Algorithmic and Non-Algorithmic Fairness:

Should We Revise our View of the Latter on our Account of Our View of the Former?

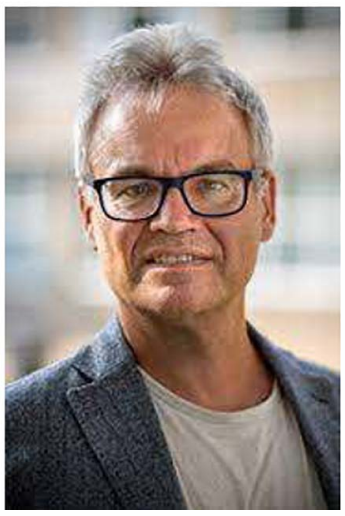

Kasper Lippert-Rasmussen, Professor of Political Science (Aarhus University) and Director of The Centre for the Experimental-Philosophical Study of Discrimination **(CEPDISC)**

Tuesday, 22 November 2022
10:00 – 12:00
Building 1453, room 415
Nobelparken, AU.

ARTS
AARHUS UNIVERSITY

The use of COMPAS for purposes of determining the risk of recidivism has generated a significant amount of work on the issue of algorithmic fairness. One fairness condition which has been defended in this connection is calibration within groups – in the context of COMPAS, roughly, this condition requires that for any given risk score, the expected percentage of individuals who are assigned that risk score who actually recidivate is the same for each and every relevant group, e.g., it is not the case that a greater proportion of high-risk white offenders will recidivate than the proportion of high-risk black offenders.

However, this condition on the absence of unfair discrimination is one that we typically do not impose in non-algorithmic settings as indicated by audit studies. For instance, a hiring committee that, in an unbiased way, assigns qualification scores to applicants such that it satisfies the condition of equal ratios of false positive rate to false negative rate for different relevant groups is generally not seen as engaging in unfair discrimination even if, due to base-rate difference across two groups, it does not satisfy the calibration requirement with respect to these two groups. Hence, if we are attracted to the calibration condition in the context of algorithmic fairness, we should either: 1) revise our view about algorithmic fairness; 2) explain why algorithmic and non-algorithmic contexts are different; or 3) revise our view about non-algorithmic fairness, e.g., in typical job hiring situations. The latter response is quite revisionary and deserves attention.