

Kantian Kantian Optimization

MATTHEW BRAHAM

Universität Hamburg

MARTIN VAN HEES

Vrije Universiteit Amsterdam

I. INTRODUCTION

In their *Theory of Games and Economic Behavior*, von Neumann and Morgenstern refer to a game-theoretic solution as a “standard of behavior” ([1944] 1953, 41). If we apply this description to all game-theoretic solution concepts and interpret the notion of a ‘standard of behavior’ as a *norm*, we can say that game theory is the study of norm-constrained behavior. Given the rich variety of ethical theories of norms, it is surprising that so much of game theory is dominated by one particular norm: Nash equilibrium. Other solution concepts may describe other norms which are worth studying.

In *How We Cooperate: A Theory of Kantian Optimization*, John Roemer (2019) sets out to develop and defend an alternative account of norm-constrained behavior. His idea is to apply Kantian moral reasoning to provide us with a new theory of social cooperation. Specifically, he tries to make use of the Categorical Imperative (CI) in an optimization model that guarantees mutually beneficial states of affairs in archetypical social dilemmas such as recycling, volunteering in times of war (‘doing one’s part’), soldiers protecting comrades in battle, voting, paying taxes, tipping, and charitable giving. Basically, Roemer wants to model and explain under what conditions we can solve the two major problems that afflict Nash equilibrium (16). These are the tragedy of the commons and the free-rider problem, which concern inefficiencies in the presence of negative and positive externalities, respectively. Roemer’s central thesis is:

AUTHORS’ NOTE: We are very grateful to the editors for their helpful comments and suggestions.

Kantian optimization ‘solves’ what must appear as the two greatest failures of Nash optimization, from the viewpoint of human welfare. (16)

Methodologically, Roemer achieves his result not by tampering with the canonical concepts of preferences and utility of rational choice theory that underpin economics, but rather by distinguishing between two different kinds of optimization strategy that are possible, given the structure of rational choice and our preferences. The two types correspond with two different questions we may raise (12):

Nash optimizer. “Given the strategy chosen by my opponent, what is the best strategy for me?”

Kantian optimizer. “What is the strategy I would like both of us to play?”

Roemer argues that the answers to these questions will usually differ. For him, social cooperation is all about the latter form of optimization. He grounds this distinction in a synthesis of work in evolutionary psychology and linguistics (Tomasello 2014, 2016) and social ontology (Gilbert 1990; Bratman 1992). The idea that emerges from this synthesis is that as a fact of evolution, humans are a “cooperative species” (1), meaning that we have evolved natural capacities ‘to do things together’ and to understand the value of doing so. We are able to form complex systems of language, behavior, and social interaction through which we can share a “union of interests” (4) and according to which we are able to judge that doing things in this way is both individually and mutually advantageous. Thus, when faced with social dilemmas, our instincts and our thoughts are not necessarily Nash optimized at all. Rather, we may understand that Kantian thinking and optimization can be more advantageous.

The heart of Roemer’s claim is that what makes Kantian optimization specifically ‘Kantian’ is that it sufficiently resembles the fundamental feature of the CI, namely, *universalization*. In Roemer’s phrasing:

Take those actions you would will be universalized. (13, emphasis added)

Roemer does not require that such a universalization be governed by altruism; in fact, far from it. The trick, rather, is to pull social cooperation out of an individualistic hat. The Kantian optimizer is still an individ-

ualistic being to the extent that she optimizes the choice of a common strategy—a strategy which is played by all—but which would be best for her. The Kantian optimizer does not consider the payoffs of others and each agent need only know their own preferences. What is required, however, is that each individual expects others to behave in a like manner and that this is based on trust or past experiences (13).

For our contribution to this symposium, we shall ignore Roemer's contribution to the study of human cooperation. The alternative solution concept that he develops is a serious challenge to orthodoxy in economics and game theory. The rigorous formal analysis as well as its application to market economies makes it a profound contribution to both normative economics and formal ethics. However, rather than expanding upon the relevance of the analysis, or its relation to other game-theoretic unorthodoxies—a discussion of the relation with models of team reasoning (Bacharach 2006) is regrettably missing—we will hone in on its theoretical embedding and, in particular, on its 'Kantian' credentials.

Roemer's use of 'Kantianism' follows an established tradition among economists. According to this tradition, an agent follows Kantian morality if she acts under the assumption that others will do the same thing that she does, and if she tries to maximize her utility under that constraint. Economists frequently overlook the fact that this interpretation of Kant differs from the core of Kantian ethical theory: the CI is about the universalization of an agent's maxims rather than her actions and it does not refer to utility maximization. Although Roemer admits that we should not afford too much importance to the reference to Kant (13), it would be interesting to examine whether his new solution concept can be grounded in such a way that it becomes compatible with the Kantian perspective in ethics.

We shall suggest a way of strengthening the Kantian pedigree of Roemer's approach. First, in section II, we will focus on the distinction between actions and maxims and explain the importance of that distinction for the Kantian perspective. It is true that Kantian optimization sometimes yields the same conclusion as a run-of-the-mill application of Kant's CI, but it can also yield both false positives (a defence of immoral behavior) and false negatives (the rejection of permissible behavior). Partly drawing on our earlier work, we then (in section III) give an interpretation of maxims that brings Roemer's analysis closer to Kantian ethics. Finally, we wrap up our analysis with a short conclusion (section IV).

II. IS KANTIAN OPTIMIZATION KANTIAN?

Roemer's main reason for attributing Kantianism to his solution concept is its reference to universalization.¹ As said, Roemer is quite aware that his theory is only loosely 'Kantian', but he nevertheless takes the Kantian optimization condition to be a "natural interpretation" of Kant's CI (viii). Also, he chooses a "Kantian' nomenclature" because "there is a history of using it in economics" (13).²

Of course, in many cases the application of a standard Kantian argument will yield an outcome that coincides with Roemer's analysis. To abide with the usual moral injunctions against theft, deceit, murder, etc., is to adopt a course of action that we all strictly prefer to one in which everyone is willing to transgress those norms. Moreover, to follow such norms means to not be tempted to change our behavior if it happens to improve our personal situation: neither the 'white lie' nor the 'perfect fraud' is an option in Kantian morality.

A fundamental difference between Kant and Roemer is that, for Kant, to act morally is to act autonomously in accordance with those *maxims* that satisfy the 'Moral Law' (for which the CI is a test). Roemer is not concerned with maxims but with actions. Moreover, in Roemer's account, a course of action is moral if we derive some advantage from it being universally adopted. This is more akin to the thought of Hobbes or Hume than it is to that of Kant. For Kant, 'advantage' at best plays an indirect role in morality; that is, it may simply make it easier to follow the moral law (Kant [1797] 1996, 519, 6:388).

Roughly speaking, Kant's CI examines whether the *underlying* reason of a person's action—the *maxim*—is one to which everyone could possibly subscribe. It is a twofold test. In the first step, the CI checks for the existence of a possible world in which everyone could act on the basis of that maxim. If such a possible world does indeed exist, a second step checks to see if an agent that adopts the maxim *can will* that world into existence. If so, acting on the basis of that maxim is morally admissible and not so otherwise.

Thus formulated, the CI is notoriously ambiguous and philosophers have spilled copious amounts of ink in their efforts to interpret it. Yet

¹ To simplify the presentation, we restrict our analysis to games with a common diagonal, allowing us to take simple Kantian equilibrium to be the relevant solution concept (cf. 23, Proposition 2.1).

² In an endnote to page 13, Roemer actually suggests that his approach is closer to Kant's 'Hypothetical Imperative' and that his use of the term 'Kantian' is "for its suggestive meaning and [I] do not wish to imply that there is a deeper, Kantian justification of my proposal" (220n7).

despite this ambiguity, two features of the CI stand out: (a) it tests one's maxims rather than one's actions; (b) it focuses on the *possibility* of everyone acting on the same maxim rather than on the *advantages* that we may derive from it.

Same Action, Different Maxims

One implication of Kant's focus on maxims is that one and the same action can be appraised differently depending on how its underlying maxim is formulated. Roemer's Kantian optimization does not capture this distinctive feature of Kantian morality. The relevance of this can be illustrated by way of the tragedy of the commons. If we exclusively focus on actions, then the farmer who brings his herd to the overgrazed commons because he needs extra earnings, however meagre they may be, to care for the well-being of his family presumably acts in accordance with the moral law. Here Kantian optimization leads to a false negative: the behavior is incorrectly condemned as a wrong. It would be correctly so rejected, if, say, the farmer lets his herd graze because he wants to make an extra buck regardless of the circumstances.

The possibility of false negatives makes it rather clear that we cannot always condemn or blame an agent for not being 'cooperative'. False positives are also possible. The original illustration of the Prisoner's Dilemma game—the case with two prisoners who are offered a deal by the District Attorney—can serve as an example of such a false positive. Assume the prisoners are members of a criminal gang and have in fact committed the crimes they are accused of. Acting on the basis of the maxim of loyalty, they both deny their guilt (play 'Cooperate'). Thus, they both play Kantian equilibrium strategies but theirs is not a play that a Kantian would be likely to endorse.

One can object, of course, that we can ignore these false negatives and false positives because their possibility merely underscores that we have to be careful in describing the moral context of the game at hand. A game to which we apply the Kantian solution concept—the Prisoner's Dilemma in the current discussion—is assumed to describe what we have called elsewhere the *moral field*: it specifies all the relevant moral features of the situation (Braham and van Hees 2012, 611). To refer to the criminal nature of the organization or to the actions that led to the prisoners' arrest means bringing in morally relevant features that are not captured by the game at hand. Yet, if we were to expand the scope of the game, and thus the moral field, we may very well see that within the resulting 'larger' game of which the Prisoner's Dilemma forms a part, cooperation between the two

	L	F
L	(0, 0)	(3, 3)
F	(-1, -1)	(0, 0)

Table 1: The Tango game.

prisoners may fail to form a Kantian equilibrium and therefore may not be justifiable. However, whereas such an expansion may indeed rule out some incorrect judgements, without further argument we cannot be sure that it will *always* do so.

Same Maxim, Different Actions

A different problem arises from the possibility that a maxim can be associated with different actions. This is the case if an individual can act upon a maxim in different ways, but also when the very same maxim corresponds with different actions for different individuals. The latter occurs, for instance, if the provision of a public good requires different inputs from different individuals because it necessitates a division of labour. The optimal outcome will then result only if all individuals act differently and within their domain of expertise. This poses no problem for Kantian morality (if all agents intend to bring about a public good) but it may complicate the application of Kantian optimization.

To see this, consider what we call the ‘Tango game’ (Table 1); a two-person game in which the players have two strategies, *Lead* (*L*) and *Follow* (*F*). The row player specializes in *L* while the column player specializes in *F*. The players’ respective utility functions are the same. The worst outcome ensues if they both try to perform the role they were not specialized to do, while the two next preferred outcomes are those in which one of them deviates from her specialization, and their most preferred outcome is the one in which they both act on the basis of their specialization.

The Pareto-optimal play (*L, F*) is the unique Nash equilibrium, whereas the two Kantian equilibria are suboptimal. In the Tango game, Kantian optimization thus generates a false positive (‘act the same way’). Being ‘cooperative’ in this context, however, means to ‘act differently’.

One could argue that this observation is a mere semantic sleight of hand in that we are simply mis-describing the actions. Suppose we re-describe the players’ actions as *Specialization* (*S*) and *Non-specialization* (*N*), respectively. This leads to the ‘Modified Tango game’ (Table 2).

Now, the problem vanishes and Kantian optimization yields the morally desirable outcome. Just as the problems following from multiple maxims suggested a move to a different type of modelling, so too the problem

	S	N
S	(3, 3)	(0, 0)
N	(0, 0)	(-1, -1)

Table 2: The Modified Tango game.

arising from multiple actions instantiating the very same maxim may be solved via re-modelling. But while this may indeed work for this particular game, it is a somewhat ad hoc solution. Why would the second game, rather than the first, describe the situation correctly? Can we simply decide how to describe the agents' actions? As Roemer notes (28), Kantian optimization requires specifying when different individuals' strategies are the same, which may not always be obvious.

Preferences and Morality

Roemer emphasizes (13) that the Kantian optimizer is only trying to realize her own preferences as well as possible. While being altruistic is compatible with Kantian optimization, it is not at all necessary for it. Crucial is that the realized outcome be the most preferred one, the character of the preferences themselves is not relevant. The difference between Kantian and Nash equilibrium lies in the comparison between their respective outcomes and not with the way in which they themselves are compared, which is preference-based. Yet the preference-based comparison does not square well with an essential characteristic of Kantian ethics. For Kant, the CI is about the *possibility* of universalizing a maxim. It is not about the desirability of the consequences of universally adopted actions or maxims.

III. RE-KANTING ROEMER

Roemer's project suggests the value of bringing Kantianism and the welfare consequentialism of economics closer together. But, can we bring them even closer? That is, is there a game-theoretic analysis of the CI that focuses on maxims but which draws on Roemer's interpretation of Kantian optimization? We shall argue that such an analysis is indeed possible.³

To do so we have to unpack the formulation of the CI—that is closest to Roemer's idea of Kantianism—which is known as the *Formula of Universal Law* (FUL) version of the CI:

³ Here, we will use some ideas from our earlier work (Braham and van Hees 2015) but which deviate from it in the way we connect maxims with preferences. The latter idea is motivated by Amartya Sen's (1974) early suggestion to model morality in terms of meta-rankings.

FUL. Act only in accordance with that maxim through which you can at the same time will that it become a universal law. (Kant [1785] 1996, 73, 4:421)⁴

FUL has two constituent parts: the aforementioned maxims and the two tests of the universalizability of the maxims. We start with the concept of a maxim. One option, which we adopt, is to view a maxim as a rule of conduct that refers to a person's intentions across a range of circumstances.⁵ We then take a maxim to be about which states of affairs are to be picked out by conduct whenever certain circumstances arise. Regimenting it a little, a maxim is a tripartite relation of the form:

an agent will do α if β in order to ϕ ,

where α ranges over actions, β over circumstances, and ϕ over states of affairs.

According to such a conception, a maxim consists of two intentions: an *act*-intention (performance of some action α) and an *outcome*-intention (realization of some state of affairs ϕ). Here, we can use the revealed-preference interpretation of utility functions according to which an agent's preference describes a hypothetical choice that she faces. These hypothetical choices can in turn be understood as describing her intentions: they specify a certain choice (α) for a state of affairs (ϕ) in a possible choice situation (β). By this interpretation '*i* prefers x over y ' means 'if the agent were to have a choice between x and y , then she would choose x ', which in turn can be interpreted as '*i* intends to choose x if the choice is between x and y '. Taking a maxim to be a collection of such conditional intentions, we arrive at a conceptual link between Kantian maxims and utility functions. In Kantian ethics the admissibility of people's behavior does not depend on the assessment of the outcomes of their actions but on the motivation underlying their behavior. By interpreting preferences and the corresponding utility functions in terms of that motivation rather than in terms of an assessment, we can apply the economic apparatus.

⁴ As is well-known, Kant also provided us with a number of other less formalistic formulations in the *Groundwork of the Metaphysics of Morals*, which he believed to be equivalent. These are known as the formulas of 'Humanity', of 'Autonomy', and 'Kingdom of Ends'.

⁵ This conception goes back to O'Neill (1975, 34–42). See also Westphal (2011, 111). There are other possible interpretations of this approach. For a recent and comprehensive analysis of what Kantian maxims are, see Herissone-Kelly (2018).

	t_1	t_2
s_1	x	y
s_2	z	v

Table 3: An example game form h .

Indeed, whereas such a route may require a stretch of imagination for the conventional Kantian, it does fit neatly within economic theory. To indicate the outlines of such ‘Kantian economics’, let us start with a game form h that models the situation that contains all the ingredients of a game (players, strategies, outcomes) except for the preferences of the players. A game g is, then, defined as a game form h plus a preference profile $\mathbf{u} = (u_1, \dots, u_n)$. Say the game form h is as in Table 3.

Thus the game form is ‘part of’ a Prisoner’s Dilemma if the outcomes correspond with the time that each prisoner has to spend in prison and if the prisoners have the intention to reduce their prison time as much as possible. In a ‘Battle of the Sexes’ game, the outcomes describe ways of spending the evening and the partners intend to be together although both would opt for different things to do together, etc. We denote a utility function of an agent i that is associated with a particular maxim m as $u_{i,m}$. Accordingly, a profile in which each individual utility function is associated with the same maxim m is denoted as \mathbf{u}_m .

We can now turn to the formulation of the CI. We will focus only on that part of it that is closest to Roemer’s analysis and which Christine Korsgaard (1996, 93) calls the ‘Practical Contradiction Interpretation’.⁶ In doing so we sidestep the issue of ascertaining which maxims can be universally adopted at all, and simply assume that the information is given exogenously.

We say that a practical contradiction arises if the universal adoption of a maxim would lead to a state of affairs that is at odds with the maxim. To see how this works, let h be the game form and let \mathcal{M} be the non-empty set of all maxims that can be adopted universally in h . For any $m \in \mathcal{M}$, g_m denotes the game (h, \mathbf{u}_m) that describes the universal adoption of m . To simplify the analysis, we assume that the game g_m associated with the universal adoption of a maxim m always has a unique and pure Nash equilibrium, the outcome of which is denoted by x_m^* .⁷

⁶ The two tests that the CI is taken to comprise are commonly referred to as the Contradiction in Conception (CC) and the Contradiction in the Will (CW) tests (O’Neill 1975). For our interpretation of the CI, we draw upon our earlier work (Braham and van Hees 2015). Note, however, that what we take to be the CW test is interpreted as the CC test by Korsgaard.

⁷ The assumption simplifies the presentation because it avoids the need to introduce preferences over lotteries or set-preferences.

Whereas a maxim in \mathcal{M} is one that *can* be a universal law, this does not yet mean that we can rationally *will* it to become a universal law. In the *Groundwork of the Metaphysics of Morals*, Kant expands on FUL as follows:

Some actions are so constituted that their maxim cannot even be *thought* without contradiction as a universal law of nature, far less could one *will* that it *should* become such. In the case of others that inner impossibility is indeed not to be found, but it is still impossible to *will* that their maxim be raised to the universality of a law of nature because such a will would contradict itself. (Kant [1785] 1996, 75, 4:424)

Given our interpretation of individual preferences, we say that a rational agent *can will* the universal adoption of a maxim m if, and only if, the outcome x_m^* resulting from the universal adoption of the maxim m is, according to m , indeed the outcome that he intends to choose in any pairwise comparison with the outcome resulting from the universal adoption of any other maxim. Or, more succinctly, it rules out the possibility of a rational agent acting on the basis of a maxim m that tells her not to act on it if everyone were to do so. This possibility is the conflict within a person's will that we take to be excluded by Kant's FUL.

We can illustrate this framework with a two-person version of the tragedy of the commons. Assume that the farmers have only two maxims available to them, which for simplicity's sake, we call *individual* and *collective*. The *individual* maxim is 'unconditional self-interest'; the *collective* maxim conditions behavior on the social optimum. Assuming with Kant ([1785] 1996, 74-75, 4:423) that universal self-interested behavior is feasible, their adoption of the individual maxim yields the Prisoner's Dilemma, whereas a game in which the cooperative outcome is the only Nash equilibrium would result if they both chose the collective maxim. Next we construct games in which the players adopt maxims. To do so, we use the notion of a Kantian game form.

Kantian game form. Given a game form h with associated \mathcal{M} , the Kantian game form \hat{h} is a particular game form in which:

1. Each individual strategy set is \mathcal{M} .
2. The outcome of a play $(m, \dots, m) \in \mathcal{M} \times \dots \times \mathcal{M}$ is x_m^* .

Each combination of a Kantian game form and a preference profile \mathbf{u}_m associated with a maxim m yields a unique game (\hat{h}, \mathbf{u}_m) . We can now

bring Roemer's idea of Kantian optimization into harmony with Kantian ethical theory. We call this '*Kantian Kantian Optimization*', and formulate it as follows:

Kantian Kantian optimization. Given a Kantian game form \hat{h} , acting on the basis of maxim m is morally admissible if, and only if, (a) $m \in \mathcal{M}$, and (b) the play (m, \dots, m) is a simple Kantian equilibrium of the game (\hat{h}, \mathbf{u}_m) .

Using again the tragedy of the commons illustration, let $\mathcal{M} = \{m^i, m^c\}$ and let (h, \mathbf{u}_{m^i}) be a Prisoner's Dilemma game with the unique Nash equilibrium (D, D) of both players defecting. Conversely, a universal adoption of the collective maxim yields a 'Prisoner's Harmony' game (h, \mathbf{u}_{m^c}) with a unique Nash equilibrium (C, C) of both players cooperating. Turning to the Kantian game form \hat{h} in which the strategies of the agents are m^i and m^c , we next examine the agents' assessments of the various games resulting from the universal adoption of a maxim. We see that in terms of their Prisoner's-Dilemma preferences (describing maxim m^i), the Prisoner's-Harmony outcome $x_{m^c}^*$ is ranked higher than the Prisoner's-Dilemma outcome $x_{m^i}^*$. Hence, (m^i, m^i) is not a simple Kantian equilibrium of the game $(\hat{h}, \mathbf{u}_{m^i})$. It is for that reason that acting on the basis of the individual maxim is said to be inadmissible. On the other hand, the play (m^c, m^c) is a simple Kantian equilibrium in the game in which m^c rather than m^i describes the preferences of the agents, that is, in $(\hat{h}, \mathbf{u}_{m^c})$. Acting on the basis of m^c is, therefore, admissible.

IV. CONCLUSION

We have suggested that Roemer's account of Kantian optimization can be brought closer to Kantian ethical theory by making certain suitable assumptions about the interpretation of maxims. Kantian optimization can then be seen as forming a solution of particular games, namely, games in which agents choose maxims on the basis of which they will act. Since such games have a very specific nature which do not coincide with the games in which Roemer analyses Kantian optimization, we called the resulting account '*Kantian Kantian optimization*'.

How exactly does this strengthen the Kantian pedigree of Roemer's solution concept? By definition, *Kantian Kantian optimization* is an instance of Kantian optimization, but the converse need not always be true. Kantian optimization may fail to be truly Kantian because there may be no maxim that, if universally adopted, would lead to the cooperative out-

come. A natural next step in our analysis would be to examine the class of games and maxims in which Kantian optimization reaches the same verdict as *Kantian* Kantian optimization. This is important for welfare economists as it will provide guidance as to how to achieve socially optimal outcomes in the morally right kind of way—which is what Kantian morality is all about. Thereby, it will introduce a dimension of moral proceduralism that is generally lacking. In this way, Roemer's *How We Cooperate* has opened up a wholly new avenue of theoretical possibilities.

REFERENCES

- Bacharach, Michael. 2006. *Beyond Individual Choice: Teams and Frames in Game Theory*. Edited by Natalie Gold and Robert Sugden. Princeton, NJ: Princeton University Press.
- Braham, Matthew, and Martin van Hees. 2012. "An Anatomy of Moral Responsibility." *Mind* 121 (483): 601–634.
- Braham, Matthew, and Martin van Hees. 2015. "The Formula of Universal Law: A Reconstruction." *Erkenntnis* 80 (2): 243–260.
- Bratman, Michael E. 1992. "Shared Cooperative Activity." *The Philosophical Review* 101 (2): 327–341.
- Gilbert, Margaret. 1990. "Walking Together: A Paradigmatic Social Phenomenon." *Midwest Studies in Philosophy* 15 (1): 1–14.
- Herissone-Kelly, Peter. 2018. *Kant on Maxims and Moral Motivation: A New Interpretation*. Cham: Springer.
- Kant, Immanuel. (1785) 1996. "Groundwork of the Metaphysics of Morals." In *Practical Philosophy*, edited by Mary J. Gregor, 37–108. Cambridge: Cambridge University Press.
- Kant, Immanuel. (1797) 1996. "The Metaphysics of Morals." In *Practical Philosophy*, edited by Mary J. Gregor, 353–604. Cambridge: Cambridge University Press.
- Korsgaard, Christine M. 1996. *Creating the Kingdom of Ends*. Cambridge: Cambridge University Press.
- O'Neill, Onora. 1975. *Acting on Principle: An Essay on Kantian Ethics*. New York, NY: Columbia University Press.
- Roemer, John E. 2019. *How We Cooperate: A Theory of Kantian Optimization*. New Haven, CT: Yale University Press.
- Sen, Amartya K. 1974. "Choice, Orderings and Morality." In *Practical Reason: Papers and Discussions*, edited by Stephen Körner, 54–67. Oxford: Blackwell.
- Tomasello, Michael. 2014. *A Natural History of Human Thinking*. Cambridge, MA: Harvard University Press.
- Tomasello, Michael. 2016. *A Natural History of Human Morality*. Cambridge, MA: Harvard University Press.
- von Neumann, John, and Oskar Morgenstern. (1944) 1953. *Theory of Games and Economic Behavior*. 3rd edition. Princeton, NJ: Princeton University Press.
- Westphal, Kenneth R. 2011. "Practical Reason: Categorical Imperative, Maxims, Laws." In *Immanuel Kant: Key Concepts*, edited by Will Dudley and Kristina Engelhard, 103–119. Durham: Acumen.

Matthew Braham is Professor of Practical Philosophy at Universität Hamburg, Germany. He is Deputy Speaker of the Graduate Program in Collective Decision-Making. He has published in moral and political philosophy, social choice, and game theory.

Contact e-mail: <matthew.braham@uni-hamburg.de>

Martin van Hees is Professor of Ethics at VU Amsterdam. His research interests concern the foundations of (moral and political) liberalism and he has worked on theories of freedom, the analysis of moral responsibility, and quality of life assessments. He is a former editor of *Economics & Philosophy*.

Contact e-mail: <m.van.hees@vu.nl>