeavours mean that while a POLYMORPHIC equilibrium in which cooperation co-exists

<sup>16</sup>Correlation can be realised variously in a social dilemma—e.g., assortative mating (Eshel and Cavalli-Sforza, 1982; Hamilton, 1971), kin selection (Hamilton, 1963; Maynard Smith and Price, 1964), homophily (McPherson et al., 2001), and network effects (Broere et al., 2017), among others. All of these support cooperation insofar as they make cooperators more likely to interact with one another, and less likely to interact with defectors. Although we lack the space to discuss these here, they constitute an important further dimension of our analysis.

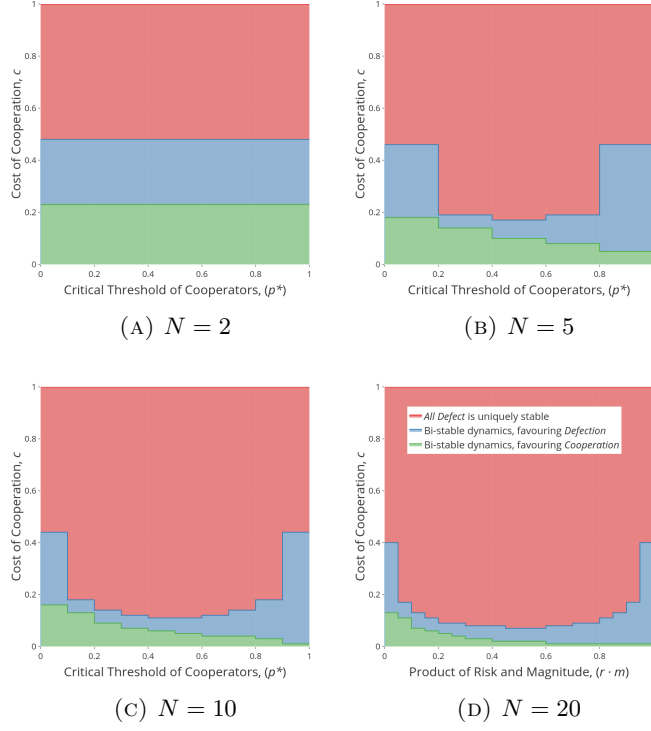


FIGURE 4. Qualitative dynamics as a function of the critical fraction of cooperators required for success,  $p^*$ , and the cost of cooperation,  $c$ . In each case,  $r \cdot m = 0.50$ .

with defection is stable under the gradient of selection, the stationary distribution demonstrates that arriving and remaining at such a cooperative state may be vanishingly unlikely.

**Robustness.** In such models, results need be robust to changes in two sets of parameters: the parameters of the payoff functions, and the parameters of the dynamics.

The results we describe in Section 4 are that the selective advantage of cooperation is monotonically *decreasing* in the cost of cooperation,  $c$ , and the size of cooperative groups,  $N$ ; and, the selective advantage of cooperation is monotonically *increasing* in the product of the perceived probability,  $r$ , and magnitude,  $m$ , of the consequences of failing to successfully cooperate.

These findings hold for the entire parameter space of the game. This is demonstrated qualitatively by Figures 2, 3, and 4; however, these facts can also be proved analytically. See Appendix C.

In addition to claims of robustness with respect to the aforementioned parameters, our results are also robust with respect to changes of the parameters of the dynamics—the Fermi process.

In particular, it is easy to demonstrate that the qualitative dynamics of the game—which determine whether a strategy is being selected for or against—are invariant for the range of conditions characterised by positive, finite intensities of selection,  $0 < \lambda < \infty$ , and positive mutation rates short of complete randomness,  $0 < \mu < 1$ .

Thus, the results presented are indeed robust under any reasonable sensitivity analysis.

## 5. DISCUSSION

*A priori*, one might have thought that inhibitive costs furnish the largest obstacle for cooperative success. Our model provides evidence for this intuition. Thus, our first moral is the obvious one.

***Moral 1:*** *Lowering the cost of cooperation increases the likelihood of cooperative success.*<sup>17</sup>

However, our model highlights that the dynamics of cooperation here are more subtle than this. We show that even for small costs, the dynamics may not guarantee population-level success relative to the demands of the cooperative challenge for a range of parameter values.<sup>18</sup> Perhaps most surprising is the magnitude of the effect of group size on these outcomes. Thus, we propose the following.

***Moral 2:*** *Small, decentralised groups may benefit sustained cooperation for responsible AI research.*

In the context of globalisation, one might have thought that the best approach for fostering AI ethics and safety was to have a universal set of principles to which

<sup>17</sup>This moral pertains to the likelihood of signing on to an agreement in the first place, but there is also a question of whether individuals who *say* they will cooperate in fact do so cooperate. When signals are cheap, they can be uninformative or dishonest (Crawford and Sobel, 1982; Farrell, 1987; Farrell and Rabin, 1996; Wärneryd, 1993). It is well-understood that costly signals can promote honesty (Johnstone, 1995; Lachmann et al., 2001; Pomiankowski, 1987; Zahavi, 1975; Zahavi and Zahavi, 1997).

<sup>18</sup>In our model, the cost for cooperation is nonnegative. So, we do not account for *incentives* to cooperate—i.e., rewards. Conversely, we could lower the payoff for defectors by introducing punishment for non-cooperation. This is already something that has been done by, e.g., the ACM (ACM, 2020). Although, empirical data suggests rewards are better than punishments for promoting cooperative behaviour in similar social dilemmas (DeSombre, 2000; Kaniaru et al., 2007). Even so, determining what costs/rewards are, how much they are, and how they are distributed is highly nontrivial.

everyone agrees. However, our model identifies that cooperation spreads as individuals observe other groups succeeding in their cooperative endeavours; this is fostered by the existence of many smaller groups where cooperation can more easily get off the ground. Hence, targeting policy agreements at smaller, independent groups, such as individual research labs or professional organisations, may constitute a more effective path to the aims of the universal adoption of ethical guidelines.

We mention how small group sizes help cooperation *in virtue of* the relationship between positive correlation between strategies and selection for pro-social behaviour. Yet another way to realise such correlation is via free partner choice. When individuals who are inclined to cooperate in adherence to norms are able to freely join one another in their efforts, this can increase their relative likelihood of success and so spread their example throughout the population. Thus, we get the following corollary.

***Moral 3:*** *Voluntary participation in AI policy agreements may catalyse the spread of cooperation.*

Individuals choosing to adhere to a set of principles (or not) creates positive correlation. So, affording individuals the option of choosing their cooperative group may be more beneficial than stipulating that one must be a part of the group. However, even in small groups cooperation can fail, as we have demonstrated. Hence,

***Moral 4:*** *It is important to accurately figure both the risks and the consequences of non-cooperation.*

The action-relevant parameters in our model are *perceived* risk and magnitude. Therefore, it is essential to accurately figure the *actual* risks and potential consequences of non-cooperation. Namely, if the perceived risk is significantly lower than the actual risk, then cooperation may fail to obtain despite its (understated) real-world importance. Failing to cooperate, when cooperation matters, incurs a greater loss than cooperating when it was not entirely necessary to do so; hence, there may be less harm in *overstating* the negative consequences of failing to cooperate.<sup>19</sup> One real-world suggestion that we might glean from this is that it may be beneficial to further incorporate education on AI ethics and safety across a breadth of educational curricula and public outreach efforts.

The insights thus far centre primarily on group dynamics; this says nothing whatsoever about the actual content of ethics guidelines for AI. Of course, the content

---

<sup>19</sup>This is a typical line of argument in much of the existential risk literature; see, e.g., [Russell \(2019\)](#).



of these guides can significantly affect the likelihood that community-wide norms are taken up. For example, if the proposed norms are impossible to be instantiated, then they will fail to be adopted. Our final moral addresses this.

**Moral 5:** *Combining many proposals may undermine their prospects for success.*

The likelihood of being willing to adhere to and able to fulfil a set of policies is bounded by the likelihood of adhering to and fulfilling the conjunction of the policies it contains. We show that policies can vary in the magnitude of the cooperative challenge they present. The combination of these facts should make us wary of how demanding we make our policy proposals, if initial and sustained adherence is our aim. Instead, a piecemeal approach in which agents can sign on and gain from participation in single policies may provide a better first step on the path to the ultimate fixation of a robust set of norms for safe and ethical AI.

In closing, we note that even disciplines and fields whose norms of self-regulation are relatively well-entrenched, as in medical practice and biomedical ethics, guiding principles can be ineffective, or at least inefficient, in guiding actual practice when there are competing incentives (Lomas et al., 1989).<sup>20</sup> This is true even when there is consensus amongst researchers and practitioners about the guidelines themselves (Lomas, 1991)—and, it should be apparent that no such consensus exists in AI/ML.

In contrast to other work that has highlighted the inefficacy of non-legislative policy agreements in AI/ML, our analysis focused on the socio-dynamics of cooperation. In doing so, we were able to discuss the circumstances under which we should expect cooperation to arise in the first place. Further, instead of providing normative suggestions for increasing the effectiveness of extant guidelines by varying their content, the normative suggestions that we forward here pertain to the social aspects of cooperation; as such, we do not fall prey to the meta-normative inefficacy described in Section 2.

We conclude with a final note about the target of our model, which addresses the proliferation of ethics principles and explores the conditions under which one might expect them to be more or less likely to garner broad adherence or adoption. We do not take our model to diagnose any of these challenges in their entirety; instead, we see it as providing tentative hypotheses that must be combined with empirical evidence and experimentation to inform our judgements adequately.

As a concrete example, the United Nations Educational, Scientific, and Cultural Organisation (UNESCO) voted, in November 2019, to begin a two-year process to

---

<sup>20</sup>Although, this is not to say that the solution is to simply impose *hard* laws—this will likely also be ineffective; see discussion in LaCroix and Bengio (2019).

create ‘the first global standard-setting instrument on the ethics of artificial intelligence in the form of a Recommendation’—i.e., a non-legislative policy instrument of the type discussed in Section 2. An *Ad Hoc Expert Group* (AHEG) prepared a first draft of the text for which they solicited public feedback from a wide range of stakeholders over the Summer, 2020. The AHEG will account for this feedback while preparing a final draft, whose adoption will be voted on by UNESCO at the end of 2021.

We suggest that there may be little reason to think that this document will be much more effective than any of the other 84 documents that have been proposed by NGOs, governments, corporations, professional associations, etc., for the reasons discussed in Section 2.

So, the target of our paper can be understood, thus. Suppose you are a member of the AHEG, and you are charged with writing a policies guideline for UNESCO. The hope is that UNESCO’s 193 member states and 11 associate members will ratify these recommendations (presumably because of wide support from researchers, labs, corporations, etc. within those member states). Ratification, then, is a decision to cooperate in the language of our model, and non-ratification is a decision to defect. (Note that ratifying such a document incurs costs to the member state, whereas refusing to do so incurs no such cost.<sup>21</sup>)

We can ask how the morals that we have gleaned from our model come to bear on this particular situation in the concrete. According to *Moral 1*, it is essential to keep the costs of cooperation low in order to attain the critical mass of cooperators required to make cooperation viable in the first place. Once such a critical mass of cooperators is acquired, the inertia of the system increases the benefits of choosing to cooperate, thus making this option more appealing.

*Moral 2* is inherently difficult in the context of a *global* ethics recommendation for AI research, which requires cooperation that spans international borders. We suggested that sustained cooperation is benefited by targeting small, decentralised groups. It should be clear that the member states of UNESCO are neither small nor decentralised. However, it may be possible for UNESCO to provide a *framework* of principles for which individual states have significant autonomy to fill in the actual details of how that framework is implemented, or what specific policies are instantiated on the road to achieving the larger goals of the global cooperative agreement. This is analogous to political federalism, like the governmental structure of Canada, wherein there is a division of power between different levels of government—federal, provincial, territorial, or municipal. Increased autonomy for how the principles are

<sup>21</sup>Assuming, of course, that the reputational costs incurred by not cooperating are smaller than the costs incurred for cooperating. However, note that reputational costs are endogenous, and are not imposed by the proposal itself.

instantiated means reduced centralisation and has the added benefit of creating smaller, cooperative subgroups.

*Moral 4* suggests that cooperative behaviour may be benefited by soliciting recommendations of both experts and the broader public. It is always the case that there may be some misalignment between those problems that the public deems important versus those that experts in the field do so. When only one of these groups is queried, the relative risks and consequences of failures to cooperate may be undervalued for certain commitments that are deemed important by the other group. As such, even though some risks and consequences may be overstated, this may serve to benefit cooperation in the long run. Furthermore, accurately figuring the risks and consequences of any given scenario will likely require broad interdisciplinary collaboration. So, an AHEG should be composed of experts from a wide range of disciplines (not just computer science), and this expert group may benefit from soliciting public feedback on their proposals. Fortunately, this is precisely the course that UNESCO has taken in this particular instance—the AHEG for the recommendation on the ethics of artificial intelligence is composed of 25 researchers from as many different countries, and these individuals have diverse backgrounds, including computer and information science, philosophy, law, and engineering (UNESCO, 2020).

Finally, *Moral 5* suggests that the proposals should start with modest goals, targeting a small number of specific aspects of safe AI research that ought to be afforded priority. It is always the case that these recommendations can be amended or strengthened in the future. Furthermore, *Moral 3* suggests that voluntary participation in such a policy agreement may catalyse the spread of cooperation. Thus, it may be possible to create a positive feedback loop wherein an organisation like UNESCO suggests a small number of modest proposals which incur only a low cost for cooperators to agree. These aspects of the agreement serve to incentivise individuals to voluntarily sign on to such an agreement which, in turn, creates positive correlation that promotes further cooperation. This creates a critical mass of cooperators which then provides further endogenous incentives for individuals to cooperate.

This concrete example demonstrates how the morals that we derive from our modelling work can be applied, and how these various considerations may have interaction effects that may help nurture population-level success relative to the demands of the cooperative challenge we discuss. Of course, some epistemic humility is a virtue in modelling work of the sort we employ. The results derived from mathematical models of socio-dynamic phenomena are best understood as alerting us to the lower-bound on the complexity of those phenomena, and as providing

tentative hypotheses that we must combine with empirical evidence and experimentation to adequately inform our judgements. It would be an error to observe the results of such a model and to become convinced that one fully apprehends the cooperative challenge it depicts.

## 6. CONCLUSION

It would be a mistake to interpret our results and arguments as a dismissal of current proposals or guidelines for the safe and ethical development of AI. To the contrary, we believe that such efforts are laudable and necessary; but to give ourselves reasonable odds for success, we must appraise ourselves of an understanding of the basic dynamics of such coordination problems. These mistaken interpretations of our results—with their concomitant negative impacts—must be avoided.

From nuclear proliferation and climate change to ethical AI and AI safety, social dilemmas characterise the gulf between us and the futures for which we hope. The potential positive social impacts of our work are plain. Artificial intelligence promises to fundamentally change nearly every facet of our lives, for better or worse. Hence, effective adherence to ethical norms throughout the process of making numerous inevitable advances in AI will make a difference to the tally of instances in which the process promotes or deranges the prospects of human flourishing.

Our contribution to this count is simple but, we believe, an important step forward for advancing a productive way of framing this cooperative problem. We deploy theoretical tools from evolutionary game theory to analyse the nature of the social dilemma at play in promoting participation in, and adherence to, the proposed policies and norms. We provide insights that, on the one hand, have not obviously informed extant guidelines and policies, and that, on the other hand, correspond to tractable changes in those proposals which may yield significant impact. If we succeed, stakeholders—research laboratories, university departments, collaborative research groups, and so on—aiming to formulate and coordinate around policy guidelines may have a richer awareness of the challenges involved that may inform the scope, scale, and content of such proposals.

**Acknowledgments.** Thanks to Assya Trofimov, Joey Bose, Sarah LaCroix, Duncan MacIntosh, and Daniel Herrmann for helpful comments on an early draft. Thanks also to Ioannis Mitliagkas, Dominic Martin, Yoshua Bengio, Brian Skyrms, Jeffrey Barrett, Gillian Hadfield, and audiences at the Schwartz Reisman Institute’s weekly seminar series in Toronto (December 2020), and the Philosophy of Science Association’s virtual and in-person poster forums (January/November 2021). The authors would like to thank the anonymous referees for their helpful comments. Thanks also to the Schwartz Reisman Institute for Technology and Society at the University of Toronto and Mila - Québec Artificial Intelligence Institute at Université de Montréal for funding this research.

## REFERENCES

- ACM (2020). ACM code of ethics enforcement procedures. <https://www.acm.org/code-of-ethics/enforcement-procedures>.
- Alexander, J. McKenzie (2007). *The Structural Evolution of Morality*. Cambridge University Press, Cambridge.
- Allison, Scott T. and Norbert L. Kerr (1994). Group Correspondence Biases and the Provision of Public Goods. *Journal of Personality and Social Psychology*, 66(4): 688–698.
- Altrock, Philipp M. and Arne Traulsen (2009). Fixation Times in Evolutionary Games under Weak Selection. *New Journal of Physics*, 11: 013012.
- Ananny, Mike (2016). Toward an Ethics of Algorithms: Convening, Observation, Probability, and Timeliness. *Science, Technology, & Human Values*, 41(1): 93–117.
- Anscombe, G. E. M. (1958). Modern Moral Philosophy. *Philosophy*, 33(124): 1–19.
- Aristotle (1995). Nichomachean ethics. In Barnes, Jonathan, editor, *The Complete Works of Aristotle, The Revised Oxford Translation*, volume 2, pages 1729–1867. Princeton University Press, Princeton.
- Ashcroft, Peter, Philipp M. Altrock, and Tobias Galla (2014). Fixation in Finite Populations Evolving in Fluctuating Environments. *Journal of the Royal Society Interface*, 11: 20140663.
- Aumann, Robert and Sergiu Hart (1992). *Handbook of Game Theory with Economic Applications*, volume 1. Elsevier, North-Holland.
- Aumann, Robert and Sergiu Hart (1994). *Handbook of Game Theory with Economic Applications*, volume 2. Elsevier, North-Holland.
- Aumann, Robert and Sergiu Hart (2002). *Handbook of Game Theory with Economic Applications*, volume 3. Elsevier, North-Holland.
- Axelrod, Robert (1981). An Evolutionary Approach to Norms. *American Political Science Review*, 80(4): 1095–1111.
- Axelrod, Robert and William D. Hamilton (1981). The Evolution of Cooperation. *Science*, 211(4489): 1390–1396.
- Barrett, Jeffrey (2007). Dynamic Partitioning and the Conventionality of Kinds. *Philosophy of Science*, 74: 527–546.
- Benkler, Yochai (2019). Don’t Let Industry Write the Rules for AI. *Nature*, 569: 161.
- Bernoulli, Jacob (1713/2005). *Ars Conjectandi: Usum & Applicationem Praecedentis Doctrinae in Civilibus, Moralibus & Oeconomicis [The Art of Conjecture]*. John Hopkins University Press, Baltimore.
- Bicchieri, Cristina (2006). *The Grammar of Society*. Cambridge University Press, Cambridge.
- Binmore, Ken G. (2004). Reciprocity and the Social Contract. *Politics, Philosophy & Economics*, 3: 5–35.
- Binmore, Ken G. and Larry Samuelson (1994). An Economist’s Perspective on the Evolution of Norms. *Journal of Institutional and Theoretical Economics*, 150(1): 45–63.
- Boehm, C. (1982). The Evolutionary Development of Morality as an Effect of Dominance Behavior and Conflict Interference. *Journal of Social and Biological Structures*, 5: 413–421.

- Brams, Steven J. and D. Marc Kilgour (1987a). Threat Escalation and Crisis Stability: A Game-theoretic Analysis. *American Political Science Review*, 81(3): 833–850.
- Brams, Steven J. and D. Marc Kilgour (1987b). Winding Down if Preemption or Escalation Occurs: A Game-Theoretic Analysis. *Journal of Conflict Resolution*, 31(4): 547–572.
- Broere, Joris, Vincent Buskens, Jeroen Weesie, and Henk Stoof (2017). Network Effects on Coordination in Asymmetric Games. *Scientific Reports*, 7: 17016.
- Campolo, Alex, Madelyn Sanfilippo, Meredith Whittaker, and Kate Crawford (2017). *AI Now 2017 Report*. AI Now Institute at New York University, New York.
- Chalub, Fabio A. C. C., Francisco C. Santos, and Jorge M. Pacheco (2006). The evolution of norms. *Journal of Theoretical Biology*, 241: 233–240.
- Chen, Xiaojie, Attila Szolnoki, and Matjaž Perc (2012). Risk-driven Migration and the Collective-risk Social Dilemma. *Physical Review E: Statistical, Nonlinear, Biological, and Soft Matter Physics*, 86: 036101.
- Claussen, Jens Christian and Arne Traulsen (2005). Non-Gaussian Fluctuations Arising from Finite Populations: Exact Results for the Evolutionary Moran process. *Physical Review E: Statistical, Nonlinear, Biological, and Soft Matter Physics*, 71(2 pt. 2): 025010.
- Crawford, Vincent P. and Joel Sobel (1982). Strategic Information Transmission. *Econometrica*, 50(6): 1431–1451.
- Crisp, Roger and Michael Slote (1997). *Virtue Ethics*. Oxford University Press, Oxford.
- Darwin, Charles (1981/1871). *The Descent of Man, and Selection in Relation to Sex*. Princeton University Press, Princeton.
- Dawes, Robyn (1980). Social Dilemmas. *Annual Review of Psychology*, 31: 169–193.
- de Caritat Condorcet, Marie Jean Antoine Nicolas (1793). *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix [Essay on the Application of Analysis to the Probability of Majority Decisions]*. L'imprimerie Royale, Paris.
- DeepMind (2017). DeepMind Ethics & Society Principles. <https://deepmind.com/applied/deepmind-ethics-society/principles/>.
- DeSombre, Elizabeth R. (2000). The Experience of the Montréal Protocol: Particularly Remarkable, and Remarkably Particular. *UCLA Journal of Environmental Law and Policy*, 19: 49–82.
- Dirac, Paul A. M. (1926). On the Theory of Quantum Mechanics. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 112(762): 661–677.
- Eshel, Ilan and L. L. Cavalli-Sforza (1982). Assortment of Encounters and the Evolution of Cooperativeness. *Proceedings of the National Academy of Sciences of the United States of America*, 79: 1331–1335.
- European Group on Ethics in Science and New Technologies (2018). Statement on artificial intelligence, robotics and ‘autonomous’ systems. [http://ec.europa.eu/research/ege/pdf/ege\\_ai\\_statement\\_2018.pdf](http://ec.europa.eu/research/ege/pdf/ege_ai_statement_2018.pdf).
- Falbo, Arianna and Travis LaCroix (2022). Est-ce que vous compute? Code-switching, cultural identity, and AI. *arXiv pre-print*, 2112.08256: 1–19. Forthcoming in *Feminist Philosophical Quarterly*. <http://arxiv.org/abs/2112.08256>.

- Farrell, Joseph (1987). Cheap Talk, Coordination, and Entry. *The RAND Journal of Economics*, 18(1): 34–39.
- Farrell, Joseph and Matthew Rabin (1996). Cheap Talk. *Journal of Economic Perspectives*, 10(3): 103–118.
- Fehl, Katrin, Daniel J. van der Post, and Dirk Semmann (2011). Co-evolution of Behaviour and Social Network Structure Promotes Human Cooperation. *Ecology Letters*, 14(6): 546–551.
- Fermi, Enrico (1926). Sulla quantizzazione del gas perfetto monoatomico [On the Quantization of the Monoatomic Ideal Gas]. *Rendiconti Lincei. Scienze Fisiche e Naturali*, 3: 181–185.
- Finus, Michael (2008). Game Theoretic Research on the Design of International Environmental Agreements: Insights, Critical Remarks, and Future Challenges. *International Review of Environmental and Resource Economics*, 2(1): 29–67.
- Fishman, Michael A. (2006). Involuntary Defection and the Evolutionary Origins of Empathy. *Journal of Theoretical Biology*, 242: 873–879.
- Fletcher, Jeffrey A. and Martin Zwick (2007). The Evolution of Altruism: Game Theory in Multilevel Selection and Inclusive Fitness. *Journal of Theoretical Biology*, 245: 26–36.
- Foot, Philippa (1978). *Virtues and Vices and Other Essays in Moral Philosophy*. Oxford University Press, Oxford.
- Fudenberg, Drew and Jean Tirole (1991). *Game theory*. The MIT Press, Cambridge, MA.
- Future of Life Institute (2017). Asilomar AI Principles. <https://futureoflife.org/ai-principles/>.
- Gabriel, Iason and Vafa Ghazavi (2021). The Challenge of Value Alignment: from Fairer Algorithms to AI Safety. *arXiv pre-print*, 2101.06060: 1–20. <http://arxiv.org/abs/2101.06060>.
- Gebru, Timnit, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Dauméé, III, and Kate Crawford (2020). Datasheets for Datasets. *arXiv pre-print*, abs/1803.09010: 1–24. <https://arxiv.org/abs/1803.09010>.
- Gintis, Herbert (2000). *Game Theory Evolving: A Problem-Centered Introduction to Modeling Strategic Behavior*. Princeton University Press, Princeton.
- Gintis, Herbert, Samuel Bowles, Robert Boyd, and Ernst Fehr (2003). Explaining Altruistic Behavior in Humans. *Evolution and Human Behavior*, 24: 153–172.
- Gokhale, Chaitanya S. and Arne Traulsen (2010). Evolutionary Games in the Multiverse. *Proceedings of the National Academy of Sciences of the United States of America*, 107(12): 5500.
- Google (2018). AI at Google: Our Principles. <https://ai.google/principles>.
- Gotterbarn, Don, Amy Bruckman, Catherine Flick, Keith Miller, and Marty J. Wolf (2018). ACM Code of Ethics: A Guide for Positive Action. *Communications of the ACM*, 61(1): 121–128.
- Government of Japan, Ministry of Internal Affairs & Communications (MIC) (2017). AI R&D principles. [http://www.soumu.go.jp/main\\_content/000507517.pdf](http://www.soumu.go.jp/main_content/000507517.pdf).
- Government of Japan, Ministry of Internal Affairs & Communications (MIC) (2018). Draft AI utilization principles. [http://www.soumu.go.jp/main\\_content/000581310.pdf](http://www.soumu.go.jp/main_content/000581310.pdf).



- Green, Ben (2019). ‘Good’ Isn’t Good Enough. *Proceedings of the AI for Social Good workshop at NeurIPS*, pages 1–7.
- Greene, Daniel, Anna Lauren Hoffmann, and Luke Stark (2019). Better, Nicer, Clearer, Fairer: A Critical Assessment of the Movement for Ethical Artificial Intelligence and Machine Learning. In *52nd Hawaii International Conference on System Sciences*, pages 2122–2131, Honolulu. Hawaii International Conference on System Sciences (HICSS).
- Grujić, Jelena, Carlos Gracia-Lazaro, Manfred Milinski, Dirk Semmann, Arne Traulsen, José A. Cuesta, Yamir Moreno, and Angel Sánchez (2015). A comparative analysis of spatial Prisoner’s Dilemma experiments: Conditional cooperation and payoff irrelevance. *Sci. Rep.*, 4(4615): srep04615.
- Grujić, Jelena, Torsten Rohl, Dirk Semmann, Manfred Milinski, and Arne Traulsen (2012). Consistent strategy updating in spatial and non-spatial behavioral experiments does not promote cooperation in social networks. *PLoS One*, 7(11): e47718.
- Hagendorff, Thilo (2019). The Ethics of AI Ethics: An Evaluation of Guidelines. *arXiv pre-print*, abs/1903.03425: 1–16. <http://arxiv.org/abs/1903.03425>.
- HAIP Initiative (2018). Harmonious Artificial Intelligence Principles (HAIP). <http://bii.ia.ac.cn/hai/index.php>.
- Hamilton, William D. (1963). The Evolution of Altruistic Behavior. *The American Naturalist*, 9: 354–356.
- Hamilton, William D. (1964a). The Genetical Evolution of Social Behaviour. I. *Journal of Theoretical Biology*, 7: 1–16.
- Hamilton, William D. (1964b). The Genetical Evolution of Social Behaviour. II. *Journal of Theoretical Biology*, 7: 17–52.
- Hamilton, William D. (1971). Selection of Selfish and Altruistic Behavior in Some Extreme Models. In Eisenberg, J. F. and W. S. Dillon, editors, *Man and Beast*, pages 59–91. Smithsonian Institution Press, Washington.
- Harari, Yuval Noah (2017). Reboot for the AI revolution. *Nature*, 550: 324–327.
- Harms, William and Brian Skyrms (2008). Evolution of Moral Norms. In Ruse, Michael, editor, *The Oxford Handbook of Philosophy of Biology*, pages 434–450. Oxford University Press, Oxford.
- Hauert, Christoph, Miranda Holmes, and Michael Doebeli (2006). Evolutionary Games and Population Dynamics: Maintenance of Cooperation in Public Goods Games. *Proceeding of the Royal Society B: Biological Science*, 273(1600): 2565–2570.
- Hausken, Kjell and Jack Hirshleifer (2008). Truthful Signalling, the Heritability Paradox, and the Malthusian Equi-Marginal Principle. *Theoretical Population Biology*, 73: 11–23.
- Helbing, Dirk (2019). *Towards Digital Enlightenment: Essays on the Dark and Light Sides of the Digital Revolution*. Springer, Cham.
- Hobbes, Thomas (1994/1651). *Leviathan, with Selected Variants from the Latin Edition of 1668*. Hackett Publishing Company, Inc., Indianapolis/Cambridge.
- Hofbauer, Josef and Karl Sigmund (1998). *Evolutionary Game and Population Dynamics*. Cambridge University Press, Cambridge.
- Hofbauer, Josef and Karl Sigmund (2003). Evolutionary Game Dynamics. *Bulletin of the American Mathematical Society*, 40: 479–519.
- House of Lords, UK (2018). AI in the UK: Ready, willing and able? <https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/100.pdf>.



- Huang, Weini and Arne Traulsen (2010). Fixation Probabilities of Random Mutants under Frequency Dependent Selection. *Journal of Theoretical Biology*, 263(2): 262–268.
- Hume, David (1739). *A Treatise of Human Nature*. John Noon, London.
- Hurd, Peter L. (1995). Communication in Discrete Action-Response Games. *Journal of Theoretical Biology*, 174: 217–222.
- IBM (2017). Principles for the Cognitive Era. <https://www.ibm.com/blogs/think/2017/01/ibm-cognitive-principles/>.
- IBM (2018). Principles for Trust and Transparency. <https://www.ibm.com/blogs/policy/trust-principles/>.
- Imhof, Lorenz A. and Martin A. Nowak (2006). Evolutionary Game Dynamics in a Wright-Fisher Process. *Journal of Mathematical Biology*, 52(5): 667–681.
- Information Technology Industry Council (2017). AI Policy Principles. <https://www.itic.org/public-policy/ITIAIPolicyPrinciplesFINAL.pdf>.
- Jäger, Gerhard (2008). Evolutionary Stability Conditions for Signaling Games with Costly Signals. *Journal of Theoretical Biology*, 253: 131–141.
- Jobin, Anna, Marcello Ienca, and Effy Vayena (2019). The Global Landscape of AI Ethics Guidelines. *Nature: Machine Intelligence*, 1: 389–399.
- Johnstone, Rufus A. (1995). Sexual Selection, Honest Advertisement and the Handicap Principle: Reviewing the Evidence. *Biological Reviews*, 7: 1–65.
- Kameda, Tatsuya and Daisuke Nakanishi (2003). Does Social/Cultural Learning Increase Human Adaptability? Rogers’s question revisited. *Evolution and Human Behavior*, 24: 242–260.
- Kaniaru, Donald, Rajendra Shende, Scott Stone, and Durwood Zaelke (2007). Strengthening the Montréal Protocol: Insurance against Abrupt Climate Change. *Sustainable Development Law & Policy*, 7(2): 3–9, 74–76.
- Kendal, Jeremy, Marcus W. Feldman, and Kenichi Aoki (2006). Cultural Coevolution of Norm Adoption and Enforcement when Punishers are Rewarded or Non-punishers Are Punished. *Theoretical Population Biology*, 70: 10–25.
- Kraig, Michael R. (1999). Nuclear Deterrence in the Developing World: A Game-Theoretic Treatment. *Journal of Peace Research*, 36(2): 141–167.
- Kurokawa, Shun and Yasuo Ihara (2009). Emergence of cooperation in public goods games. *Proceedings of the Royal Society B: Biological Science*, 276(1660): 1379–1384.
- Lachmann, Michael, Szabolcs Szamado, and Carl T. Bergstrom (2001). Cost and Conflict in Animal Signals and Human Language. *Proceedings of the National Academy of Sciences*, 98(23): 13189–13194.
- LaCroix, Travis (2020). *Complex Signals: Reflexivity, Hierarchical Structure, and Modular Composition*. PhD thesis, University of California, Irvine.
- LaCroix, Travis (2022). Using logic to evolve more logic: Composing logical operators via self-assembly. *British Journal for the Philosophy of Science*, 73(2): 407–437.
- LaCroix, Travis and Yoshua Bengio (2019). Learning from Learning Machines: Optimisation, Rules, and Social Norms. *arXiv pre-print*, abs/2001.00006: 1–24. <https://arxiv.org/abs/2001.00006>.
- LaCroix, Travis and Cailin O’Connor (2020). Power by Association. *PhilSci Archive pre-print*, 14318: 1–26. Forthcoming in *Ergo*. <http://philsci-archive.pitt.edu/14318/>.

- Littman, Michael L. (1994). Markov Games as a Framework for Multi-agent Reinforcement Learning. *ICML'94: Proceedings of the Eleventh International Conference on International Conference on Machine Learning*, pages 157–163.
- Liu, Xuesong, Mingfeng He, Yibin Kang, and Qiuhui Pan (2017). Fixation of Strategies with the Moran and Fermi Processes in Evolutionary Games. *Physica A*, 484: 336–344.
- Liu, Xuesong, Qiuhui Pan, Yibin Kang, and Mingfeng He (2015). Fixation Probabilities in Evolutionary Games with the Moran and Fermi Processes. *Journal of Theoretical Biology*, 364: 242–248.
- Liu, Yongkui, Xiaojie Chen, Long Wang, Bin Li, Wenge Zhang, and Huifeng Wang (2011). Aspiration-based Learning Promotes Cooperation in Spatial Prisoner's Dilemma Games. *EPL (Europhysics Letters)*, 94(6): 60002.
- Lomas, Jonathan (1991). Words without Action? The Production, Dissemination, and Impact of Consensus Recommendations. *Annual Review of Public Health*, 12(1): 41–65.
- Lomas, Jonathan, Geoffrey M. Anderson, Karin Domnick-Pierre, Eugene Vayda, Murray W. Enkin, and Walter J. Hannah (1989). Do Practice Guidelines Guide Practice? *New England Journal of Medicine*, 321(19): 1306–1311.
- Luccioni, Alexandra and Yoshua Bengio (2019). On the Morality of Artificial Intelligence. *arXiv pre-print*, abs/1912.11945: 1–12. <http://arxiv.org/abs/1912.11945>.
- Madani, Kaveh (2010). Game Theory and Water Resources. *Journal of Hydrology*, 381(3–4): 225–238.
- Makridakis, Spyros (2017). The Forthcoming Artificial Intelligence (AI) Revolution: Its Impact on Society and Firms. *Futures*, 90: 46–60.
- Maynard Smith, John (1982). *Evolution and the Theory of Games*. Cambridge University Press, Cambridge.
- Maynard Smith, John and George R. Price (1964). Group Selection and Kin Selection. *Nature*, 201: 1145–1147.
- McNamara, Andrew, Justin Smith, and Emerson Murphy-Hill (2018). Does ACM's Code of Ethics Change Ethical Decision Making in Software Development? In Leavens, Gary T., Alessandro Garcia, , and Corina S. Păsăreanu, editors, *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering - ESEC/FSE 2018*, pages 1–7. ACM Press, New York.
- McPherson, Miller, Lynn Smith-Lovin, and James M. Cook (2001). Birds of a Feather: Homophily in Social Networks. *Annual Review of Sociology*, 27: 415–444.
- Miceli, Milagros, Julian Posada, and Tianling Yang (2022). Studying Up Machine Learning Data: Why Talk About Bias When We Mean Power? *arXiv pre-print*, 2109.08131: 1–14. <http://arxiv.org/abs/2109.08131>.
- Microsoft (2018). Microsoft AI Principles. <https://www.microsoft.com/en-us/ai/our-approach-to-ai>.
- Milinski, M., R. D. Sommerfeld, H. J. Krambeck, F. A. Reed, and J. Marotzke (2008). The Collective-risk Social Dilemma and the Prevention of Simulated Dangerous Climate Change. *Proceedings of the National Academy of Sciences of the United States of America*, 105(7): 2291–2294.

- Mittelstadt, Brent (2019). Principles alone cannot guarantee ethical AI. *Nature Machine Intelligence*, 1(11): 501–507.
- Mohseni, Aydin (2019). Stochastic Stability & Disagreement in Evolutionary Dynamics. *Philosophy of Science*, 86(3): 497–521.
- Morley, Jessica, Luciano Floridi, Libby Kinsey, and Anat Elhalal (2019). From What to How: An Initial Review of Publicly Available AI Ethics Tools, Methods and Research to Translate Principles into Practices. *arXiv pre-print*, abs/1905.06876: 1–28. <https://arxiv.org/abs/1905.06876>.
- Moyano, Luis G. and Angel Sánchez (2009). Evolving Learning Rules and Emergence of Cooperation in Spatial Prisoner’s Dilemma. *Journal of Theoretical Biology*, 259(1): 84–95.
- Nakahashi, Wataru (2007). The Evolution of Conformist Transmission in Social Learning when the Environment Changes Periodically. *Theoretical Population Biology*, 72: 52–66.
- Neumann, John Von and Oskar Morgenstern (2007/1944). *Theory of Games and Economic Behavior*. Princeton university press, Princeton.
- Nowak, Martin A. (2012). Evolving Cooperation. *Journal of Theoretical Biology*, 299: 1–8.
- Nowak, Martin A., Joshua B. Plotkin, and David C. Krakauer (1999). The Evolutionary Language Game. *Journal of Theoretical Biology*, 200: 147–162.
- Nowak, Martin A., Akira Sasaki, Christine Taylor, and Drew Fudenberg (2004). Emergence of Cooperation and Evolutionary Stability In Finite Populations. *Nature*, 428: 646–650.
- Nowak, Martin A. and Karl Sigmund (2004). Evolutionary Dynamics of Biological Games. *Science*, 303: 793–799.
- Ohtsuki, Hisashi, Pedro Bordaob, and Martin A. Nowak (2007). The One-third Law of Evolutionary Dynamics. *Journal of Theoretical Biology*, 249(2): 289–295.
- Ohtsuki, Hisashi and Martin A. Nowak (2006). Evolutionary Games on Cycles. *Proceedings of the Royal Society B: Biological Science*, 273(1598): 2249–2256.
- Ohtsuki, Hisashi and Martin A. Nowak (2008). Evolutionary Stability on Graphs. *Journal of Theoretical Biology*, 251: 698–707.
- OpenAI (2018). OpenAI Charter. <https://blog.openai.com/openai-charter/>.
- Ostrom, Elinor (2000). Collective Action and the Evolution of Social Norms. *Journal of Economic Perspectives*, 14(3): 137–158.
- Pacheco, Jorge M, Francisco C. Santos, Max O. Souza, and Brian Skyrms (2009). Evolutionary Dynamics of Collective Action in n-Person Stag Hunt Dilemmas. *Proceedings of the Royal Society B: Biological Sciences*, 276(1655): 315.
- Pacheco, Jorge M., Vítor V. Vasconcelos, and Francisco C. Santos (2014). Climate Change Governance, Cooperation and Self-organization. *Physics of Life Reviews*, 11(4): 573–586.
- Page, Karen M. and Martin A. Nowak (2002). Empathy leads to fairness. *Bulletin of Mathematical Biology*, 64: 1101–1116.
- Partnership on AI (2016). Tenets. <https://www.partnershiponai.org/tenets/>.
- Pawlowitsch, Christina (2007). Finite Populations Choose an Optimal Language. *Journal of Theoretical Biology*, 249: 606–616.
- Pawlowitsch, Christina (2008). Why Evolution Does Not Always Lead to an Optimal Signaling System. *Games and Economic Behavior*, 63(1): 203–226.

- Poisson, Siméon Denis (1837). *Recherches sur la probabilité des jugements en matière criminelle et en matière civile, précédées des règles générales du calcul des probabilités*. Bachelier, Paris.
- Pomiankowski, Andrew (1987). Sexual Selection: The Handicap Principle Does Work – Sometimes. *Proceedings of the Royal Society B: Biological Science*, 231: 123–145.
- Rand, David G. and Martin A. Nowak (2013). Human Cooperation. *Trends in Cognitive Science*, 17(8): 413–425.
- Rapoport, Anatol and Albert M. Chammah (1966). The Game of Chicken. *American Behavioral Scientist*, 10(3): 10–28.
- Rogers, Alan R. (1988). Does Biology Constrain Culture? *American Anthropologist*, 90: 819–831.
- Ross, Don (2019). Game Theory. In Zalta, Edward N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2019 edition.
- Royal Statistical Society and the Institute and Faculty of Actuaries (2019). A Guide for Ethical Data Science: A Collaboration between the Royal Statistical Society (RSS) and the Institute and Faculty of Actuaries (IFoA). <https://www.actuaries.org.uk/system/files/field/document/An%20Ethical%20Charter%20for%20Date%20Science%20WEB%20FINAL.PDF>.
- Russell, Stuart (2019). *Human Compatible: Artificial Intelligence and the Control Problem*. Viking, New York.
- Sage (2017). The Ethics of Code: Developing AI for Business with Five Core Principles. <https://www.sage.com/ca/our-news/press-releases/2017/06/designing-AI-for-business>.
- Sánchez, Angel and José A. Cuesta (2005). Altruism May Arise from Individual Selection. *Journal of Theoretical Biology*, 235: 233–240.
- Sandholm, William H. (2007). Simple Formulas for Stationary Distributions and Stochastically Stable States. *Games and Economic Behavior*, 59(1): 154–162.
- Santos, Francisco C. and Jorge M. Pacheco (2011). Risk of Collective Failure Provides an Escape from the Tragedy of the Commons. *Proceedings of the National Academy of Sciences of the United States of America*, 108(26): 10421–10425.
- SAP (2018). Sap’s guiding principles for artificial intelligence. <https://news.sap.com/2018/09/sap-guiding-principles-for-artificial-intelligence/>.
- Serrano, Roberto and Allan M. Feldman (2013). *A Short Course in Intermediate Microeconomics with Calculus*. Cambridge University Press, Cambridge.
- Shapley, Lloyd S. (1953). Stochastic Games. *Proceedings of the National Academy of Sciences of the United States of America*, 39: 1095–1100.
- Sigmund, Karl (2010). *The Calculus of Selfishness*. Cambridge University Press, Cambridge.
- Skyrms, Brian (1994). Darwin Meets the Logic of Decision: Correlation in Evolutionary Game Theory. *Philosophy of Science*, 61: 503–528.
- Skyrms, Brian (2004). *The Stag Hunt and the Evolution of Social Structure*. Cambridge University Press, Cambridge.
- Skyrms, Brian (2010). *Signals: Evolution, Learning, & Information*. Oxford University Press, Oxford.
- Skyrms, Brian (2014/1996). *Evolution of the Social Contract*. Cambridge University Press, Cambridge.

- Sony (2018). Sony group AI ethics guidelines. [https://www.sony.net/SonyInfo/csr\\_report/humanrights/hkrfmg0000007rtj-att/AI\\_Engagement\\_within\\_Sony\\_Group.pdf](https://www.sony.net/SonyInfo/csr_report/humanrights/hkrfmg0000007rtj-att/AI_Engagement_within_Sony_Group.pdf).
- Sossin, Lorne and Charles W. Smith (2003). Hard Choices and Soft Law: Ethical Codes, Policy Guidelines and the Role of the Courts in Regulating Government. *Alberta Law Review*, 40: 867–893.
- Souza, Max O., Jorge M. Pacheco, and Francisco C. Santos (2009). Evolution of Cooperation under N-person Snowdrift Games. *Journal of Theoretical Biology*, 260(4): 581–588.
- Stanford University (2018). The Stanford Human-Centered AI Initiative (HAI). <http://hai.stanford.edu/news/introducing-stanfords-human-centered-ai-initiative/>.
- Szabo, Gyorgy, Attila Szolnoki, and Jeromos Vukov (2009). Selection of Dynamical Rules in Spatial Prisoner’s Dilemma Games. *EPL (Europhysics Letters)*, 87(1): 18007.
- Szolnoki, Attila, Jeromos Vukov, and Gyorgy Szabo (2009). Selection of Noise Level in Strategy Adoption for Spatial Social Dilemmas. *Physical Review E: Statistical, Nonlinear, Biological, and Soft Matter Physics*, 80(5 pt. 2): 056112.
- Taylor, Christine, Drew Fudenberg, Akira Sasaki, and Martin A. Nowak (2004). Evolutionary Game Dynamics in Finite Populations. *Bulletin of Mathematical Biology*, 66(6): 1621–1644.
- Taylor, Christine, Yoh Iwasa, and Martin A. Nowak (2006). A Symmetry of Fixation Times in Evolutionary Dynamics. *Journal of Theoretical Biology*, 243(2): 245–245.
- Taylor, Peter D. and Leo B. Jonker (1978). Evolutionarily Stable Strategies and Game Dynamics. *Mathematical Biosciences*, 40: 145–156.
- The Future Society (2017). Principles for the Governance of AI. <http://www.thefuturesociety.org/science-law-society-sls-initiative/#1516790384127-3ea0ef44-2aae>.
- The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems (2017). Ethically Aligned Design, Version 2. [http://standards.ieee.org/develop/indconn/ec/autonomous\\_systems.html](http://standards.ieee.org/develop/indconn/ec/autonomous_systems.html).
- The Japanese Society for Artificial Intelligence (2017). The Japanese Society for Artificial Intelligence Ethical Guidelines. <http://ai-elsi.org/wp-content/uploads/2017/05/JSAI-Ethical-Guidelines-1.pdf>.
- The Public Voice (2018). Universal Guidelines for Artificial Intelligence. <https://thepublicvoice.org/ai-universal-guidelines/>.
- Traulsen, Arne and Christoph Hauert (2009). Stochastic Evolutionary Game Dynamics. In Schuster, H. G., editor, *Reviews of Nonlinear Dynamics and Complexity*, volume 2, pages 25–62. Wiley-VCH, Weinheim.
- Traulsen, Arne, Martin A. Nowak, and Jorge M. Pacheco (2006a). Stochastic Dynamics of Invasion and Fixation. *Physical Review E: Statistical, Nonlinear, Biological, and Soft Matter Physics*, 74(1 pt. 1): 011909.
- Traulsen, Arne, Jorge M. Pacheco, and Lorens Imhof (2006b). Stochasticity and Evolutionary Stability. *Physical Review E: Statistical, Nonlinear, Biological, and Soft Matter Physics*, 74(2 pt. 1): 021905.
- Traulsen, Arne, Jorge M. Pacheco, and Martin A. Nowak (2007). Pairwise Comparison and Selection Temperature in Evolutionary Game Dynamics. *Journal of*

- Theoretical Biology*, 246(3): 522–529.
- Traulsen, Arne, Dirk Semmann, Ralf D. Sommerfeld, Hans-Jürgen Krambeck, and Manfred Milinski (2009). Human Strategy Updating in Evolutionary Games. *Proceedings of the National Academy of Sciences of the United States of America*, 107(7): 2962–2966.
- Trivers, Robert L. (1971). The Evolution of Reciprocal Altruism. *The Quarterly Review of Biology*, 46(3): 35–57.
- UNESCO (2020). Composition of the Ad Hoc Expert Group (AHEG) for the Recommendation on the Ethics of Artificial Intelligence / Composition du Groupe d’experts ad hoc (GEAH) pour la Recommandation sur l’éthique de l’intelligence artificielle. *United Nations Educational, Scientific, and Cultural Organization*, 0000372991: 1–8.
- UNI Global Union (2017). Top 10 Principles For Ethical Artificial Intelligence. [http://www.thefutureworldofwork.org/media/35420/uni\\_ethical\\_ai.pdf](http://www.thefutureworldofwork.org/media/35420/uni_ethical_ai.pdf).
- Université de Montréal (2017). The Montreal Declaration for a Responsible Development of Artificial Intelligence. <https://www.montrealdeclaration-responsibleai.com/the-declaration>.
- US Public Policy Council, Association for Computing Machinery (2017). Principles for Algorithmic Transparency and Accountability. [https://www.acm.org/binaries/content/assets/public-policy/2017\\_usacm\\_statement\\_algorithms.pdf](https://www.acm.org/binaries/content/assets/public-policy/2017_usacm_statement_algorithms.pdf).
- Wagner, Ben (2018). Ethics as an escape from regulation: From ‘ethics-washing’ to ethics-shopping? In Bayamlioglu, E., I. Baraliuc, L. A. W. Janssens, and M. Hildebrandt, editors, *Being Profiled: Cogitas Ergo Sum: 10 Years of Profiling the European Citizen*, pages 84–89. Amsterdam University Press, Amsterdam.
- Wagner, Ulrich J. (2001). The Design of Stable International Environmental Agreements: Economic Theory and Political Economy. *Journal of Economic Surveys*, 15(3): 377–411.
- Wakano, Joe Yuichiro and Kenichi Aoki (2006). A Mixed Strategy Model for the Emergence and Intensification of Social Learning in a Periodically Changing Natural Environment. *Theoretical Population Biology*, 70: 486–497.
- Wakano, Joe Yuichiro, Kenichi Aoki, and Marcus W. Feldman (2004). Evolution of Social Learning: A Mathematical Analysis. *Theoretical Population Biology*, 66: 249–258.
- Wang, Jing, Feng Fu, Te Wu, and Long Wang (2009). Emergence of Social Cooperation in Threshold Public Goods Games with Collective Risk. *Physical Review E: Statistical, Nonlinear, Biological, and Soft Matter Physics*, 80: 016101.
- Wärneryd, Karl (1993). Cheap Talk, Coordination and Evolutionary Stability. *Games and Economic Behavior*, 5(4): 532–546.
- Weibull, Jörgen M. (1997). *Evolutionary Game Theory*. The MIT Press, Cambridge, MA.
- Whittaker, Meredith, Kate Crawford, Roel Dobbe, Genevieve Fried, Elizabeth Kazianas, Varoon Mathur, Sarah Mysers West, Rashida Richardson, Jason Schultz, and Oscar Schwartz (2018). *AI now report 2018*. AI Now Institute at New York University, New York. [https://ainowinstitute.org/AI\\_Now\\_2018\\_Report.pdf](https://ainowinstitute.org/AI_Now_2018_Report.pdf).
- Whittlestone, Jess, Rune Nyrop, Anna Alexandrova, and Stephen Cave (2019). The Role and Limits of Principles in AI Ethics: Towards a Focus on Tensions. In



- Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '19, page 195–200, New York. Association for Computing Machinery (ACM).
- Wu, Bin, Philipp M. Altrock, Long Wang, and Arne Traulsen (2010). Universality of Weak Selection. *Physical Review E: Statistical, Nonlinear, Biological, and Soft Matter Physics*, 82: 046106.
- Wu, Bin, Benedikt Bauer, Tobias Galla, and Arne Traulsen (2015). Fitness-based Models and Pairwise Comparison Models of Evolutionary Games are Typically Different—Even in Unstructured Populations. *New Journal of Physics*, 17: 023043.
- Young, H. Peyton and Shmuel Zamir (2014). *Handbook of Game Theory*, volume 4. Elsevier, North-Holland.
- Zagare, Frank C. (1987). *The Dynamics of Deterrence*. University of Chicago Press, Chicago.
- Zahavi, Amotz (1975). Mate Selection: A Selection for a Handicap. *Journal of Theoretical Biology*, 53(1): 205–214.
- Zahavi, Amotz and Avishag Zahavi (1997). *The Handicap Principle*. Oxford University Press, New York.
- Zhang, Kaiqing, Zhuoran Yang, and Tamer Başar (2019). Multi-Agent Reinforcement Learning: A Selective Overview of Theories and Algorithms. *arXiv pre-print*, abs/1911.10635. <https://arxiv.org/abs/1911.10635>.
- Zollman, Kevin J. S. (2005). Talking to Neighbors: The Evolution of Regional Meaning. *Philosophy of Science*, 72(1): 69–85.

## APPENDIX A. GAME THEORY

In this brief appendix, we provide some further game-theoretic background than we had space to discuss in Section 3. For more comprehensive introductions to game theory, see, e.g., Aumann and Hart (1992, 1994, 2002); Maynard Smith (1982); Neumann and Morgenstern (1944); Weibull (1997); Young and Zamir (2014).

**Game-Theoretic Analysis of Cooperation and Conflict.** Cooperative behaviour persists in human and non-human animal populations alike, but it provides something of an evolutionary puzzle (Axelrod and Hamilton, 1981; Darwin, 1871; Hauert et al., 2006; Moyano and Sánchez, 2009; Nowak, 2012; Nowak et al., 2004; Taylor et al., 2004): How can cooperation be maintained despite incentives for non-cooperative behaviour (i.e., defection)? Evolutionary game theory provides useful tools for analysing the evolution of cooperative behaviour quantitatively in both human and non-human animals.<sup>22</sup>

Game theory can be used to study the ways in which independent choices between actors interact to produce outcomes.<sup>23</sup> In game theory, a game is determined by

<sup>22</sup>See, for example, Ashcroft et al. (2014); Fehel et al. (2011); Gintis (2000); Grujić et al. (2015, 2012); Hofbauer and Sigmund (1998, 2003); Imhof and Nowak (2006); Kurokawa and Ihara (2009); Maynard Smith (1982); Nowak and Sigmund (2004); Ohtsuki and Nowak (2006, 2008); Rand and Nowak (2013); Traulsen et al. (2009); Weibull (1997).

<sup>23</sup>See Ross (2019) for a philosophical overview.

the *payoffs*. For example, the payoff matrix for a generic,  $2 \times 2$ , symmetric, normal form game is displayed in Figure 5.

		Player 2	
		<i>C</i>	<i>D</i>
Player 1	<i>C</i>	$(a, a)$	$(b, c)$
	<i>D</i>	$(c, b)$	$(d, d)$

FIGURE 5. Payoff matrix for a generic,  $2 \times 2$ , symmetric, normal form game.

Each actor (Player 1 and Player 2) in this example can choose one of two strategies, *C* or *D*. The payoffs to each of the players are given by the respective entries in each cell—i.e., the first number in the top-right cell (*b*) is the payoff afforded to Player 1 when she plays *C* and her partner plays *D*; the second number (*c*) is the payoff afforded to Player 2 in the same situation (i.e., when Player 2 plays *D* and Player 1 plays *C*).

As discussed in the paper, social dilemmas are games where (i) the payoff to each individual for non-cooperative behaviour is higher than the payoff for cooperative behaviour, and (ii) every individual receives a lower payoff when everyone defects than they would have, had everyone cooperated (Dawes, 1980).

When  $c > a > d > b$ , in Figure 5, we have a *Prisoner's Dilemma*.<sup>24</sup> Note that when both actors cooperate (i.e., both play *C*), their payoff is higher than if they both defect ( $a > d$ ), thus satisfying criterion (ii) mentioned above. However, each actor has an individual incentive to defect (i.e., play *D*) *regardless* of what the other actor does; Player 1 would prefer to defect when Player 2 cooperates ( $c > a$ ), and she would prefer to defect when Player 2 defects ( $d > b$ )—and *mutatis mutandis* for Player 2. This satisfies criterion (i) above.

In this case, we say that *defect* is a *strictly dominant* strategy for each player, which leads to the unique *Nash equilibrium*:  $\langle D, D \rangle$ —that is, a combination of strategies where no actor can increase her payoff by unilateral deviation from her strategy. The ‘dilemma’ is that mutual cooperation yields a better outcome for all parties than mutual defection, but, from an individual perspective, it is never rational to cooperate.

**Evolutionary Game Dynamics.** In an evolutionary context, the payoffs are identified with *reproductive fitness*, so that more-successful strategies are more likely

<sup>24</sup>Named and formalised by Canadian mathematician Albert W. Tucker in 1952, based on Merrill M. Flood and Melvin Dresher’s 1950 model; see Serrano and Feldman (2013).



to propagate, reproduce, be replicated, be imitated, etc. This provides a natural way to incorporate *dynamics* to the underlying game.

There are two natural interpretations of evolutionary game dynamics. The first is *biological*, where strategies are encoded in the genome of individuals, and those who are successful pass on their genes at higher rates; the second is *cultural*, where successful behaviours are reproduced through learning and imitation. We are primarily concerned with processes of cultural evolution. This process should be familiar to those in AI/ML who work on multi-agent reinforcement learning (MARL) (Littman, 1994; Shapley, 1953; Zhang et al., 2019).

In addition to the game, an evolutionary model requires a specification of the *dynamics*—namely, a set of rules for determining how the strategies of actors in a population will update (under a cultural interpretation), or how the proportions of strategies being played in the population will shift as they proliferate or are driven to extinction (under a biological interpretation). Evolutionary game dynamics are often studied in infinite populations using deterministic differential equations. For example, the *replicator dynamic* (Taylor and Jonker, 1978) captures how strategies with higher-than-average fitness tend to increase, and strategies with lower-than-average fitness tend to decrease. A population state is *evolutionarily stable* only if it is an asymptotically stable rest point of the dynamics (Maynard Smith, 1982).

**Stochastic Game Dynamics.** In finite populations, *stochastic* game dynamics are used to study the selection of traits with frequency-dependent fitness (Liu et al., 2011; Ohtsuki et al., 2007; Sigmund, 2010; Szabo et al., 2009; Szolnoki et al., 2009; Traulsen et al., 2006a).

A standard stochastic game dynamics that is used extensively is the *Moran process*. This is a simple birth-death process where an individual is chosen proportional to their fitness and replaces a randomly chosen individual with an offspring of its own type (Altrock and Traulsen, 2009; Claussen and Traulsen, 2005; Huang and Traulsen, 2010; Liu et al., 2017, 2015; Taylor et al., 2006; Traulsen et al., 2007; Wu et al., 2010, 2015).

In the standard *Fermi process*, which we discuss in Section 3, an individual is chosen randomly from a finite population, and its reproductive success is evaluated by comparing its payoff to a second, randomly-selected individual from the population (Liu et al., 2017; Traulsen et al., 2006b, 2007).

As mentioned in Section 3, the pairwise comparison of payoffs of the focal individual and the role model informs the probability,  $p$ , that the focal individual copies the strategy of the role model; the probability function, called the *Fermi* function, was presented in Equation 3, and repeated here for convenience:

$$p = \left[ 1 + e^{\lambda(\pi_f - \pi_r)} \right]^{-1}.$$

Again, if both individuals have the same payoff, the focal individual randomises between the two strategies. Note, then, that the focal individual does not always switch to a better strategy; the individual may switch to one that is strictly worse.

When the intensity of selection  $\lambda = 0$ , selection is random; when selection is *weak* ( $\lambda \ll 1$ ),  $p$  reduces to a linear function of the payoff difference; when  $\lambda = 1$ , our model gives us back the *replicator dynamic*; and, when  $\lambda \rightarrow \infty$ , we get the *best-response* dynamic (Fudenberg and Tirole, 1991).

Evolutionary game dynamics have been used to shed light upon many aspects of human behaviour, including altruism,<sup>25</sup> moral behaviour,<sup>26</sup> empathy,<sup>27</sup> social learning,<sup>28</sup> social norms,<sup>29</sup> and the evolution of communication, proto-language, and compositional syntax,<sup>30</sup> among many others. See Ross (2019) for further details.

## APPENDIX B. TECHNICAL DETAILS

In this brief appendix, we provide some further formal details for our model than we had space to discuss in Section 3.

**Mean Payoffs.** Recall that the payoffs to each cooperators,  $C$ , and defectors,  $D$ , in a group of size  $N$  are given as a function of the number of cooperators in that group,  $n_C$ , as follows:

$$\begin{aligned}\pi_C(n_C) &= b \cdot \Theta(n_C - n^*) + b \cdot (1 - rm) \cdot \Theta(n_C - n^*) - cb, \\ \pi_D(n_C) &= \pi_C(n_C) + cb,\end{aligned}$$

where  $\Theta$  is the Heaviside step function. The mean payoffs to each type in a population of size  $Z$ , where groups are determined by random mixing, is then given as a function of the total fraction of cooperators in the population,  $x_C = n_C^Z/Z$ , as

<sup>25</sup>See, e.g., Fletcher and Zwick (2007); Gintis et al. (2003); Sánchez and Cuesta (2005); Trivers (1971).

<sup>26</sup>See, e.g., Alexander (2007); Boehm (1982); Harms and Skyrms (2008); Skyrms (2004, 1996).

<sup>27</sup>See, e.g., Fishman (2006); Page and Nowak (2002).

<sup>28</sup>See, e.g., Kameda and Nakanishi (2003); Nakahashi (2007); Rogers (1988); Wakano and Aoki (2006); Wakano et al. (2004).

<sup>29</sup>See, e.g., Axelrod (1981); Bicchieri (2006); Binmore and Samuelson (1994); Chalub et al. (2006); Kendal et al. (2006); LaCroix and O'Connor (2020); Ostrom (2000).

<sup>30</sup>See, e.g., Barrett (2007); Hausken and Hirshleifer (2008); Hurd (1995); Jäger (2008); LaCroix (2022, 2020); Nowak et al. (1999); Pawlowitsch (2007, 2008); Skyrms (2010); Zollman (2005).

follows:

$$\begin{aligned}\Pi_C(x_C) &= \sum_{n_C=0}^N \frac{n_C}{N} \binom{N}{n_C} x_C^{n_C} (1-x_C)^{N-n_C} \pi_C(n_C), \\ \Pi_D(x_C) &= \sum_{n_C=0}^N \frac{N-n_C}{N} \binom{N}{n_C} x_C^{n_C} (1-x_C)^{N-n_C} \pi_D(n_C).\end{aligned}$$

**Fermi Dynamics.** The Fermi dynamics uses the average payoffs to each type to determine the probability that a randomly-chosen individual from the population will imitate the strategy of a second randomly-chosen individual from the population. Such a change will produce one of three outcomes: the number of cooperators in the population,  $k = n_C^Z$ , will increase, decrease, or remain the same. This is captured by the following transition probabilities, which yield a tri-diagonal transition matrix,  $T$ , for our birth-death process:

$$\begin{aligned}T^+(k) &= (1-\mu) \frac{k}{Z} \frac{Z-k}{Z-1} \left(1 + e^{\lambda(\Pi_C - \Pi_D)}\right)^{-1} + \frac{\mu}{2} \\ T^-(k) &= (1-\mu) \frac{Z-k}{Z} \frac{k}{Z-1} \left(1 + e^{\lambda(\Pi_D - \Pi_C)}\right)^{-1} + \frac{\mu}{2} \\ T^0(k) &= 1 - T^+(k) - T^-(k)\end{aligned}$$

where  $\lambda$  is the inverse temperature associated with the influence of selection versus drift, and  $\mu$  is the rate of mutation. This produces an ergodic Markov process.

**Gradient of Selection.** The gradient of selection of the process captures the expected direction of selection as a function of the number of cooperators in the population,  $k$ , in a way that is analogous to the mean-field dynamics for the infinite-population case. This is given by

$$G(k) = T^+(k) - T^-(k) = \frac{k}{Z} \frac{Z-k}{Z-1} \left( \tanh \frac{\lambda}{2} (\Pi_C(k) - \Pi_D(k)) \right),$$

where  $G(k) > 0$  implies that selection favours cooperation, and  $G(k) < 0$  implies that defection is favoured.

**Stationary Distribution.** The stationary distribution of the process captures the long run distribution of time the process spends at each state. For an ergodic process, the stationary distribution is known to be unique and independent of initial conditions of that process. We compute it as follows.

$$\sigma_k = \frac{\prod_{i=1}^k \frac{T^+(i-1)}{T^-(i)}}{\sum_{i=1}^Z \prod_{j=1}^i \frac{T^+(j-1)}{T^-(j)}} \quad \text{for } k \in \{1, \dots, Z\}.$$

The stationary distribution can also be approximated via the Chapman-Kolmogorov equation which states that  $n$ th step transition matrix corresponds to the  $n$ th power

of the one-step transition matrix,  $T_t = T^t$ . Thus, we get that  $\sigma$  corresponds to any row of the matrix given by  $\lim_{t \rightarrow \infty} T^t$ .

### APPENDIX C. PROOFS

Here we demonstrate several propositions which elucidate the general relationship between selection for cooperation under the Fermi dynamics and the parameters of the strategic interaction.

We say that *selection for* a strategy,  $\sigma$ , under the dynamics is increasing in parameter  $x$  if the transition probability  $T^+(k)$  from a state with  $k$  individuals playing strategy  $\sigma$  to one with  $k + 1$  individuals playing  $\sigma$  increases as  $x$  increases, for every interior state. That is  $x < x'$  implies  $T^+(k; x) < T^+(k; x')$  for all  $k \in \{1, \dots, Z - 1\}$ .

For the following proofs, we fix the initial endowment of agents as some positive constant,  $b > 0$ , without loss of generality, and we assume non-extreme values of the strategic parameters of interest:  $N \leq Z \in \mathbb{N}$ ;  $r \in (0, 1)$ ;  $m \in (0, 1)$ ;  $c \in (0, 1)$ ;  $p^* \in (\frac{1}{N}, \frac{N-1}{N})$ ;  $\mu \in (0, 1)$ ; and  $\lambda \in \mathbb{R}_{>0}$ . Note that allowing for extreme values of the parameters makes it so the following inequalities hold only weakly.

**Lemma C.1.** *Selection for a cooperation under the Fermi dynamics increases (decreases) as the differences of its mean payoff with that of defection increases (decreases).*

*Proof.* Recall that the transition probability from a state with  $k$  cooperators to one with  $k + 1$  cooperators is given by

$$T^+(k) = (1 - \mu) \frac{k}{Z} \frac{Z - k}{Z - 1} \left(1 + e^{\lambda(\Pi_C - \Pi_D)}\right)^{-1} + \frac{\mu}{2}.$$

So, for any (non-extremal) values of  $k$ ,  $\lambda$ , and  $\mu$ , we have it that  $T^+$  is proportional to the logit function,  $(1 + e^{\lambda(\Pi_C - \Pi_D)})^{-1}$ , which in turn clearly increases (decreases) as the difference of mean payoffs,  $\Pi_C - \Pi_D$ , increases (decreases).  $\square$

**Proposition C.2.** *Selection for cooperation decreases as the cost to cooperation increases.*

*Proof.* Consider the difference in mean payoffs between cooperators and defectors,  $\Pi_C(k; c) - \Pi_D(k; c)$ , as a function of the cost of cooperation,  $c \in (0, 1)$ , and then fix each of the other parameters at some non-extremal values.

Observe that for any fixed number of cooperators,  $k \in \{1, \dots, Z-1\}$ , the difference in mean cost of cooperation

$$\Pi_C(k; c) - \Pi_D(k; c) \propto \sum_{n_C=0}^N \pi_C(n_C; c) - \pi_D(n_C; c) \propto -cb.$$

Since  $b > 0$ , increasing cost of cooperation,  $c$ , decreases  $\Pi_C - \Pi_D$ , as required.

By lemma C.1, it follows that selection for cooperation decreases as the cost of cooperation increases.  $\square$

**Proposition C.3.** *Selection for cooperation decreases as the size of cooperative groups increases.*

*Proof.* Consider the differences in mean payoffs between cooperators and defectors,  $\Pi_C(k; N) - \Pi_D(k; N)$ , as a function of the size of cooperative groups,  $N \in \mathbb{N}$ , and fix each other parameter at some non-extremal value.

Reformulate the difference in mean payoffs in terms of the fraction of the expected fractions of each cooperators and defectors in successful cooperative groups,  $p, q$  respectively:

$$\Pi_C - \Pi_D = (p\Pi_{C;s} + (1-p)\Pi_{C;f}) - (q\Pi_{D;s} + (1-q)\Pi_{D;f}),$$

where the subscripts  $s$  and  $f$  denote when the payoff is for success or failure. We pair the payoffs terms to get

$$(p\Pi_{C;s} - (q\Pi_{D;s}) - ((1-p)\Pi_{C;f}) - (1-q)\Pi_{D;f}),$$

and then use the fact that  $\Pi_{D;s} = \Pi_{C;s} + cb$  and  $\Pi_{D;f} = \Pi_{C;f} + cb$ , and some algebra, to simplify the expression to:  $(p-q)\Pi_{C;s} + (q-p)\Pi_{C;f} - cb$ .

Since the payoff for success is greater than failure,  $\Pi_{C;s} > \Pi_{C;f}$ , it follows that the difference in average payoffs is decreasing in the difference of the fraction  $p - q$  of successful cooperators and defectors. Taking the derivative of the difference of fractions with respect to  $N$  yields

$$\begin{aligned} & \frac{d}{dN}[p(N) - q(N)] \\ &= \frac{d}{dN} \left[ \sum_{n_C=\lceil p^* N \rceil}^N \binom{N}{n_C} (k/Z)^{n_C} (1-k/Z)^{N-n_C} \frac{2n_C - N}{N} \right], \\ &\propto - \sum_{n_C=\lceil p^* N \rceil}^N N^{-2} \end{aligned}$$

which is strictly negative.

Hence the difference in the fractions of cooperators and defectors who succeed and fail is decreasing in group size, and so the difference in mean payoffs between cooperators and defectors is decreasing. By Lemma C.1, it follows that selection for cooperation decreases as group size increases.  $\square$

**Proposition C.4.** *Selection for cooperation increases as the product of the perceived risk and magnitude of the consequences of failing to successfully cooperate increases.*

*Proof.* Consider the differences in mean payoffs to between cooperators and defectors,

$$\Pi_C(k; r, m) - \Pi_D(k; r, m),$$

as a function of the the product of the perceived probability,  $0 < r < 1$ , and magnitude,  $0 < m < 1$ , such that  $0 < rm < r'm' < 1$ .

Observe that

$$\begin{aligned} \Pi_C(k; r, m) - \Pi_D(k; r, m) &\propto \sum_{n_C=0}^N \pi_C(n_C; r, m) - \pi_D(n_C; r, m) \\ &= \pi_C(0; r, m) - \pi_D(0; r, m) + (N-1)cb + \pi_C(N; r, m) - \pi_D(N; r, m) \\ &\propto \pi_D(N; r, m) - \pi_C(0; r, m). \end{aligned}$$

When  $c_N = N$ , we have  $\pi_C(N; r, m) = b(1-c)$  (all cooperators; cooperation succeeds) and when  $c_N = 0$ , we have  $\pi_C(0; r, m) = b(1-rm-c)$  (all defectors; cooperation fails). Thus

$$\pi_D(N; r, m) - \pi_C(0; r, m) = brm,$$

and since  $b > 0$ , this is increasing in  $rm$ .

Hence the difference in mean payoffs,  $\Pi_C(k; r, m) - \Pi_D(k; r, m)$ , is also increasing in  $rm$ . By lemma C.1, it follows that selection for cooperation increases as the product of the risk and magnitude of the failure to cooperate increases.  $\square$