



Groningen – Wharton PPE Workshop

Monday, Sept. 12, and Tuesday, Sept 13
University Library, Groningen (Tammessaal, 4th floor)

General information

This is a workshop to discuss work in progress from the realm of PPE (Philosophy, Politics, and Economics). Papers will be circulated in advance and participants are expected to read them beforehand. Authors will give a short introduction, followed by a commentary and then a general discussion. The workshop is supported by the Zicklin Center for Business Ethics Research of The Wharton School at the University of Pennsylvania and the Center for Philosophy, Politics and Economics at the University of Groningen. It has been organized by Brian Berkey (bberkey@wharton.upenn.edu), Julian Jonker (jonker@wharton.upenn.edu), and Lisa Herzog (l.m.herzog@rug.nl), and will be held in person at the University of Groningen.

Registration

If you are interested in participating, please contact the local organizer, Lisa Herzog, at l.m.herzog@rug.nl. Places will be allocated on a first come, first serve basis.

Timetable

Monday, Sept. 12, 9-17

- 9.30-10.00 Welcome and round of introductions
- 10.00-11.00 Sanjana Govindarajan and Boudewijn de Bruin (Groningen): Artificial Intelligence and Refugee Status Determination: The Role of Businesses in Mitigating Epistemic Injustice in Asylum Systems
Commentator: Santiago Mejia (Fordham)
- 11.00-11.30 Coffee break
- 11.30-12.30 Robert Hughes (Harvard Safra Center): The Ethical Duty to Refrain from Wage Exploitation
Commentator: Boudewijn de Bruin (Groningen)
- 12.30-13.30 Lunch break (catering)
- 13.30-14.30 Frank Hindriks (Groningen): How Do Social Structures Empower and Disempower?
Commentator: Julian Jonker (Wharton)
- 14.30-15.30 Brookes Brown (Clemson/Wharton): The Romantic Lie in the Brain: Collective Agency, Moral Responsibility, and the State
Commentator: Frank Hindriks (Groningen)
- 15.30-16.00 Coffee break
- 16.00-17.00 Andreas Schmidt (Groningen): Long-term liberalism and the freedom of future people
Commentator: Brian Berkey (Wharton)

17.30 Drinks (place t.b.c.)
19.00 Dinner (place t.b.c.)

Tuesday, Sept. 13, 9-13

9.30-10.30 Samuel Mortimer (Wharton): What's special about collective action?
Commentator: Lisa Herzog (Groningen)
10.30-10.45 Coffee break
10.45-11.45 Laetitia Mulder (Groningen): When do incentives become "perverse"?
Commentator: Vikram Bhargava (Wharton)
11.45-12.45 Rosemarie Monge (Wharton): The Wrong in Robinhood Markets, Inc.:
Illustrating and Expanding upon the Everyday Business Ethics
Approach
Commentator: Laetitia Mulder (Groningen)
13.00 Lunch (place t.b.c.)

Abstracts (in alphabetical order of speakers)

Brookes Brown (Clemson/Wharton): The Romantic Lie in the Brain: Collective Agency, Moral Responsibility, and the State

Many political philosophers have recently argued that states should be held accountable for their actions as moral agents in themselves. To date, debate has focused on the possibility of moral agency: can collections of people constitute agents in their own right? My interest here is in what has remain hidden in these discussions. Implicit in such claims is not only a theory of collective agency but also a theory of states—of what they are and what they are like—such that they qualify as agents capable of being held accountable, if groups can qualify as such. Unlike questions about group agency as such, this theory has received relatively little attention. My aim in this essay is to expose this oversight as a mistake. I show that our failure to carefully conceptualize the state has blinded us to the fact that political institutions predictably lack the features that theories of collective agency hold they need to be properly held to account: they are too disaggregated, with too little awareness of their circumstances, and too little capacity to control their actions. Failing to adequately theorize the state has thus led advocates of collective agency to advance conclusions that their own views cannot support. Since such inattention to the state is quite widespread, this revelation has serious implications for political philosophy more broadly.

Commentator: Frank Hindriks (Groningen)

Sanjana Govindarajan and Boudewijn de Bruin (Groningen): Artificial Intelligence and Refugee Status Determination: The Role of Businesses in Mitigating Epistemic Injustice in Asylum Systems

Credibility assessments form a central component of asylum determinations worldwide, wherein the decision of a host state to grant or withhold humanitarian protection often hinges on the amount of trust placed by a refugee decision-maker on an asylum-seeker's testimony. In recent years, a growing number of countries have begun to rely on artificial intelligence (AI) technologies to carry out aspects of asylum decision-making. Specifically, they use these technologies as decision-guides to help determine how much credibility to apportion to individual claimants. Due to factors like algorithmic bias and opacity, AI technologies often reinforce and perpetuate pre-existing epistemic vulnerabilities of asylum-seekers. This renders individuals at risk

of being unfairly disbelieved and as a consequence, being denied protection. Literature in the intersection of credibility and forced migration focuses predominantly on the dysfunctional interpersonal and structural dynamics characterizing the relationship between asylum-seeker and decision-maker. In contrast, there is little discussion to be found on the responsibilities of businesses, which essentially mediate these epistemic-legal encounters through the use of their technologies for decision-making. In this paper, we argue that in light of the growing role of AI in migration management, it is essential to reconceptualize the epistemic relationship between the host state and asylum-seeker. Companies developing AI technologies that aid in credibility-related decision-making should be viewed as a third, invisible party in these interactions, bearing their own set of responsibilities vis-à-vis epistemic justice. Our argument, while citing the risks of using AI for credibility-based decision-making, advances a positive account in support of the value-based design of epistemic injustice-sensitive AI technologies. It has meaningful implications both for businesses at the helm of AI innovation as well as philosophical literature on epistemic injustice.

Commentator: Santiago Mejia (Fordham)

Frank Hindriks: How Do Social Structures Empower and Disempower?

Just as a river can prevent someone from traveling onwards, a flow of cars and a stream of rejection letters can constrain her from crossing the street and getting a job. These constraints can be resolved, respectively, by a bridge, a crosswalk and positive discrimination. These examples illustrate that social structures constrain and enable. This platitude is often mentioned, sometimes described, but rarely explained. To do exactly this, I present the Triple AR account of social structures to explain their effects on people's powers to act and achieve things. This account features agents, attitudes and actions as well as rules, roles and resources. It allows me to explain how constraining/disempowering and enabling/empowering are sometimes two sides of the same coin, that constraints on actions can enable people to achieve outcomes, and that signaling rules can create and thereby enable new actions. The Triple AR account also reveals that, when social structures are unjust, raising awareness does little to solve the problem. Activation, assurance and a salient and superior alternative are needed as well.

Commentator: Julian Jonker (Wharton)

Robert Hughes (Harvard Safra Center): The Ethical Duty to Refrain from Wage Exploitation

Much of the philosophical debate about the ethics of wage exploitation presupposes that low wage labor benefits workers compared with a non-transaction baseline. The debate has been about whether some employers who pay low wages ought to benefit their workers more than they do. This paper argues that a non-transaction baseline cannot be used to distinguish beneficial transactions from harmful ones. The non-transaction baseline test cannot explain how harm occurs in cases of overdetermination and preemption. An account of harm that can handle such cases will entail that an employer's choice to pay a low wage can harm the worker even if other prospective employers would have paid even less.

A labor relationship can generate a welfare surplus over which the employer has agency. The employer can choose whether to claim the largest possible share of the surplus (i.e., to pay the lowest wage the market allows) or to divide the surplus with the worker in another way. If the welfare surplus from hiring is great enough to provide a living wage to the worker while still benefiting the employer, the employer's

decision not to pay a living wage is an act that by its very nature prevents the worker from earning a decent living. The resulting deprivation for the employee is attributable to the employer's choice, and it is attributable as an active and intentional harm, not as a mere omission of a benefit. Employers are thus morally required to pay workers a living wage if they can pay a living wage while still benefiting from hiring.

Commentator: Boudewijn de Bruin (Groningen)

Rosemarie Monge (Wharton): The Wrong in Robinhood Markets, Inc.: Illustrating and Expanding upon the Everyday Business Ethics Approach

This paper extends Nien-hê Hsieh and Rosemarie Monge's everyday business ethics approach to a new area of analysis: the scandal that plagued Robinhood Financial, Inc. in the early months of 2021, owing to their decision to stop trading of various stocks that were experiencing heightened levels of volatility. It is argued that the wrong cannot be understood apart from prior choices surrounding Robinhood's risk management and customer practices. The goal of this analysis is threefold. The first is to expand upon Hsieh and Monge's project of 'everyday business ethics.' The analysis does this by illustrating how the everyday business ethics approach might incorporate consequence-based thinking, a necessity if the approach is to live up to its own claims of being suitably pluralist. The second goal of the analysis is to offer an exposition of the proportionality condition of double effect as a way of incorporating consequence-based thinking into the everyday business ethics approach. The third goal is to illustrate how the everyday business ethics approach necessitates conversation about the social values that are brought into tension by participation in the market, thereby demonstrating how everyday business ethics is properly understood as complementing other levels of analysis, such as social ethics or debates in corporate governance.

Commentator: Laetitia Mulder (Groningen)

Samuel Mortimer (Wharton): What's special about collective action?

Some philosophers have suggested that there is something philosophically interesting about the difference between individual and collective actions. Popular explanations of the difference appeal to a range of philosophical constructs—from joint or shared intentions, we-intentions, and participatory intentions, to collective beliefs and desires, mutual obligations, etc. I believe this is a mistake. In this paper, I defend a deflationary account of collective action, which holds that the difference between individual and collective actions can be explained purely in terms of the behavior of the participants. This account faces the challenge of explaining how collective action can be intentional if it is to be reduced to the intentional behavior of individuals: intentionality is usually held to distinguish actions from accidents (and thus, one would expect, collective actions from collective accidents), but according to some philosophers, individuals cannot intend to do collective actions. To respond to this concern, I show that individual agency can extend beyond the individual to encompass the actions of others.

Commentator: Lisa Herzog (Groningen)

Laetitia Mulder (Groningen): When do incentives become “perverse”?

What is the exact effect of incentives/gains for unethical behavior? A large body of literature supports the logical idea that, when there is a lot to gain from unethical behavior, people are more likely to engage in the unethical behavior. For example, it has been found that when people earn more money from lying, they are more likely to lie (Gneezy, 2005; Gneezy, Rockenbach, & Serra-Garcia, 2013), and the prospect of

stock option rewards for CEO's contributed to financial misrepresentation of company results (Harris & Bromiley, 2007). On the other hand, there is literature (o.a. Mazar et al; Shalvi etc) who make the point that people are also motivated to uphold their moral self-image and are therefore *unlikely* to go "all the way" in their unethical acts. Rather, they only cheat to a moderate extent. In this sense, they are more likely to steal €1 rather than €5. So from that literature one would conclude that gains for unethical behavior would not increase the likelihood of engaging in unethical behavior. So the question is: when do high incentives evoke ethical behavior and when do they not? The answer may lie in whether the *harm* of the unethical behavior is (clearly) dependent of the *gain* of the unethical behavior. For some UB's there is a clear link between the two. For example, the higher amount you steal or embezzle, the greater the harm for the victim. For other UB's the link is less clear. For example, engaging in science fraud because you gain tenure with a new publication or because it is "merely" a new publication of an already long list, does not relate to the societal harm done by the fraud. I propose that personal gain increases engagement in UB *only* for those UB's where the personal gain of UB is independent of to the harm of the UB, and not for those UB's where the personal gain is clearly lined to the harm of the UB. I will experimentally test this in a lab setting in which participant play a cheating game in which personal gain and harm of cheating are manipulated independently of each other.

Commentator: Vikram Bhargava (Wharton)

Andreas Schmidt (Groningen): Long-term liberalism and the freedom of future people

What happens to liberal political philosophy, if we consider not only the freedom of present but also future people? Curiously, liberal theorists to date have said little on the freedom of future people. In this article, I defend what I call long-term liberalism: freedom should be a central goal in institutional design and policy, and we should often be particularly concerned with the effect on long-term future distributions of freedom. I argue, first, that liberal arguments to value freedom give us reason to be concerned with the freedom of future people, including those in the far future. Second, engaging with recent discussions on longtermism and our duties towards future people, I argue that reasons to be 'longtermist' in turn imply we should care about the freedom of existing and future people. Third, I outline some implications. For example, long-term liberalism implies liberals should place far greater emphasis on preventing freedom-reducing catastrophic risks. Moreover, long-term liberalism in principle countenances heavy freedom trade-offs between generations, for example by constraining our freedom now for the sake of future people's freedom. However, I also argue that catastrophic risk prevention often does not carry such a trade-off: for both empirical and conceptual reasons, reducing catastrophic risks can often increase rather than decrease the freedom of both future and existing people.

Commentator: Brian Berkey (Wharton)