

Persuasive Privacy

ICML 2026

Joshua J Bon[†] James Bailie^{*} Judith Rousseau[◇] Christian P
Robert^{◇○}

[†]Adelaide University ^{*}Chalmers University [◇]Université Paris–Dauphine PSL [○]University
of Warwick

Persuasive Privacy

A **semantics-first** understanding of data privacy where

- Privacy definitions are systematically constructed from an **agent-based game**
- Game assumptions have real-world interpretations
- Probabilistic and pure **DP special cases**
- Rényi and f -divergence DP also by changing assumptions
- Privacy developed for **deterministic mechanisms**

A Game of Persuasion

Sender chooses a mechanism that **persuades** Receiver to make decisions limiting privacy loss.

Sender

Designs mechanism

Worst-case analysis in

- *data value*
- *Receiver's data-prior*

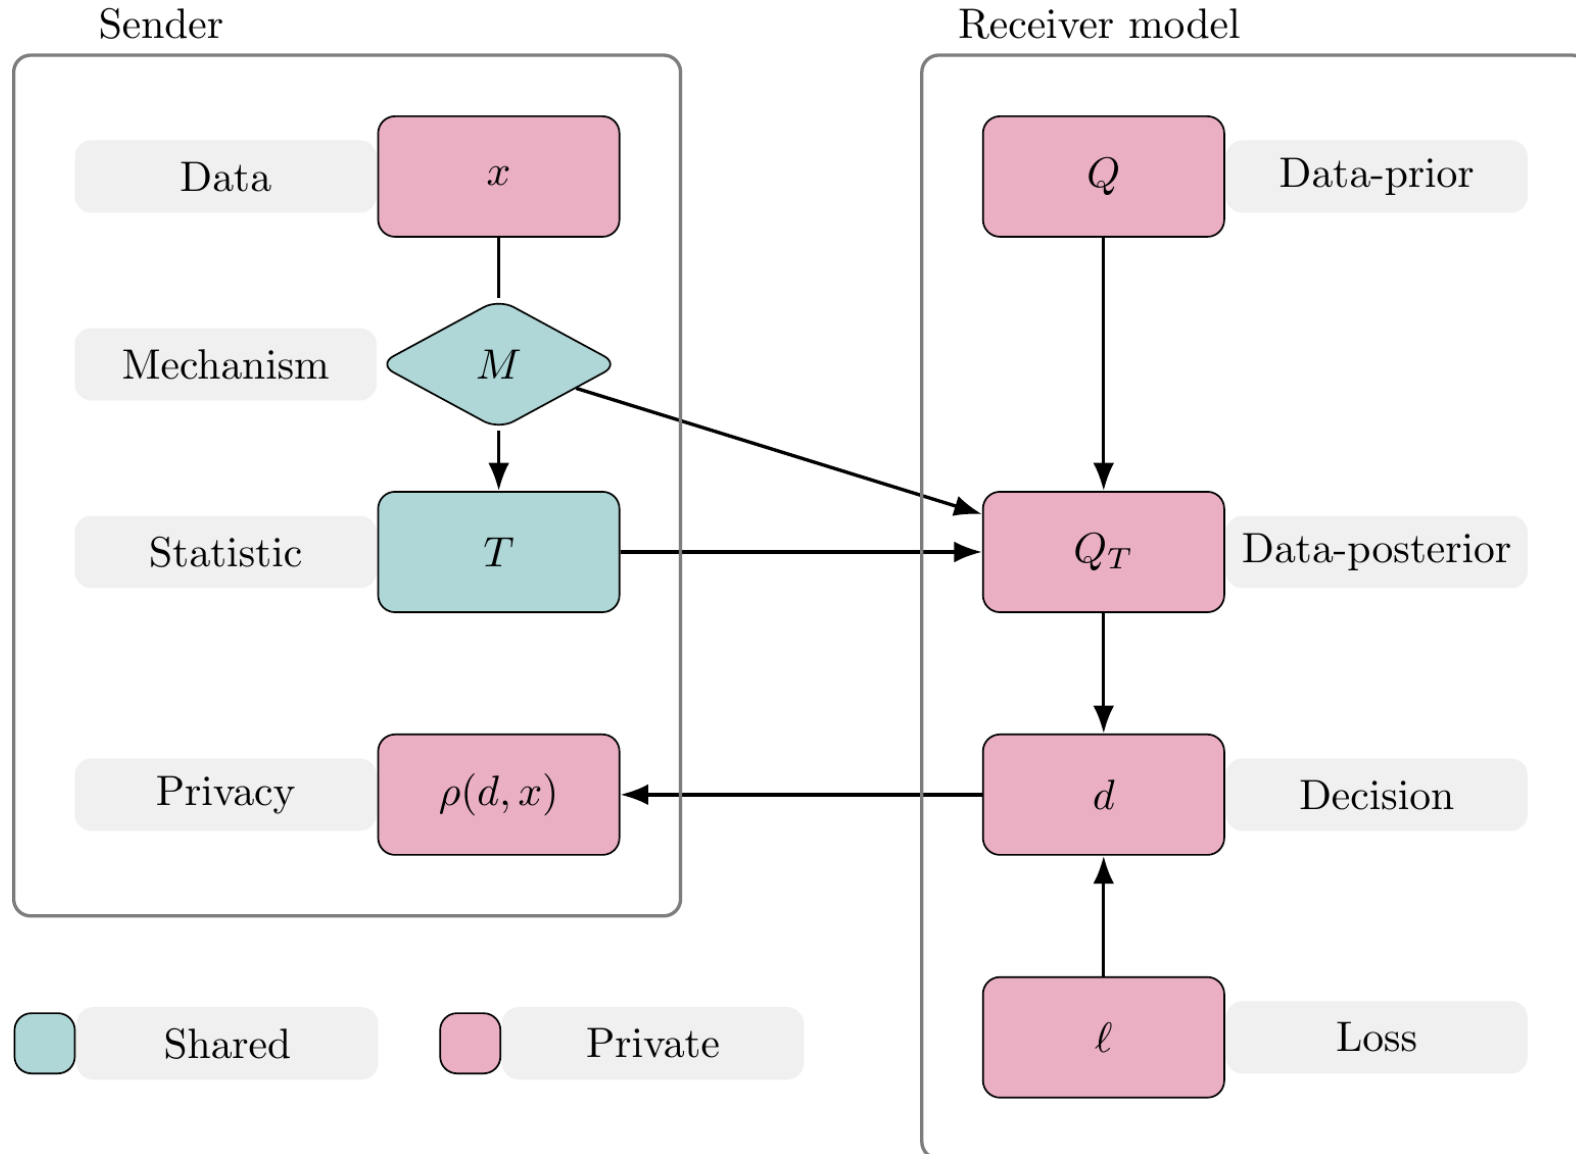
Receiver

Observes mechanism output

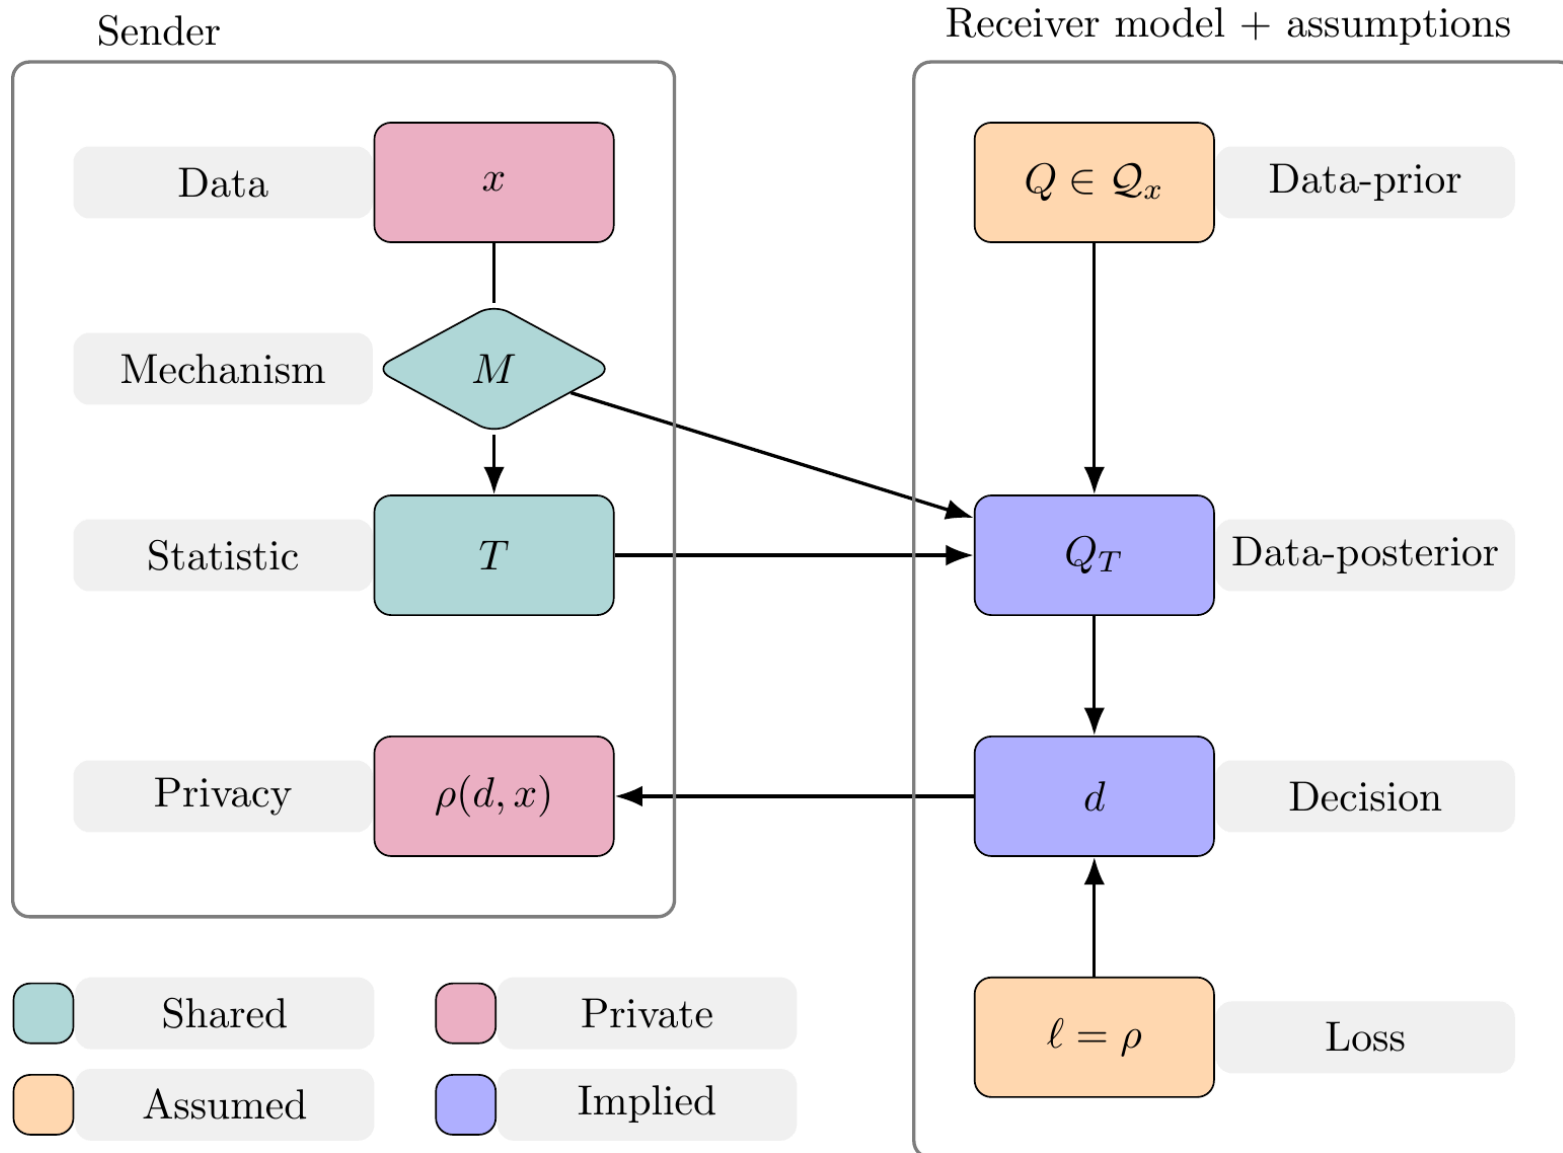
Minimizes expected loss w.r.t. data-posterior,

Bayesian agent¹

Data privacy game



Game under persuasive privacy



Persuasive Privacy Paper

- New interpretations of post-processing property
- Properties: **composition** and **Receiver post-processing**
- Deterministic mechanisms: empirical mean, cell suppression
- Recovery of existing DP privacy definitions
- arXiv: <https://arxiv.org/abs/2601.22945>

References

Kamenica, Emir, and Matthew Gentzkow. 2011. “Bayesian Persuasion.” *American Economic Review* 101 (6): 2590–2615.