

Game-Theoretic Models of Moral and Other-Regarding Agents (extended abstract)

Gabriel Istrate

West University of Timișoara

gabrielistrate@acm.org

We investigate Kantian equilibria in finite normal form games, a class of non-Nashian, morally motivated courses of action that was recently proposed in the economics literature. We highlight a number of problems with such equilibria, including computational intractability, a high price of miscoordination, and problematic extension to general normal form games. We give such a generalization based on concept of *program equilibria*, and point out that that a practically relevant generalization may not exist. To remedy this we propose some general, intuitive, computationally tractable, other-regarding equilibria that are special cases Kantian equilibria, as well as a class of courses of action that interpolates between purely self-regarding and Kantian behavior.

1 Introduction

Game Theory is widely regarded as the main conceptual foundation of strategic behavior. The promise behind its explosive development (at the crossroads of Economics and Computer Science) is that of understanding the dynamics of human agents and societies and, equally importantly, of guiding the engineering of artificial agents, ultimately capable of realistic, human-like, courses of action. Yet, it is clear that the main models of Game Theory, primarily based on the self-interested, rational actor model, and exemplified by the concept of Nash equilibria, are not realistic representations of the richness of human interactions. Concepts such as *bounded rationality* [63], and the limitations they impose on the computational complexity of agents' cognitive models [56] can certainly account for some of this difference. But this is hardly the only possible explanation: People behave differently from ideal economic agents not because they would be irrational [2], but since many human interactions are cooperative, rather than competitive [68], guided by social norms such as *reciprocity*, *fairness* and *inequity-aversion* [13], often involving *networked minds*, rather than utility maximization performed in isolation [27], driven by moral considerations [69] or by other not purely self-regarding behaviors, e.g. *altruism* [36] and *spite* [17, 16].

Moral considerations (should) interact substantially with game theory: indeed, the latter field has been used to propose a reconstruction of moral philosophy [9, 8, 10]; conversely, some philosophers have gone as far as to claim that we need a *moral equilibrium theory* [65]. Whether that's true or not, it is a fact that *homo economicus*, the Nash optimizer of economics, is increasingly complemented by a rich emerging typology of human behavior [28], that also contains (in Gintis's words) "*homo socialis*, the other-regarding agent who cares about fairness, reciprocity, and the well-being of others, and *homo moralis*"¹ ... the Aristotelian bearer of nonconsequentialist character virtues".² These claims are well-documented experimentally: for instance, Fischbacher et al. [23] investigated the percent of people

¹since our agents are not necessarily human, we will use alternate names such as "moral agent" for this type of behavior.

²Gintis proposes a taxonomy of behavior with three distinct types of preferences: *self-regarding*, *other regarding* and *universalist*; a further relevant distinction is between so-called *private* and *public personas*, that leads to further types of behavior such as *homo Parochialis*, *homo Universalis* and *homo Vertus*. See [28] for further details.

having self-regarding preferences in a public goods game, showing that it is in the range of 30-40%, while the remaining were either other-regarding or moral agents. Since artificial agents (will) interact with humans, such concerns are highly relevant to the design of multiagent systems and justify the study of alternative other-regarding notions, e.g. Rong and Halpern's [32, 55] "cooperative equilibria" or *dependency theory* [62, 30]. Other-regarding considerations could be encoded (e.g. [22]) as *externalities* into agents' perceived utilities, that may lead them away from straightforward maximization of their material payoffs. However, keeping them explicit may be important for agent implementations.

The purpose of this paper is to contribute to the emerging literature on non-Nashian, morally inspired game theoretic concepts and, equally important, to bring its concerns and methods to the attention of the various communities represented in TARK. We are inspired by what we believe is one of the most intriguing classes of equilibrium concepts that can be seen as morally grounded: *Kantian (a.k.a. Hofstadter) equilibria* [54]. This notion emerged from three separate lines of research converging on an identical mathematical definition, but justifying it, however, from several very different perspectives: *superrationality* [38, 25], *team reasoning* [4], and *Kantian optimization*, respectively [54].

The common framework (most crisply developed for symmetric coordination games) only considers as relevant the action profiles where all agents choose *the same action*, choosing the action x that, if played by everyone, maximizes agents' (identical) utility functions. The justification of this restriction depends on the perspective: *superrationality* assumes that if rationality constrains an agent to choose a specific course of action x , then the same reasoning compels **all** agents (at least in the case of symmetric games, when all agents are positionally indistinguishable from the original agent) to also choose x .³ In contrast, *Kantian optimization* justifies the limitation to symmetric profiles in a very different manner: Roemer [52] suggested that agents often ignore the potential for action of the other players, acting instead according to the *Kantian categorical imperative* [58] "act only according to that maxim whereby you can, at the same time, will that it should become a universal law", that is, choose a course of action that, if adopted by every agent, would bring all agents the highest payoff.⁴ One way to formalize this idea, employed e.g. in [1], is to decouple the *material payoffs* agents receive from their (perceived) utility, which agents maximize in order to select the action. Specifically, assume the given agent i plays strategy x against action profile y . We assume that the material payoff the agent receives is $\pi_i(x, y)$. On the other hand the utility the agent uses to evaluate alternative x may not be equal to $\pi_i(x, y)$ and may in fact, have in fact nothing to do with y at all! Instead, $u_i(x, y) = \pi_i(x, \bar{x}_{-i})$, where \bar{x}_{-i} is the action profile where all agents other than i play x as well. That is, the agent evaluates the desirability of action x in isolation from the actions of the other players, as if choosing x could somehow "magically" determine the other players to adopt the same strategy.⁵ Alternatively, $\pi(x, \bar{x}_{-i})$ measures the extent to which action x is "the morally best course of action". **Such a justification is cognitively plausible:** experiments have shown [45] that people often employ such "universalization" arguments when judging the morality of a given behavior.

³To cite Hofstadter: "If reasoning dictates an answer, then everyone should independently come to that answer. Seeing this fact is itself the critical step in the reasoning toward the correct answer [...]". Though superrationality does away with the assumption of counterfactual independence of Nash equilibria, it is otherwise compatible with a particular version of homo economicus that requires some very strong assumptions on agent rationality (see [25] for a discussion).

⁴As recognized by Roemer himself and discussed e.g. in [14], the connection of Kantian equilibria to actual Kantian ideas is quite loose. Another possible interpretation is that Kantian equilibria embody *rule utilitarianism* [35]. Finally, see [60] for a discussion of the normative aspects of Kantian equilibria.

⁵Frank [26] refers to this as *voodoo causation*. Elster [21] argues that Kantian optimization seems to be rooted in a form of *magical thinking*, "causing agents to act on the belief (or act as if they believed) that they can have a causal influence on outcomes that are effectively outside their control". We take a descriptive, rather than normative position: such reasoning is something people **simply do**; understanding its implications is strategically valuable.

The questions we attempt to start answering in this paper are:

1. *Can we extend the definition of Kantian equilibria to cover all natural cases of "Kantian behavior"?*

However, we are **not** simply looking in our generalization for yet another equilibrium notion of primarily mathematical interest, but for one satisfying specific tractability requirements that ensure easy implementation in computational agents. Specifically, a target concept should at least be:

- I. **expressive**, i.e. indicative of realistic behavior of human agents in sufficiently typical situations
- II. **cognitively plausible**: the equilibrium should **not** be justifiable in terms of expensive epistemic assumptions (the way common knowledge of rationality can be used to justify Nash equilibria [3]);
- III. **logically tractable**: proposed equilibria should be easy to specify formally, in a way that translates to efficient implementations.
- IV. **computationally tractable**: equilibria should be easy to compute [56], since bounded rational agents are assumed to compute (and play) them.

The main message of the paper is that such an extension is possible, but any general notion of Kantian equilibria may be of theoretical interest only: while we give an interesting extension for certain symmetric games (Sec. 5) inspired by the concept of *program equilibria*, it's not clear how to further extend it. Together with intractability (Thm. 2) this suggests that **a general, practically relevant, notion of Kantian equilibrium might not exist.**

2. *What is the relation between Kantian equilibria and Bacharach's (informally defined) team-reasoning equilibria [4]?* The answer is that Kantian equilibria are a proper subset of team-reasoning equilibria.
3. *Given that the answer to Q1 is negative, are there more specialized equilibria related to Kantian optimization that satisfy (I)-(IV)?* We will show that there exist, indeed, several more restrictive equilibrium notions, satisfying tractability and plausibility constraints, and relate them to Kantian equilibria.
4. *Real people are seldom purely selfish or purely Kantian. (How) can we formalize this?* We give such a definition, and motivate it through the case of Prisoners' Dilemma.

The outline of the paper is as follows: In Section 3 we review some basic notions. In Section 4 we obtain some further results on (and highlight some limitations of) Kantian equilibria: first of all, we point out that finding a mixed Kantian equilibrium is computationally intractable even for two-player symmetric games (Theorem 2). Second, multiple Kantian equilibria may exist, and lack of coordination on the same equilibrium may be detrimental to players, even with all of them playing a common linear combination of Kantian actions. In Section 5 we discuss the problem of extending Kantian equilibria to non-symmetric games, giving a proposal based on the concept of program equilibria. As such, our proposal inherits the problems of this concept. Given these problems, in Section 6 we propose several other-regarding equilibria.⁶ We show (Theorem 6) that these equilibria can be computed efficiently, that they are indeed Kantian equilibria (according to our generalized definition), and that they yield Kantian equilibria for symmetric coordination games. Finally, in Section 7 we relax the assumption that the agents are other-regarding: we assume that agents have a degree of greed, zero for Kantian agents, infinite for Nashian agents. We show (Theorem 8) how our definition applies to Prisoners' Dilemma.

For reasons of space, **most proof details are deferred to a longer version of the paper, available on arXiv [40]**. So have we done, for reasons of abundance of technical details, with some of the results: e.g. the ones the proper definition and characterization of Kantian program equilibria (Theorems 9, 10 in [40]), which also clarify the connection between Kantian and *team reasoning*.

⁶Generally, ethical egoism and its variant, rational egoism, are not accepted as a basis of moral behavior; counterexamples exist, [50]; however, it's fair to say that such positions are controversial, and somewhat marginal. In contrast, moral and other-regarding behaviors are better aligned, with other-regarding behavior often a consequence of moral play.

2 Related Work

The literature on other-regarding game-theoretic models is quite large, and a short section like this one cannot do justice to all the related, relevant work. Instead we have chosen to highlight a modicum of references directly relevant to our work.

The major impetus for this work was *Kantian optimization*. It was developed in [52, 53], developing early ideas of Laffont [43]. The current status of the theory is consolidated in the recent book [54]. A recent special issue of the *Erasmus Journal of Economics* is devoted to discussing and situating Roemer's contribution. Particularly valuable articles in this collection include [14, 60].

The other strand of ideas relevant to our work concerns the concept of *program equilibria*, defined in [66] and further investigated in [24, 41, 37, 44, 6, 19, 48]. There are several other related (and relevant) models, such as the *translucent player* model of [15, 31], or *mediated equilibria* [46].

The two other paradigms leading to the same concept for two-player symmetric coordination games, *superrationality* and (especially) *team reasoning* are, of course, relevant to our approach. Superrationality is rather different, though, and we only reiterate recommendations of [38, 25]. The main reference for team reasoning is still [5]. We also recommend papers [64, 18, 29].

Notions of symmetry in games have been insufficiently investigated, and they play an important role in defining Kantian programs. We refer to [33, 67] for such studies.

Finally, an impressive amount of work on behaviorally relevant game-theoretic notions related to moral behavior is summarized in [20]. While it is by no means comprehensive (especially with respect to the computer science literature), it is an excellent starting point.

3 Preliminaries

We assume knowledge of basic results of game theory at the level of a textbook such as, e.g. [49], in particular with concepts such as normal form games, best response strategy, and mixed (Nash) equilibria. All the games G we consider are normal form and, unless mentioned otherwise, have identical action sets Act_G for all players. Given a finite set S , we will define $\Delta(S)$ to be the set of probability distributions on s . Elements of $\Delta(S)$ are functions $c : S \rightarrow [0, 1]$ satisfying $\sum_{i \in S} c(i) = 1$. $\Delta^n := \Delta(\{1, 2, \dots, n\})$ is, geometrically, a $(n - 1)$ -dimensional simplex. When G is a normal-form game and k a player in the game we will denote by Δ_G^k the set of mixed actions available to player k , identified with some simplex Δ^n with a suitable dimension. We will occasionally drop k from the notation and simply write Δ_G instead when the player is clear from the context, or when all agents have the same action set. Given vectors $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_n)$, we say that x *dominates* y iff $x_i \geq y_i$ for all $i = 1, \dots, n$. The domination is *strict* if at least one inequality is. When comparing (mixed) action profiles \mathbf{a} and \mathbf{b} , the domination relation may apply to the vectors of agent utilities $(u_1(\mathbf{a}), \dots, u_n(\mathbf{a}))$ and $(u_1(\mathbf{b}), \dots, u_n(\mathbf{b}))$, respectively. Action profiles that are strictly dominated may be assumed not to occur in game play.

A game with identical action sets is *diagonal* if every pure action profile is Pareto dominated by some profile on the diagonal, the set of action profiles where all players play the same action. A particular class of diagonal games are *coordination games*, where all player utilities are zero outside the diagonal. Such a game is *symmetric* if, additionally, agent utilities are identical for all action profiles on the diagonal.

Definition 1. Let G be a game with common action set A . A variation function is a function $\phi : \Xi \times A \rightarrow A$, for some set of parameters Ξ .⁷ A Kantian (Hofstadter) equilibrium is a pure strategy profile

⁷The precise form of this definition follows [60], and is motivated by Roemer's definition of *additive/multiplicative Kantian*

		Player 2	
		C	D
Player 1	C	2,2	0,3
	D	3,0	1,1

		Player 2	
		B	S
Player 1	B	2,3	0,0
	S	1,1	3,2

Figure 1: a. Prisoners' Dilemma. b. BoS game as modified by Roemer.

$x^{opt} = (x_1^{opt}, \dots, x_n^{opt})$ that maximizes the material payoff of each agent, should everyone deviate similarly. Formally, for every agent i and $r \in \Xi$,

$$V_i(x_1^{opt}, \dots, x_n^{opt}) \geq V_i(\phi(r, x_1^{opt}), \dots, \phi(r, x_n^{opt})).$$

Example 1. One of the original applications of Kantian equilibria was Prisoners' Dilemma (PD, Fig. 1 (a)). Kantian equilibria provide an elegant solution to the paradox: Kantian agents coordinate on action profile (C,C), as jointly doing so gives them a higher payoff than the Nash equilibrium (D,D).

Kantian equilibria are easiest to justify for symmetric diagonal games, since in this case they dominate all other action profiles, thus can be properly seen as "best course of action for all". There are symmetric (nondiagonal) games, though, where no pure strategy Kantian equilibrium is adequate, and which seem to compel us to considering mixed-strategy Kantian equilibria. An example is Hofstadter's "Platonia's Dilemma" [38], a special case of the *market entry games* of Selten and Güth [59]:

Definition 2. In Platonia Dilemma n agents (say, $n = 20$) are offered to win a prize. Agents may choose to send their name to a referee. An agent wins the prize if and only if it is **the only one submitting their name**: if zero or at least two agents send their names then noone wins anything.

It is easy to see that both pure strategies, sending/not sending their name, are equally bad if adopted by all agents: they get zero payoff. A better option is to allow independent randomization:

Definition 3. Given a game G with identical actions, a **mixed Kantian agent** will choose a mixed strategy $X^{OPT} \in \Delta_G$ that maximizes its expected utility, should everyone play X . For two-player symmetric games with game matrix A and variation function $\phi(b, a) = b$ we have $X^{OPT} = \operatorname{argmax}\{y^T A y : y \in \Delta_G\}$.

Lemma 1. In Platonia Dilemma the probabilistic strategy where each agent independently submits their name with probability p brings an expected profit to every agent equal to $p(1-p)^{n-1}$. This quantity is maximized for $p = \frac{1}{n}$. Thus the strategy with $p = \frac{1}{n}$ is a mixed Kantian equilibrium.

Proof. Let $f(p) = p(1-p)^{n-1}$. $f'(p) = (1-p)^{n-1} - (n-1)p(1-p)^{n-2} = (1-p)^{n-2}(1-np)$, so f is increasing on $[0, 1/n]$ and decreasing on $[1/n, 1]$. \square

4 Limitations of Mixed Kantian Equilibria

In this section we note some properties of mixed Kantian equilibria. They are mostly negative: finding a mixed Kantian equilibrium is intractable. Also, such equilibria may be vulnerable to miscoordination.

equilibria, with action set $A = \mathbb{R}_+$ and variation functions $\phi(r, a) = a+r$, $a \cdot r$, respectively. **We will mostly be concerned with variation functions of the type "change (everyone's) current action to b "** (for $b \in A$). Formally, $\Xi = A$ and $\phi(b, a) = b$.

4.1 The Computational Intractability of Mixed Kantian Equilibria

We make an easy observation concerning the computational complexity of mixed Kantian equilibria in symmetric two player games. To our knowledge this has not been discussed before. Note: such equilibria are guaranteed to exist, since the $(n - 1)$ -dimensional simplex of mixed strategies is a compact set and the common utility function is continuous. First of all, finding a mixed Kantian equilibrium is easy in symmetric coordination games, as all such equilibria coincide with pure Kantian equilibria.

Theorem 1. *Consider a finite symmetric coordination game. Then mixed Kantian equilibria coincide with pure Kantian equilibria. Hence one can compute mixed Kantian equilibria in polynomial time.*

Platonian Dilemma with $n = 2$ shows that Theorem 1 does not extend to general symmetric games. This is no coincidence: in this case finding (or just detecting) the optimal mixed strategy is intractable:

Theorem 2. *The following problem, called MIXED KANTIAN EQUILIBRIUM, is NP-hard:*

INPUT: A two-player symmetric game G , and an aspiration level $r \in \mathbb{Q}$.

TO DECIDE: Is there a mixed strategy profile $x = (x_1, \dots, x_N)$ such that the utility of every player under common mixed action $x_1a_1 + x_2a_2 + \dots + x_m a_m$ is at least r ?

Proof. We point out to the existence of a reduction from CLIQUE to MIXED KANTIAN EQUILIBRIUM, that shows that the latter problem is NP-hard. In fact the reduction will only consider symmetric games with 0/1 payoffs.

Consider, indeed, a graph g . Let k be an integer and (g, k) be the corresponding instance of CLIQUE.

Define the symmetric two-player game G whose payoff matrix is the adjacency matrix A of g .

Mixed Kantian equilibria $x = (x_1, \dots, x_N)$ of G correspond to optimal solutions of the following quadratic program:

$$\begin{cases} \max(x^T A x) \\ x_1 + \dots + x_N = 1 \\ x_1, \dots, x_N \geq 0. \end{cases} \quad (1)$$

This is a problem that has been called [12] *the standard quadratic optimization problem*, and has been investigated substantially in the global optimization literature (see e.g. [11]). A beautiful result due to Motzkin and Straus [47] can be restated as claiming that for programs whose matrix A is the adjacency matrix of a graph g , if o is the optimum of problem (1) then $\frac{1}{1-o}$ is the size of the maximum clique in g .

Hence $(g, k) \in \text{CLIQUE}$ if and only if $(G, \frac{k-1}{k}) \in \text{MIXED-KANTIAN-EQUILIBRIUM}$. □

4.2 Multiple Equilibria and miscoordination

Optimal diagonal action profiles may fail to be unique. If the agents are not communicating (and no implicit coordination mechanisms are acting, e.g., one of the action profiles being a *focal point*, such as in the Hi-Lo game from [5]), agents may reach a suboptimal action profile due to their lack of coordination on the same optimal action: Consider, indeed, the game in Figure 2. (C, C) and (E, E) are equally good pure (and mixed) Kantian equilibria. But if one player plays C and the other plays E the resulting outcomes are the worst possible for both of them, being dominated by every single possible strategy profile! Randomizing among Kantian actions might not help either: miscoordination impacts even "Kantian" scenarios, where players, lacking a salient equilibrium to coordinate on, play a joint mixed strategy formed of Kantian actions.⁸ We quantify the degradation in performance as follows:

⁸Such a scenario is, of course, not justifiable from a usual rational choice perspective. But **it is justifiable in a Kantian setting where every player believes that choosing a pure action a will immediately make all other players do the same:** a

		P1 2		
		C	D	E
P1 1	C	5, 5	3, 6	1, 2
	D	6, 3	4, 4	6, 3
	E	2, 1	3, 6	5, 5

		P1 2	
		C	D
C	C	10, 1	0, 0
	D	0, 0	4, 2

		P1 2	
		B	S
B	B	6, 1	0, 0
	S	0, 0	3, 2

		P1 2	
		C	S
C	C	10, 10	100, 200
	S	200, 100	6, 6

Figure 2: (a). A game with multiple Kantian equilibria. (b). Modified BoS (Example 3). (c). Modified BoS (Example 2). (d). An anti-coordination game.

Definition 4. For a symmetric game G with strictly positive payoffs let NC be the set of mixed action profiles composed of Kantian actions only. The price of miscoordination of G is the ratio $p(G) = \sup_{a \in NC} \frac{u_i(X^{OPT})}{u_i(a)}$. Because of symmetry this does not depend on the particular choice of player i .

The following result shows that the price of miscoordination can be arbitrarily large:

Theorem 3. Let G be a symmetric diagonal game with $k \geq 2$ players and $r \geq 1$ pure Kantian actions. Then the price of miscoordination of G is in the range $[1, r^{k-1}]$. Both bounds are tight and can be reached in settings where players choose a Kantian action uniformly at random.

The merit of this simple result is to point out that the definition of generalized Kantian equilibria needs to include scenarios where randomness is correlated, as in *correlated equilibria* (see e.g. [61]).

Proof. The price of miscoordination is insensitive to dividing all utilities by the same factor λ , so w.l.o.g. one may assume that the utilities agent receive on pure Kantian equilibrium profiles is 1. For the mixed action \mathbf{a} where players play the r Kantian actions (w.l.o.g. $1, 2, \dots, r$) with probabilities p_1, p_2, \dots, p_r (which add up to 1), its expected utility is $E[u_i(\mathbf{a})] = \sum u_i(i_1, i_2, \dots, i_k) \cdot p_{i_1} p_{i_2} \dots p_{i_k} \geq \sum_{i=1}^r p_i^k \geq r \cdot \frac{1}{r} = \frac{1}{r^{k-1}}$, by Jensen’s inequality. The upper bound is obtained when off-diagonal action profiles formed of Kantian actions only have utilities equal to 0. As for the lower bound, for diagonal games by domination we have $u_i(i_1, i_2, \dots, i_k) \leq 1$, so $E[u_i(\mathbf{a})] = \sum u_i(i_1, i_2, \dots, i_k) \cdot p_{i_1} p_{i_2} \dots p_{i_k} \leq \sum p_{i_1} p_{i_2} \dots p_{i_k} = (p_1 + p_2 + \dots + p_r)^k = 1$. A game realizing the lower bound is the one where agent utilities on all pure action profiles are equal to 1. □

5 Kantian Program Equilibria in (Pareto) Symmetric Games

Definition 1 of Kantian equilibria makes the most sense in symmetric coordination games, but does not capture all the intuitive cases of Kantian behavior. Indeed, let us consider the BoS game, as modified by Roemer (Fig. 1 (b)).⁹ Intuitively, agents would perhaps agree that the following **protocol** could be called Kantian, in that it is symmetric and both players benefit if they both follow it: flip a fair coin; if it comes out *heads*, they (both) play B , else both play S . As described, the protocol requires the centralized choice of a random bit, but it could easily be implemented in a distributed manner by making each of the

player may use the Kantian imperative to restrict itself to pure Kantian equilibria, then use the assumption to justify playing a convex combination of pure Kantian equilibria it is indifferent between.

⁹Roemer ([54], Proposition 2.3) argues that (S,B) is a simple Kantian equilibrium. His argument is, however, ad-hoc, based on making this profile "diagonal" by flipping the order of B and S for the second player, and the conclusion that (B,S) is Kantian is, we feel, unintuitive, since $(B,B), (S,S)$ strictly dominate it. Our protocol plays $\frac{1}{2}(B,B) + \frac{1}{2}(S,S)$, different from (and better than) what Roemer calls the mixed Kantian equilibrium, where row player plays $\frac{3}{8}B + \frac{5}{8}S$ and the column player plays $\frac{3}{8}S + \frac{5}{8}B$.

two agents flip a (fair) coin and taking their XOR. The implementation of the protocol (Algorithm 5.1) assumes that each agent is parameterized by an *agent ID* $i \in \{1, 2\}$ (not needed in this particular case) and vector $(myb, otherb)$ of random bit choices, one for each player, and shared between players.

An even more dramatic case is that of the game from Figure 2 (d), where the best outcomes are not symmetric. In these cases it even seems irrational for the agents to play symmetric action profiles, since these action profiles are dominated by all the other action profiles! Rather, it is plausible that agents would agree that they need to **anticoordinate**, but they have different preferences for the joint action profile to coordinate upon. A "best for all" solution would jointly play a random anti-coordinated profile $\frac{1}{2}(C, S) + \frac{1}{2}(S, C)$. As in the previous example, this course of action can be implemented by the two agents in a distributed manner, by jointly playing according to the protocol in Algorithm 5.2. In this example, in addition to the extra bit *otherb* communicated by the other player, the protocol of each agent **makes explicit use of the agents' own *id***, $i \in \{1, 2\}$.

Algorithm 5.1: (*BoS*)

$BOS(i :: ID, myb :: BIT, otherb :: BIT)$

Randomly choose a bit $myb \in \{0, 1\}$
 communicate *mybit* to the
 other player as its *otherb*.
if $[myb \oplus otherb == 0]$
 then play *B*
 else play *S*

Algorithm 5.2: (*Anticoord*)

$Anticoord(i :: ID, myb :: BIT, otherb :: BIT)$

Randomly choose a bit $myb \in \{0, 1\}$
 communicate *myb* to the
 other player as its *otherb*.
if $[myb \oplus otherb \equiv i \pmod{2}]$
 then play *C*
 else play *S*

The intuitive conclusion of these two examples is simple: Definition 1 is not sufficient. Some simple games may have coordinated **protocols** that could properly be called "Kantian". In this section **we give a somewhat more general definition¹⁰ of Kantian equilibria, but not for general games, only for a class of "symmetric" games**. There are multiple definitions of game symmetry in the literature [33, 67]; the most important one requires that for every player i , action profile (x_1, x_2, \dots, x_n) and permutation $\sigma \in S_n$, we have $u_{\sigma(i)}(x_1, x_2, \dots, x_n) = u_i(x_{\sigma(1)}, x_{\sigma(2)}, \dots, x_{\sigma(n)})$. We use a slightly less demanding definition (we call our version *Pareto symmetry*, see Definition 16 in the longer version [40]) to capture some asymmetric games like Romer's version of BoS. This game is asymmetric, but in an inconsequential way: all the asymmetries concern dominated profiles. Our extension is inspired by the concept of **program equilibria** [39, 66]. To define these equilibria we associate to every finite normal-form game an extended game whose actions correspond to *programs*. Agents' programs have access to own and others' sources¹¹, and can act on them. A program equilibrium is a Nash equilibrium in the extended game. We extend this idea to Kantian equilibria:

Definition 5. *Given a (Pareto symmetric) game G with identical actions sets for all players, a Kantian program equilibrium in G is a probability distribution p on the action profiles of G such that: (a). p has*

¹⁰other plausible alternatives are discussed (and ruled out) in Section 13 of the longer version [40]. Of special interest is the relation between Kantian equilibria and *team-reasoning equilibria*.

¹¹this aspect was important for Nashian optimization. It will be less so for us, since deviations are not present in the definition of Kantian equilibria. On the other hand, a Kantian agent playing program P can make sure it is not taken advantage upon by the other players, either alone, as in [66], by reading other players' programs and only playing P when all do, or with the help of a *mediator*, which implements on behalf of all players the following protocol: if all agents follow P then the mediator will simulate P on behalf of the agent; otherwise it will play in a Nashian way.

its support on the set of action profiles that are Pareto optimal, that is not Pareto dominated by any other action profile. (b). p is implemented by **agents playing a common program P** in the extended game. (c). there exists no probability distribution q with properties a. and b. such that the vector of expected utilities $(E[u_i(q)])$ Pareto dominates the vector $(E[u_i(p)])$.

The assumption that *players have the same action sets* is motivated by the "common program" requirement of Def. 5 (b). Point (a). encodes a simple rationality condition. Points (b). and (c). embody a generalized version of the Kantian categorical imperative: (b). encodes the constraint of identical behavior, (c). encodes the fact that implementing p is "a best action" for all players.

We only need to formalize what we mean by program in this definition. The semantics is inspired by the one in [66], but the full formalization is somewhat subtle. We defer a full presentation to Section 15 of the longer version [40]. A couple of technical points are, however, worth stating:

- The semantics of programs in [66] does not allow for any synchronization between different versions of the same program, other than testing whether they are syntactically equivalent. Since (as recognized in Theorem 3) we need to include correlated randomness, we need to extend the semantics of programs from [66], where this is not possible. There are many ways to do it, but one way is to allow *correlated sampling from distributions*: all the programs get an identical sample from a given distribution.
- **However, simply adding correlated sampling of action profiles to the semantics of programs from [66] leads (see Theorem 9 in the long version [40]) to paradoxical results:** every convex combination of Pareto optimal strategy profiles would be a Kantian program equilibrium. This is an issue: in Prisoners' Dilemma profiles $(C, D), (D, C)$ could perhaps be justified from a "team reasoning" perspective where one player "sacrifices itself" so that the other one walks free. To accept them as "Kantian equilibria" seems, however, problematic (see also the discussion in sections 13.3 and 14 of [40])
- If, on the other hand, **agent programs didn't communicate at all, used no private randomness, or used no specific ID/payoff information** then they would run identically for all agents, coordinating on the same action (excluding, thus, scenarios like that of Example 3, that we want to model).
- We will take a middle-ground approach, and assume that agents can use their ID and the information about the game payoffs in a very limited way, that makes the program act "identically with respect to a group of symmetries acting transitively on the set of agents". This requires us to restrict ourselves to the class of Pareto symmetric games of Definition 16 in the longer version [40], whose set of Pareto dominant action profiles has such symmetries. The precise technical details are spelled out in Section 15 of the longer version [40], where we prove (Theorem 10) a characterization of Kantian equilibria for Pareto symmetric games which also shows that Kantian program equilibria are a strict subset of the class of team-reasoning equilibria.

Algorithms 4.1 and 4.2 lend some credibility to the intuition that Kantian equilibria are somehow related to some "symmetric" notion of correlated equilibria. This intuition is correct: in Definition 19 in the longer version [40] we define a notion of "correlated symmetric equilibrium". We then prove:

Theorem 4. *Correlated symmetric equilibria of symmetric games are Kantian program equilibria.*

Kantian program equilibria allow players to obtain a better expected payoff in Platonia Dilemma:

Theorem 5. *Algorithm 5.3 implements a Kantian program equilibrium for Platonia Dilemma.*

Proof. Points (a). and (b). from the definition of Kantian program equilibria are clear, the only one that merits a discussion is point (c).

The expected utility of each player under Algorithm 5.3 is equal to $1/n$. Since the sum of utilities of all players under a particular set of random choices is equal to 1, no vector of expected utilities can strictly dominate the vector $(1/n, 1/n, \dots, 1/n)$ of expected utilities for the Algorithm. \square

Algorithm 5.3: CHOOSE-WINNER(i, b_1, b_2, \dots, b_n)

```

Randomly choose an integer  $b_i \in \mathbb{Z}_n$ 
if  $[\sum_{j=1}^n b_j \equiv i \pmod{n}]$ 
  then S(UBMIT)
  else D(ON'T)
```

6 Some computationally efficient other-regarding equilibria

As defined in the previous section, Kantian program equilibria for games with identical action sets inherit some of the definitional problems of "ordinary" program equilibria. Among them:

- *fragility*: (Kantian) program equilibria are sensitive (see e.g. [48]) to the precise specification of programs: do we insist that all agent programs are syntactically identical, or just "do the same thing"? See [44, 37] for some attempted solutions for program equilibria that could be adapted to our setting.
- *lack of generality*: Definition 5 it is only applicable to (some of the) games with identical action sets. To further generalize it to all finite normal-form games one would need to specify what it means for two agents to "take the same course of action" in settings with differing action sets.
- *lack of predictive power*: There may be multiple (even infinitely many) Kantian program equilibria.

Given these objections, and with constraints (I)-(IV) in mind, we propose in the sequel a substantially more modest approach: Rather than seeking a general definition of Kantian equilibria we propose instead *several* other-regarding equilibria. **They all correspond intuitively to real-life situations, are tractable, can be justified by team reasoning and are related, for symmetric coordination games, to Kantian equilibria.** One was independently suggested in [42], the other ones are first introduced here:

Definition 6. A Rawlsian equilibrium is a probability distribution over Pareto optimal profiles maximizing the egalitarian social welfare (the expected utility of the worst-off player) and is strictly dominated by no other profile with this property. Such equilibria implement the idea of justice as fairness [51].

Example 2. We modify the BoS example as in Fig. 2 (c): perhaps I is a classical music lover, that gets a higher utility than the other player by going, together with its partner, to any of the two concerts. Then (S,S) is the (unique) Rawlsian equilibrium. Choosing such an equilibrium is an example of altruistic behavior from player I, since it maximizes the payoff of its non-music-lover partner.

Definition 7. A Bentham-Harsányi equilibrium is a probability distribution on Pareto optimal profiles maximizing the sum of expected payoffs. See [34] for a philosophical motivation. A best-off equilibrium is a prob. distrib. on Pareto optimal profiles maximizing the largest expected payoff, and strictly dominated by no profile with this property. E.g., in Exp. 2 (B,B) is the unique Bentham-Harsányi/best-off equilibrium.

Although a best-off equilibrium may not seem "fair", there exist real-life "team reasoning" situations that elicit behavior suggestive of such an equilibrium: one such example is, for instance, scenarios where members of a team "sacrifice" for one of their members (e.g. parents for a child).

The equilibrium notions we introduced so far implicitly assumed that player utility is given by material payoffs. Sometimes the frustration a player feels is derived by counterfactually comparing its realized payoff with all possible ones. There are many implementations of this idea. The following notion quantifies the extent to which a given profile is worse for the given player than a random profile.

Definition 8. The percentile index of profile a for player i is the percentage of Pareto optimal profiles that would get i a strictly better payoff than a . A Rawlsian percentile equilibrium is a profile minimizing the largest expected percentile index of all players, and strictly dominated by no profile with this property.

Example 3. Consider the game shown in Figure 2 (b). Then percentile indices of Pareto optimal profiles are $(0, 100)$ for (C, C) , and $(100, 0)$ for (D, D) , respectively. Profile $\frac{1}{2}(C, C) + \frac{1}{2}(D, D)$ is a Rawlsian percentile equilibrium. Player 1 gets average utility 7 while player 2 gets average utility $\frac{3}{2}$.

An even less cognitively sophisticated model of agent frustration relies on classifying outcomes as "happy/not happy". The following is a simple example of such a notion:

Definition 9. The natural expectation point of player i is the median (over all undominated pure strategy profiles) payoff. If there are two medians then the average value is taken. A player is happy in a pure strategy profile a iff its payoff is larger or equal than its natural expectation point and unhappy otherwise.

An aspiration equilibrium is a mixed strategy profile that minimizes the largest probability of unhappiness among all players and is strictly dominated by no other profile with this property.

Example 4. Take a coordination game with payoffs $(C, C) \rightarrow (10, 1)$, $(D, D) \rightarrow (9, 2)$, $(E, E) \rightarrow (8, 3)$, $(F, F) \rightarrow (4, 7)$. The natural expectation points of players are 8.5 and 2.5, respectively. The first player is happy in (C, C) and (D, D) , the second in (E, E) , (F, F) . Hence in $\frac{1}{4}(C, C) + \frac{1}{4}(D, D) + \frac{1}{4}(E, E) + \frac{1}{4}(F, F)$ the players are happy 50% of the time and no mixed action profile can do any better.

Unlike general Kantian program equilibria, the equilibria we defined are computationally tractable:

Theorem 6. Rawlsian, Rawlsian percentile, Bentham-Harsányi, best-off, aspiration equilibria exist and can be found by solving a sequence of linear programs (hence in polynomial time).

We now connect our other regarding equilibria to Kantian equilibria in symmetric coordination games. We call an equilibrium point *extremal* if it cannot be written as a nontrivial convex combination of other (similar) equilibria. We show that extremal self-regarding equilibria generalize Kantian pure equilibria. Extremality is needed, since our equilibria are closed under convex combinations (such combinations are justifiable from a magical thinking perspective, see footnote 7), while pure Kantian equilibria are not. Because of Thm. 2 no similar connection is likely for mixed Kantian equilibria:

Theorem 7. In symmetric diagonal games Rawlsian, Bentham-Harsányi, best-off, Rawlsian percentile, aspiration equilibria coincide with convex combinations of Kantian pure equilibria.

7 Agents with bounded greed

So far we have assumed that people are other-regarding. In reality people are not unrestricted optimizers, nor are they perfect Kantian moralists. Alger and Weibull [1] attempted to interpolate between utilitarian agents and Kantian ones, by defining *homo moralis* to be an agent whose utility has the form $u_i(x, y) = (1 - k)\pi(x, y) + k\pi(x, x)$, where $k \in [0, 1]$ is the so-called *degree of morality* of the agent. They showed that evolutionary models with assortative mixing and incomplete information favor a particular kind of *homo moralis*, those whose degree of morality coincides with the degree of assortativity of the matching process. Interesting as this result is, it has some weaknesses. For instance [1], *homo moralis* behaves like *homo economicus* in Prisoners' Dilemma and all constant-sum games when $k \neq 1$. In other words, agent behavior is not sensitive to the degree of morality, as long as the agent is not Kantian.

We give (for symmetric games, but the idea can be extended to general ones, via Kantian program equilibria) a definition with the same overall intention, but capturing a slightly different agent behavior:

Definition 10. Let $\lambda \in [1, \infty]$. Agent i is called λ -utilitarian if, for every action profile (a_i, b) , its utility $u_i(a_i, b)$ is (a). $\pi_i(a_i, (\bar{a}_i)_{-i})$ if a_i is a Kantian action. (b). 0 if a_i is not Kantian and $\pi_i(a_i, b) \leq \lambda \cdot \pi_i(X^{OPT})$; (c). $\pi_i(a, b)$ if a_i is not Kantian and $\pi_i(a_i, b) \leq \lambda \cdot \pi_i(X^{OPT})$. I.e., a λ -utilitarian agent deviates from its Kantian action X^{OPT} **only** if the utility it obtains is more than λ times larger.

We call the number $\frac{1}{\lambda-1}$ the *greed index* of i . It varies between 0 (Kantian agents) and ∞ (purely utilitarian ones). The natural equilibrium concept for such agents is no longer Kantian, but Nash equilibrium. Definition 10 allows giving an empirically plausible justification of all possible outcomes in PD:

Theorem 8. All pure action profiles in PD are Nash equilibria of agents with varying degrees of greed.

Proof. Bounded-greed agents still coordinate on the Kantian equilibrium (C, C) as long as both their greed indices are < 2 (i.e. they would need at least a twofold increase in payoff to deviate). If one of them has greed index < 2 and the other one has greed index ≥ 2 , then the latter one will defect. If both agents have greed indices ≥ 2 , then they will coordinate, just as if utilitarian agents would do, on the Nash equilibrium (D, D) . \square

8 Conclusions

Our main contribution is bringing Kantian equilibria (and related concepts) to the attention of TARK community, showing that this notion is theoretically interesting, but that the road to implementable behaviors probably goes through less general equilibrium concepts. Many of the notions we introduced, on the other hand, including Kantian program equilibria and bounded greed agents, deserve further investigation. For instance a justification like that of Theorem 8 could be used as a rationality criterion. One could look for evolutionary justifications of bounded greed agents along the lines of [1]. One could use such agents in relation to work on the concept of *price of anarchy* [57]. On a more conceptual level, the use of *frames* in game theory [5, 7] and how this interacts with equilibrium notions deserves further study. Finally, several open problems remain: Can we find algorithms for our equilibria that bypass the need for solving multiple LP's? Is the problem from Theorem 2 NP-complete (i.e. in NP)?

References

- [1] Ingela Alger & Jörgen W Weibull (2013): *Homo moralis - preference evolution under incomplete information and assortative matching*. *Econometrica* 81(6), pp. 2269–2302, doi:10.3982/ECTA10637.
- [2] Dan Ariely (2010): *Predictably irrational*. Harper.
- [3] Robert Aumann & Adam Brandenburger (1995): *Epistemic conditions for Nash equilibrium*. *Econometrica: Journal of the Econometric Society*, pp. 1161–1180, doi:10.2307/2171725.
- [4] Michael Bacharach (1999): *Interactive team reasoning: A contribution to the theory of co-operation*. *Research in economics* 53(2), pp. 117–147, doi:10.1006/reec.1999.0188.
- [5] Michael Bacharach (2006): *Beyond individual choice: teams and frames in game theory*. Princeton University Press, doi:10.1515/9780691186313.
- [6] Mihaly Barasz, Paul Christiano, Benja Fallenstein, Marcello Herreshoff, Patrick LaVictoire & Eliezer Yudkowsky (2014): *Robust Cooperation in the Prisoner's Dilemma: Program Equilibrium via Provability Logic*. *arXiv preprint https://arxiv.org/abs/1401.5577*.
- [7] José Luis Bermúdez (2021): *Frame it Again: New Tools for Rational Decision-making*. Cambridge University Press.
- [8] Kenneth G. Binmore (1994): *Game theory and the social contract: just playing*. M.I.T. Press.

- [9] Kenneth G. Binmore (1994): *Game theory and the social contract: playing fair*. M.I.T. Press.
- [10] Kenneth G. Binmore (2005): *Natural justice*. Oxford University Press, USA, doi:10.1093/acprof:oso/9780195178111.001.0001.
- [11] Immanuel M Bomze (1997): *Evolution towards the maximum clique*. *Journal of Global Optimization* 10(2), pp. 143–164, doi:10.1023/A:1008230200610.
- [12] Immanuel M Bomze (1998): *On standard quadratic optimization problems*. *Journal of Global Optimization* 13(4), pp. 369–387, doi:10.1023/A:1008369322970.
- [13] Samuel Bowles & Herbert Gintis (2013): *A cooperative species: Human reciprocity and its evolution*. Princeton University Press.
- [14] Matthew Braham & Martin van Hees (2020): *Kantian Kantian Optimization*. *Erasmus Journal for Philosophy and Economics* 13(2), pp. 30–42, doi:10.23941/ejpe.v13i2.513.
- [15] Valerio Capraro & Joseph Y Halpern (2019): *Translucent players: Explaining cooperative behavior in social dilemmas*. *Rationality and Society*, pp. 371–408, doi:10.1177/2F1043463119885102.
- [16] Jing Chen & Silvio Micali (2016): *Auction revenue in the general spiteful-utility model*. In: *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, pp. 201–211, doi:10.1145/2840728.2840741.
- [17] Po-An Chen & David Kempe (2008): *Altruism, selfishness, and spite in traffic routing*. In: *Proceedings of the 9th ACM conference on Electronic Commerce*, pp. 140–149, doi:10.1145/1386790.1386816.
- [18] Andrew M Colman & Natalie Gold (2018): *Team reasoning: Solving the puzzle of coordination*. *Psychonomic Bulletin & Review* 25(5), pp. 1770–1783, doi:10.1080/10002003098538748.
- [19] Andrew Critch (2019): *A parametric, resource-bounded generalization of Löb’s theorem, and a robust cooperation criterion for open-source game theory*. *The Journal of Symbolic Logic*, pp. 1–15, doi:10.1017/jsl.2017.42.
- [20] Sanjit Dhami (2016): *The foundations of behavioral economic analysis*. Oxford University Press.
- [21] Jon Elster (2017): *On seeing and being seen*. *Social Choice and Welfare* 49(3-4), pp. 721–734, doi:10.1007/s00355-017-1029-9.
- [22] Ernst Fehr & Klaus M Schmidt (1999): *A theory of fairness, competition, and cooperation*. *The quarterly journal of economics* 114(3), pp. 817–868, doi:10.1162/003355399556151.
- [23] Urs Fischbacher, Simon Gächter & Ernst Fehr (2001): *Are people conditionally cooperative? Evidence from a public goods experiment*. *Economics Letters* 71(3), pp. 397–404, doi:10.1016/S0165-1765(01)00394-9.
- [24] Lance Fortnow (2009): *Program equilibria and discounted computation time*. In: *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge*, pp. 128–133, doi:10.1145/1562814.1562833.
- [25] Ghislain Fourny (2020): *Perfect Prediction in normal form: Superrational thinking extended to non-symmetric games*. *Journal of Mathematical Psychology* 96, p. 102332, doi:10.1016/j.jmp.2020.102332.
- [26] Robert H Frank (2004): *What Price the Moral High Ground? Ethical Dilemmas in Competitive Environments*. Princeton University Press.
- [27] Herbert Gintis (2016): *Individuality and entanglement: the moral and material bases of social life*. Princeton University Press, doi:10.2307/j.ctvc779cx.
- [28] Herbert Gintis (2016): *A Typology of Human Morality*. In David S. Wilson & Alan Kirman, editors: *Complexity and Evolution: Towards a New Synthesis for Economics*, M.I.T. Press, doi:10.7551/mitpress/9780262035385.003.0007.
- [29] Natalie Gold & Andrew M Colman (2020): *Team reasoning and the rational choice of payoff-dominant outcomes in games*. *Topoi* 39(2), pp. 305–316, doi:10.1007/s11245-018-9575-z.
- [30] Davide Grossi & Paolo Turrini (2012): *Dependence in games and dependence games*. *Autonomous Agents and Multi-Agent Systems* 25(2), pp. 284–312, doi:10.1007/s10458-011-9176-3.

- [31] Joseph Y Halpern & Rafael Pass (2018): *Game theory with translucent players*. *International Journal of Game Theory* 47(3), pp. 949–976, doi:10.1007/s00182-018-0626-x.
- [32] Joseph Y Halpern & Nan Rong (2010): *Cooperative equilibrium*. In: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*, pp. 1465–1466. Available at <http://ifaamas.org/Proceedings/aamas2010/pdf/02%20Extended%20Abstracts/Red/R-49.pdf>.
- [33] Nicholas Ham (2013): *Notions of Symmetry for Finite Strategic-Form Games*. arXiv preprint <https://arxiv.org/abs/1311.4766>.
- [34] John C Harsanyi (1955): *Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility*. *Journal of Political Economy* 63(4), pp. 309–321, doi:10.1086/257678.
- [35] John C Harsanyi (1977): *Rule utilitarianism and decision theory*. *Erkenntnis* 11(1), pp. 25–53, doi:10.1007/BF00169843.
- [36] Martin Hoefer & Alexander Skopalik (2013): *Altruism in atomic congestion games*. *ACM Transactions on Economics and Computation (TEAC)* 1(4), pp. 1–21, doi:10.1145/2542174.2542177.
- [37] Wiebe van der Hoek, Cees Witteveen & Michael Wooldridge (2013): *Program equilibrium—a program reasoning approach*. *International Journal of Game Theory* 42(3), pp. 639–671, doi:10.1007/s00182-011-0314-6.
- [38] Douglas Hofstadter (1985): *Dilemmas for Superrational Thinkers, Leading up to a Luring Lottery*. In: *Metamagical Themas: Questing for the Essence of Mind and Pattern*, Basic Books.
- [39] John V Howard (1988): *Cooperation in the Prisoner’s Dilemma*. *Theory and Decision* 24(3), p. 203, doi:10.1007/BF00148954.
- [40] Gabriel Istrate (2021): *Game-theoretic Models of Moral and Other-Regarding Agents*. arXiv preprint <http://arxiv.org/abs/2012.09759v2>.
- [41] Adam Tauman Kalai, Ehud Kalai, Ehud Lehrer & Dov Samet (2010): *A commitment folk theorem*. *Games and Economic Behavior* 69(1), pp. 127–137, doi:10.1016/j.geb.2009.09.008.
- [42] Ioannis Kordonis (2020): *A Model for Partial Kantian Cooperation*. In: *Advances in Dynamic Games*, Springer, pp. 317–346, doi:10.1007/978-3-030-56534-3_13.
- [43] Jean-Jacques Laffont (1975): *Macroeconomic constraints, economic efficiency and ethics: An introduction to Kantian economics*. *Economica* 42(168), pp. 430–437, doi:10.2307/2553800.
- [44] Patrick LaVictoire, Benja Fallenstein, Eliezer Yudkowsky, Mihaly Barasz, Paul Christiano & Marcello Herreshoff (2014): *Program equilibrium in the Prisoner’s Dilemma via Löb’s theorem*. In: *Workshops at the Twenty-Eighth AAAI Conference on Artificial Intelligence*. Available at https://mipc.inf.ed.ac.uk/2014/papers/mipc2014_lavictoire_etal.pdf.
- [45] Sydney Levine, Max Kleiman-Weiner, Laura Schulz, Joshua Tenenbaum & Fiery Cushman (2020): *The logic of universalization guides moral judgment*. *Proceedings of the National Academy of Sciences* 117(42), pp. 26158–26169, doi:10.1073/pnas.2014505117.
- [46] Dov Monderer & Moshe Tennenholtz (2009): *Strong mediated equilibrium*. *Artificial Intelligence* 173(1), pp. 180–195, doi:10.1016/j.artint.2008.10.005.
- [47] Theodore S Motzkin & Ernst G Straus (1965): *Maxima for graphs and a new proof of a theorem of Turán*. *Canadian Journal of Mathematics* 17, pp. 533–540, doi:10.4153/CJM-1965-053-6.
- [48] Caspar Oesterheld (2019): *Robust program equilibrium*. *Theory and Decision* 86(1), pp. 143–159, doi:10.1007/s11238-018-9679-3.
- [49] Martin Osborne & Ariel Rubinstein (1994): *A Course in Game Theory*. M.I.T. Press.
- [50] Ayn Rand (1964): *The virtue of selfishness*. Penguin.
- [51] John Rawls (2001): *Justice as fairness: A restatement*. Harvard University Press.

- [52] John E Roemer (2010): *Kantian equilibrium*. *Scandinavian Journal of Economics* 112(1), pp. 1–24, doi:10.1111/j.1467-9442.2009.01592.x.
- [53] John E Roemer (2015): *Kantian optimization: A microfoundation for cooperation*. *Journal of Public Economics* 127, pp. 45–57, doi:10.1016/j.jpubeco.2014.03.011.
- [54] John E Roemer (2019): *How We Cooperate: A Theory of Kantian Optimization*. Yale University Press, doi:10.2307/j.ctvfc52jk.
- [55] Nan Rong & Joseph Y Halpern (2013): *Towards a deeper understanding of cooperative equilibrium: characterization and complexity*. In: *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, pp. 319–326. Available at <http://www.ifaamas.org/Proceedings/aamas2013/docs/p319.pdf>.
- [56] Iris van Rooij, Mark Blokpoel, Johan Kwisthout & Todd Wareham (2019): *Cognition and intractability: A guide to classical and parameterized complexity analysis*. Cambridge University Press, doi:10.1093/comjnl/bxm038.
- [57] Tim Roughgarden (2005): *Selfish Routing and the Price of Anarchy*. M.I.T. Press.
- [58] Sally Sedgwick (2008): *Kant's groundwork of the metaphysics of morals: an introduction*. Cambridge University Press, doi:10.1017/CBO9780511809538.
- [59] Reinhard Selten & Werner Güth (1982): *Equilibrium point selection in a class of market entry games*. In: *Games, economic dynamics, and time series analysis*, Springer, pp. 101–116, doi:10.1007/978-3-662-41533-7_6.
- [60] Itai Sher (2020): *Normative Aspects of Kantian Equilibrium*. *Erasmus Journal for Philosophy and Economics* 13(2), pp. 43–84, doi:10.23941/ejpe.v13i2.514.
- [61] Yoav Shoham & Kevin Leyton-Brown (2009): *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press.
- [62] Jaime Simão Sichman & Rosaria Conte (2002): *Multi-agent dependence by dependence graphs*. In: *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 1*, pp. 483–490, doi:10.1145/544741.544855.
- [63] Herbert Alexander Simon (1997): *Models of bounded rationality: Empirically grounded economic reason*. 3, M.I.T. Press, doi:10.7551/mitpress/4711.001.0001.
- [64] Robert Sugden (2003): *The logic of team reasoning*. *Philosophical explorations* 6(3), pp. 165–181, doi:10.1080/10002003098538748.
- [65] William J Talbott (1998): *Why We Need a Moral Equilibrium Theory*. In P. Danielson, editor: *Modeling Rationality, Morality and Evolution*, Oxford University Press.
- [66] Moshe Tennenholtz (2004): *Program equilibrium*. *Games and Economic Behavior* 49(2), pp. 363–373, doi:10.1016/j.geb.2004.02.002.
- [67] Fernando A Tohmé & Ignacio D Viglizzo (2019): *Structural relations of symmetry among players in strategic games*. *International Journal of General Systems* 48(4), pp. 443–461, doi:10.1080/03081079.2019.1573228.
- [68] Michael Tomasello (2009): *Why we cooperate*. M.I.T. Press, doi:10.7551/mitpress/8470.001.0001.
- [69] Michael Tomasello (2016): *A natural history of human morality*. Harvard University Press, doi:10.4159/9780674915855.