

## Abstract

Mechanism design deals with distributed algorithms that are executed with self-interested agents. The designer, whose objective is to optimize some function of the agents' private *types*, needs to construct a computation that takes into account agent incentives which are not necessarily in alignment with the objective of the mechanism. Traditionally, agents are modeled to only care about the utility they derive from the mechanism outcome, and mechanisms designed for such agents often fully or partially disclose the agents' declarations. When agents are *privacy aware*, i.e., their loss of privacy can adversely affect their utility, such mechanisms may become inadequate - ignoring privacy-awareness in the design of a mechanism may render it not incentive compatible, and hence inefficient. Interestingly, and somewhat counter-intuitively, Xiao [eprint 2011] has recently showed that this can happen even with mechanisms that preserve a strong notion of privacy.

Towards constructing mechanisms for privacy-aware agents, we put forward and justify a model of privacy-aware mechanism design. We then show that privacy-aware mechanisms are feasible. The following is a summary of our contributions:

- **Modeling privacy-aware agents:** We propose a new model of privacy-aware agents where agents need only have a conservative upper bound on how loss of privacy adversely affects their utility. This is in deviation from prior modeling which required a complete characterization.
- **Privacy of the privacy loss valuations:** Agent privacy valuations are often sensitive on their own. Our model of privacy-aware mechanisms takes into account the loss of utility due to information leaked about these valuations.
- **Guarantees for agents with high privacy valuations:** As it is impossible to guarantee incentive compatibility for agents that have arbitrarily high privacy valuations, we require a privacy-aware mechanism to set a threshold such that the mechanism is incentive compatible w.r.t. agents whose privacy valuations are below the threshold, and differential privacy is guaranteed for all other agents.
- **Constructing privacy-aware mechanisms:** We first construct a privacy-aware mechanism for a simple polling problem, and then give a more general result, based on recent generic construction of approximately additive mechanisms by Nissim, Smorodinsky, and Tennenholtz [ITCS 2012]. We show that under a mild assumption on the distribution of privacy valuations (namely, that valuations are bounded for all but a vanishing fraction of the population) these constructions are incentive compatible w.r.t. almost all agents, and hence give an approximation of the optimum. Finally, we show how to apply our generic construction to get a mechanism for privacy-aware selling of digital goods.