6th edition of

Formal Ethics

Book of abstracts

Date: June 19-21, 2019
Venue: Royal Academy of Dutch Language and Literature (KANTL)

Organization:

Centre for Logic and Philosophy of Science
(Ghent University)

Chairs
Frederik Van De Putte (Ghent & Bayreuth) & Federico Faroldi (Ghent & Frankfurt)

Local Organizing Committee
Thijs De Coninck, Federico Faroldi, Stef Frijters, Kristian González Barman, Bert Leuridan, Karina Makhnev, Joke Meheus, Frederik Van De Putte & Nathan Wood.

About the workshop
Formal Ethics is a common denominator for the application of tools from logic, decision theory, game theory, and social choice theory to the analysis of concepts in moral and political philosophy and to the development of ethical theory. This edition of Formal Ethics is sponsored by the Flemish Research Foundation (FWO-Vlaanderen), the Special Research Fund of Ghent University (BOF UGent) and by the European Research Council (ERC), through the Marie Skłodowska Curie project DYCODE.

Abstracts of Invited Talks

Ray Briggs
What Are the Options?

According to traditional expected utility theory, you rationally ought to choose the option with the greatest expected utility. But in some decision situations, choosing following this advice looks not rational at all, but short-sighted. To borrow a famous example from Michael Bratman, after a post-dinner glass of wine, the utility of a second glass of wine might temporarily increase, so that you feel tempted to go back on your policy of drinking one glass at a time. Or to borrow another famous example from Derek Parfit, you may find it appealing to go back on your promises, once the advantage of doing what you promised has evaporated. We are tempted to act in ways that benefit us in the short term and cost us in the long term, and to defect in prisoner's dilemma games and tragedies of the commons in ways that disadvantage everyone. I defend the view that we should respond to problems of temptation by rethinking the concept of a decision problem, and in particular, the concept of an option. Traditionally, decision theory has it that options are given before deliberation begins, that they are mutually exclusive (you can carry out at most one of them) and that they are jointly exhaustive (you must carry out at least one of them). I argue for a much more expansive view of the options. Your options include both acts that you can perform now, and policies or strategies that you can participate in, but that are ultimately carried out over time, or by groups of people. This means that your options are many, and may stand in interesting logical relations of inclusion or overlap. In Bratman's wine example, for instance, one of your options is to drink a second glass of wine right now, another is to refuse a second glass of wine right now, and a third option, distinct from both, is to carry out your long-term policy of never drinking more than one glass of wine after dinner. And on at least some ways of filling in the detail, carrying out this policy is the right option.

Campbell Brown
The Valence of Lives

A life may be either good, bad, or neutral. This is the "valence" of the life. Valence is an important consideration in many problems of Population Ethics. For example, whether or not it is permissible to create a new person plausibly depends, at least in part, on whether this person would have a good life, or a "life worth living". In this talk, I evaluate some approaches to defining valence. I focus in particular on what may be called "temporal" approaches. A life has a temporal structure: it persists and changes over time. And its value, too, may change over time: it may get better or worse. Temporal approaches exploit this structure in order to define valence. I argue that such definitions face serious problems, which may be avoided only by restricting the "domain" of the definition. This restriction has the effect of excluding certain substantive theories of well-being. However, as I argue, this may be a tolerable consequence.

Edith Elkind
Justified Representation: From Axioms to Algorithms

Suppose that a group of voters wants to select $k > 1$ alternatives from a given set, and each voter indicates which of the alternatives are acceptable to her: the alternatives could be conference submissions, applicants for a scholarship or locations for a fast food chain. In this setting it is natural to require that the winning set represents the voters fairly, in the sense that large groups of voters with similar preferences have at least some of their approved

alternatives in the winning set. We describe several ways to formalize this idea, and show how to use it to classify voting rules; surprisingly, two voting rules proposed in the XIXth century turn out to play an important role in our analysis.

Abstracts of Contributed Talks

Jon Marc Asper
Subjective Values should be Sharp

Adam Elga (2010) has argued that even when no particular subjective probability is required by one's evidence, perfectly rational people will have sharp subjective probabilities. Otherwise, they would be rationally permitted to knowingly turn down some sure gains. In mostly a parallel manner, I argue that even when one's practical reasons are underdetermined (due to optionality or ignorance), perfectly rational people will have sharp subjective values. Some might be inclined to reject this application of Elga's argument to subjective values for at least two kinds of reasons. First, some might object that our evaluations should match the (purported) unsharpness of objective values (e.g., clumpy value, vague value, or parity). I respond that a requirement for sharp subjective values is consistent with some (although not all) purported kinds of unsharp objective values. Second, some might object that value sharpenings would be arbitrary and thus should not be action-guiding—one's 'real' (i.e., practically authoritative) evaluations would remain unsharp. I respond that sharpened subjective values must be 'real' because otherwise perfect rationality would be impossible in principle.

György Barabás & Andras Szigeti
Collectivizing Justice: A Novel Argument for Quota-Based Affirmative Action

It seems reasonable to assume that if a workplace hires its employees from representatives of two equally numerous, equally qualified groups (e.g., women and men), and does so in a group-blind manner, then the workplace representation of those groups will equalize in the long run. However, this intuition relies on the tacit assumption that each hired employee is equally likely to remain at the workplace. In reality, members of the minority group are often more likely to quit prematurely, for reasons of marginalization caused by being a minority at work. Using a simple mathematical model, we show that taking this into account generates certain "points of recalcitrance", preventing the minority group from reaching equal representation. The surprising upshot is that a hiring policy which is in principle fair can conserve the minority status of the smaller group, and this will remain the case even if all relevant stakeholders are well-intentioned and committed to the aims of affirmative action. We see this finding as strongly supporting quota-based affirmative action, because it shows that only preferential hiring of the minority group will make it possible to push the intragroup distribution past the points of recalcitrance.

Constanze Binder
Moral Responsibility and Individual Choice Sets.

The agency paradox shows that a rational person often cannot be held morally responsible for her actions, if the availability of a reasonable alternative in a person's choice set is a necessary condition for moral responsibility (Binder and van Hees 2015). In this article we explore a route to avoid the agency paradox by taking moral responsibility as a criterion for the assessment of the quality of choice sets rather than a consideration to account

for in a person's decision making. More specifically, we identify a universal set of reasonable alternatives and argue that a person can be held morally responsible for all choices made from a choice set whose intersection with the universal set of reasonable alternatives is non-empty. We show how the focus on choice sets instead of decisions allows to avoid the agency paradox. We conclude with a discussion of possible problems that can occur if the reasonableness of alternative options is taken to be context- or choice-set independent.

JUSTIN BRUNER
Inequality and Majority Rule

While intuitive, majority rule (MR) has a rather peculiar status in democratic theory. Although MR is thought to be compatible with democracy and political equality, a variety of authors have argued that MR is not unique in this regard. This realization has in part prompted serious interest in familiar alternatives to MR as well as less popular voting schemes (lottery voting). I provide a novel defence of MR. As is the case when groups must collectively decide, some portion of the voting body is frustrated while others 'get what they want'. I consider the distribution of voter satisfaction in response to the outcome of a vote and show that under certain conditions MR minimizes the level of inequality present in the distribution of voter satisfaction. Thus, there is a sense in which MR ensures equality at both the level of procedure and outcome.

ILARIA CANAVOTTO & ALESSANDRO GIORDANI
Causation and Accountability in Dynamic Action Logic

While several deontic and action logics have recently been advanced to clarify the notion of mens rea and its connection with legal responsibility (see, e.g., Broersen, 2011a and 2011b; De Lima et al., 2010; Lorini et al., 2013), less attention has been payed to the notion of actus reus, and specifically to the distinction between causally contributing to a state of affairs and being accountable for it. The aim of this paper is to make a first step towards filling this gap. To do this, we implement a refined version of the INUS account of causality [Baumgartner, 2013] in a dynamic setting, with actions interpreted as transitions, equipped with a distinction between default and deviant conducts (see Halpern and Hitchcock, 2015). By focusing on criminal cases involving medical negligence, we then define a procedure to formally elucidate the legal principle that the defendant is accountable for a result if her act was an operating and substantial cause of it. In the last part of the paper, we investigate to what extent the proposed procedure constitutes a general method to identify guilty acts in scenarios involving sequences of actions of different agents.

NICOLAS COTE
Liberalism and Social Choice

One of the core commitments of liberalism is to the protection of individuals against intolerable intrusions upon their liberty. It enshrines these protections through rights, which safeguard us from coercion. Sen's celebrated "Impossibility of a Paretian Liberal" result proved that welfarism, as expressed by the Pareto principle, is inconsistent with even the most minimal concern for liberal rights. However, though Sen is acknowledged to have shown the existence of a conflict between liberalism and welfarism, his conception of liberal rights generates paradoxes. I present a novel approach to the representation of rights in the language of social choice, one which avoids the paradoxes of Sen's original view, while allowing us to express a whole range of liberal commitments. Most notably,

my framework allows us to state precisely how liberal rights constrain welfarism, and even to define meaningful measures of illiberalism, allowing us to rank social states with respect to how deeply illiberal they are.

Mike Deigan
Rational Partiality and Objective Value.

I'm inclined to think it's a rational requirement to prefer outcomes with greater expected objective value to those with less. But I'm also inclined to think it can be rational to have preferences partial towards oneself and one's loved ones. Can these be consistently held together? I argue that they can be, but only if we take outcomes to include self-locating information (and so model them with centered worlds rather than worlds) and objective value to be something had by individuals, rather than worlds.

Hein Duijf, Allard Tamminga & Frederik Van De Putte
An Impossibility Result on Methodological Individualism

We prove that the statement that a group of individual agents performs a deontically admissible group action cannot be expressed in a well-established multi-modal deontic logic of agency that models every conceivable combination of actions, omissions, abilities, obligations, prohibitions, and permissions of finitely many individual agents. Our formal result has repercussions for methodological individualism, the methodological precept that any social phenomenon is to be explained ultimately in terms of the actions and interactions of individuals. We show that by making the relevant social and individualistic languages fully explicit and mathematically precise, we can precisely determine the situations in which methodological individualism is tenable.

Tomi Francis
A Strengthened Impossibility Theorem in Population Axiology

Population axiology is the study of betterness, in contexts where the number and identities of those individuals who ever exist can change depending on which action is chosen. A chief concern of population axiology is to avoid the "Repugnant Conclusion", which says that for any population consisting of excellent lives, there is some, much larger, better population containing only lives barely worth living. However, a number of impossibility theorems in the literature show that there several combinations of highly plausible premises which jointly entail the Repugnant Conclusion. Most such theorems rely on weakenings of the "Mere Addition Principle", which says adding good lives makes an outcome at least as good as before. But I argue that these principles are not significantly more compelling than Mere Addition: if we deny Mere Addition, we can already argue on this basis that the logically weaker principles also fail. Despite this, we are not out of the woods. I present an impossibility result for "Repugnant Addition" - a variant of the Repugnant Conclusion - which does not rely on any principle similar to Mere Addition.

Stef Frijters, Joke Meheus & Frederik Van De Putte
Quantifying over the indexes of obligation operators

General deontic statements are of paramount importance in ethical reasoning. Consider 'all parents have an obligation towards their children to take care of them' and 'only doctors are allowed to prescribe drugs.' These statements contain universal quantification over the bearers and counterparties of so-called personal obligations and permissions. The goal of our talk is to present term-modal deontic logic, TMDL, that aims to capture reasoning with

these statements. We will present a sound and strongly complete Hilbert-style axiomatisation of TMDL. TMDL itself is a conservative extension of SDL. However, we show that the term-modal extension is easily generalisable to other deontic logics. In this talk we show that, in spite of their central role in ethical reasoning, many general deontic statements have so far not been treated properly in deontic logic. In contrast, we illustrate that TMDL does capture reasoning with these statements. In addition, we will discuss how TMDL can be used to capture Hohfeldian rights relations. Earlier work on Hohfeldian rights in deontic logic was focussed on paucital rights. We will show that TMDL can capture not only paucital rights, but also classes of multital rights. These classes are central to the original Hohfeldian analysis.

JOBST HEITZIG & SARAH HILLER

Measures of individual and groupwise ex-post and ex-ante responsibility in extensive-form games with unquantifiable uncertainty

Problems like climate change involve heterogeneous agents, complex interactions, diverging information, and various forms of uncertainty. The dominant methodology for such problems is game theory, explaining failure of collective action by collective inefficiency of equilibria arising from selfishness and rationality, giving pessimistic predictions. However, evidence shows that people are also guided by moral principles like responsibility, ignored by game-theoretic analyses. Literature suggests different forms of responsibility, such as being a responsible person or being blameworthy. Here we distinguish ex-ante and (factual or counterfactual) ex-post responsibility, individual vs. collective, absolute vs. gradual, and assigned vs. perceived responsibility, related to ethical obligation, blame, moral luck, and the diffusion of responsibility effect. Since most of these are not sufficiently formalized to be clearly distinguishable, generally applicable, or measurable. We explore an approach that bases related quantitative measures on probabilities and evaluations of possible consequences of actions, properly treating different forms of uncertainties, in particular regarding other agents' behaviour, by extending the data structure of extensive-form games, using paradigmatic situations to derive desirable properties of responsibility measures, and deriving candidate formulae fulfilling these axioms that may serve as a proof of concept for our approach and as a starting point for further discussions.

ALEKS KNOKS & ERIC PACUIT

Reason-Based Choice for Groups

Rational choice theory usually assumes that an agent's preferences over a set of alternatives are fixed, unchangeable, and without structure. Dietrich and List (Nous 47(1), 2013; IJGT 42, 2013) have recently enriched the standard model by explicitly representing an agent's motivational state. In this new model, preferences are subject to change and are determined by the agent's reasons. In my talk I will discuss the results of applying this enriched model to some outstanding questions in group choice. The extra structure of the model allows for a more subtle perspective on (dis)agreement between the agents' preference. And this, in turn, allows for more intelligent conflict resolution. One surprising result of our investigation is that agreement does not always entail consensus.

KACPER KOWALCZYK

Yet another argument against anti-aggregation

Consider Scanlon's (1998) classic case Transmitter Room. One person falls on the cables in the transmitter room as the World Cup final is going on. If you let her be, she will be painfully electrocuted for the rest of the

match. If you help her, the transmission will be cut off for billions of viewers worldwide. What should you do? Anti-aggregationists thinks you should help the one person, regardless of the number of the viewers watching the match. This is because the one's complaint will be much greater than the complaints of the many, and in this kind of case you should let the person with the greatest complaint decide. Cf. Voorhoeve (2014). Aggregationists disagree. I will provide new arguments against anti-aggregationism. First I show that anti aggregationism is incompatible with the following principles:

Individual Continuity: Let x, y, z be wellbeing levels with x higher than y higher than z. Then there is some probability π strictly between zero and one such that getting z with probability π and x with probability $(1 - \pi)$ is better for a given individual than getting y for sure.

Ex Ante Pareto: If everyone's prospects are better in X than in Y, you should choose X rather than Y.

Statewise Dominance: If in each state of nature you should choose X rather than Y, then you should choose X rather than Y.

It is well known that a structurally similar problem arises for egalitarianism, as described in Otsuka & Voorhoeve (2009), for example. Indeed, this problem is closely related to Harsanyi's aggregation theorem, as detailed in Thomas (2016: ch. 3). Nonetheless, I show that anti aggregationism and egalitarianism are distinct doctrines, with different theoretical commitments. I also spell out the character of their similarity. I turn this result into an argument against anti-aggregationism.

CHRISTOPHER MEACHAM

Utilitarianism, Consent, and the Self-Other Asymmetry

A number of authors, including Sider (1993), Portmore (2008) and Vessel (2010), have suggested modifying consequentialist theories like Utilitarianism in order to accommodate a "self-other asymmetry" – a moral difference between imposing burdens on oneself and imposing burdens on others. And they've suggested doing so by adopting theories that employ a disjunctive account of permission. I'll argue that these proposals are problematic for two reasons. First, I'll argue that the disjunctive accounts of permission these approaches employ are implausible. A plausible account should be sensitive to tradeoffs in a way in which these disjunctive proposals are not. I'll then show how to construct a theory that incorporates the self-other asymmetry which is sensitive to such tradeoffs. Second, I'll suggest that it's a mistake to directly appeal to a self-other asymmetry. The morally important distinction is not whether the sacrifice is made by oneself or by others, it is whether the sacrifice is made by subjects who consents to it. I'll then show how to construct a plausible theory that's sensitive to tradeoffs involving consent.

ABELARD PODGORSKI

The Self-Esteem Theorem

In a range of normative domains, there are attractive views about how to evaluate our options against each other which are decision-dependent – where some crucial input in the evaluation may be affected or determined by the decision itself. On one class of views about how to deal with such decision-dependent cases, instead of looking for some privileged perspective from which to evaluate, we begin with a set of decision-relative evaluations, one from the point of view of each action to be performed, and generate an overall evaluation as a function of all of these. In this paper, I prove a result called the Self-Esteem Theorem, which shows that there is only one way to make good on this kind of approach consistent with an appealing version of a principle called the independence of irrelevant alternatives – the overall evaluation of each option must correspond to its evaluation from its own perspective. I go on to describe possible lessons we may draw from this result.

Soroush Rafiee Rad & Olivier Roy
Deliberation, Single-Peakedness, and Voting Cycles

Deliberation can help avoiding voting cycles. This claim is usually made via the notion of "meta-agreement" which, it has been claimed, entails that the group members' preferences are single-peaked, which in turn guarantees the existence of a Condorcet winner. We provide evidence from simulations that this detour through meta-agreement, or even single-peaked preferences, is in general unnecessary. To the extent that deliberation induces a process of rational preference change, it almost completely avoids voting cycles. Whether it does so through the creation of single-peaked preferences depends on whether the participants rank only strictly the alternatives or are allowed to be indifferent. In the former case deliberation does eliminate cycles by creating single-peaked preferences. In the latter case, however, single-peaked preferences play little to no role in the elimination of voting cycles.

Korbinian Rüger
Aggregation and Equality

Moral aggregation is the combining of claims or the well-being of individual people in order to decide what we morally ought to do. There is an influential strand in contemporary moral philosophy that objects to the aggregative element in Utilitarianism and other theories. This kind of scepticism about interpersonal aggregation is a staple of contemporary non-consequentialist moral thought. Many sceptics of aggregation are also on record as Egalitarians. However, usually, issues regarding aggregation are treated separately from equality. I here argue that they should be treated together. Aggregation is most plausibly rejected on egalitarian grounds. This is because the most powerful positive arguments for aggregation rely on premises that seem to deny egalitarianism. If one is committed to Egalitarianism, one will want to reject these premises. And since, as I will argue, these are the weakest premises of the arguments, aggregation can most plausibly be resisted by taking on egalitarian commitments.

Kai Spiekermann
Epistemic Network Injustice

To find out what is in one's own best interest, it is helpful to ask one's epistemic peers. However, identifying one's epistemic peers is not a trivial task. I consider a stylized political setting, an electoral competition of 'Masses' and 'Elites'. To succeed, the Masses need to know which alternative on offer is truly in their interest. To find out, the Masses can pool their privately held information in a pre-election ballot, provided that they can reliably find out with whom they should pool information. I investigate the process of finding the relevant peer group for information pooling by modeling this group formation process as dynamic network change. The simulations show that the Masses can succeed in finding the right peers, but they also suggest reasons why the Elites may often be more successful. This phenomenon generalizes to the notion of epistemic network injustice. Such injustice arises when a subset of citizens is systematically deprived of connections to helpful epistemic peers, leading to their reduced political influence. Epistemic network injustice is a new form of epistemic injustice, related to but distinct from the notion introduced by Miranda Fricker.

Satoru Suzuki
Measurement-Theoretic Foundations of Weighted Utilitarianism

Harsanyi (1955) attempts to develop expected utility theory of von Neumann and Morgenstern (1944) to provide a formalization of (weighted) utilitarianism. Weymark (1991) refers to this result as Harsanyi's Aggregation Theorem. There are at least two well-known criticisms on Harsanyi's Aggregation Theorem. In order to escape these two criticisms, we might resort to Harvey's Aggregation Theorem (1999) that has preference intensity relations as primitive which can be represented by utility differences, and concerns only with preferences on outcomes but does not concern with preferences on lotteries in Harsanyi's Aggregation Theorem. The first aim of this paper is to offer two criticisms on Harvey's Aggregation Theorem from a measurement-theoretic point of view. The second aim is to prove a new theorem, which escapes these two criticisms, inspired by Harvey's Aggregation Theorem. One of key techniques for proving this theorem is a version of Moment Theorem in abstract linear spaces in Domotor (1979).

Christian Tarsney
Vive la Différence? Structural Diversity as a Challenge for Metanormative Theories.

Decision-making under normative uncertainty requires an agent to aggregate the assessments of options given by rival normative theories into a single assessment that tells her what to do in light of her uncertainty. But what if the assessments of rival theories differ not just in their content but in their structure − e.g., some are merely ordinal while others are cardinal? This paper describes and evaluates three general approaches to this "problem of structural diversity": structural enrichment, structural depletion, and multi-stage aggregation. All three approaches have notable drawbacks, but I tentatively defend multi-stage aggregation as least bad of the three.