Differential Privacy for Strategic Information Sharing and Learning: An Annotated Reading List

M. AMIN RAHIMIAN
University of Pittsburgh
and
YUXIN LIU
University of Pittsburgh

Differential privacy has evolved from a technical framework for protecting individual-level data into a useful language for studying privacy in economic and strategic environments. Beyond limiting what can be inferred about any single data contributor, differential privacy can shape incentives, participation, information disclosure, and data-driven decisions. This annotated reading list highlights a small set of papers that connect the foundations of differential privacy to recent work on data acquisition, platform design, and operational decision-making under privacy constraints.

General Terms: Algorithms, Economics, Theory
Additional Key Words and Phrases: Differential Privacy, Mechanism Design, Data Acquisition, Platform Decisions, Operational Decisions

1. INTRODUCTION

Differential privacy (DP) was originally developed as a rigorous framework for protecting individual-level data while enabling statistical analysis. The standard formulation of DP requires that a randomized mechanism produce similar output distributions on neighboring datasets [Dwork et al. 2006]. We call two datasets neighboring if they differ in the data of a single individual. Formally, a randomized mechanism M is ε-differentially private if, for any neighboring datasets D and D' and for any possible output set S,

$$\Pr[M(D) \in S] \leq e^{\varepsilon} \Pr[M(D') \in S].$$

This definition limits how much the participation of any single individual can change the distribution of observable outcomes.

For the economics and computation community, this stability property is useful not only as a technical privacy guarantee, but also as a way to reason about incentives and strategic behavior. When the influence of any single report is limited, an agent's ability to manipulate outcomes is also limited. This observation connects differential privacy to mechanism design, large games, and other settings in which agents strategically decide what information to reveal.

Building on this perspective, recent work has used differential privacy to study a broader range of economic and operational questions. Privacy constraints can affect who is willing to contribute data, how much data must be purchased, how platforms use information for personalization, and how firms make data-driven pric-

Authors' addresses: rahimian@pitt.edu, yul435@pitt.edu

ing, recommendations, and inventory decisions. Thus, DP is not only a constraint on statistical release; it can also be a design feature that shapes participation, information use, and operational performance.

This annotated reading list highlights a small set of papers that connect the foundations of differential privacy to recent work on data acquisition, platform design, and operational decision-making under privacy constraints. The selected papers are not intended to be comprehensive. Rather, they offer a concise path from foundational ideas to applications in economic, strategic, and operational settings.

2. ANNOTATED READING LIST

I. Foundations and Game-Theoretic Perspectives

1. *Calibrating Noise to Sensitivity in Private Data Analysis* [Dwork et al. 2006]

This paper introduces differential privacy and the Laplace mechanism, establishing the sensitivity-based approach to rigorous privacy guarantees. It provides the formal starting point for understanding privacy as a stability constraint: the participation of any single individual should have only a limited effect on the distribution of released outcomes.

2. *Mechanism Design via Differential Privacy* [McSherry and Talwar 2007]

This paper made one of the earliest connections between differential privacy and incentives. By introducing the exponential mechanism, it showed that DP can be used not only to protect data but also to stabilize agents' incentives by limiting the influence of any single report on the outcome. Because the potential gain from manipulation is bounded to be near zero, truthful reporting becomes an approximate dominant strategy.

3. *Selling Privacy at Auction* [Ghosh and Roth 2015]

This paper studies agents who have explicit concerns about the privacy loss associated with the use of their personal data. It formulates private-data collection as a mechanism design problem, where agents may need to be compensated for privacy loss before their data can be used to estimate a population statistic. This perspective connects privacy concerns with participation, payments, and the cost of obtaining useful data.

II. Privacy, Data Acquisition, and Participation

4. *Optimal Data Acquisition with Privacy-Aware Agents* [Cummings et al. 2023]

This paper studies data acquisition when agents have heterogeneous privacy costs and derive value from the quality of the learned model. It highlights a central trade-off: stronger privacy can reduce statistical accuracy by adding noise, but it can also increase participation by making agents more willing to share data.

5. *Optimal and Differentially Private Data Acquisition: Central and Local Mechanisms* [Fallah et al. 2024]

This paper studies data acquisition from privacy-sensitive agents who have heterogeneous and privately known privacy costs. The platform must choose privacy

losses, payments, and an estimator to induce truthful participation while estimating an underlying parameter. The paper compares central and local privacy architectures. In the central model, agents share data with a trusted curator, who forms a weighted average and adds aggregate Laplace noise. In the local model, each agent adds Laplace noise before sharing data, so the platform averages already privatized reports. This distinction changes both estimation and mechanism design: the central model supports a more direct score-based mechanism, whereas the local model leads to a harder optimization problem over independently privatized reports.

6. *The Privacy Paradox and Optimal Bias–Variance Trade-offs in Data Acquisition* [Liao et al. 2024]

This paper studies data acquisition when privacy concerns may be correlated with the data being collected. It shows how participation decisions can create bias and develops mechanisms that balance the bias introduced by selective participation against the variance introduced by privacy noise.

III. Differential Privacy in Operational and Platform Decisions

7. *Privacy-Preserving Personalized Revenue Management* [Lei et al. 2024]

This paper studies differential privacy in personalized revenue management. It shows how firms can use historical customer data to make personalized revenue decisions while limiting the disclosure of individual information. A central insight is that, when sufficient historical data are available, privacy protection can sometimes be achieved with only a small additional loss relative to the statistical cost of learning demand.

8. *An Algorithmic Approach to Managing Supply Chain Data Security: The Differentially Private Newsvendor* [Chen and Chua 2026]

This paper brings differential privacy into supply-chain and inventory decisions. In a data-driven newsvendor problem, the ordering decision itself may reveal sensitive demand information. The paper studies how to design private decision rules that protect data while maintaining near-optimal operational performance.

9. *Privacy-Preserving Dynamic Personalized Pricing with Demand Learning* [Chen et al. 2022]

This paper studies dynamic personalized pricing when a firm learns demand from consumer data over time. It shows how privacy constraints interact with demand learning and revenue maximization in terms of regret and develops pricing policies that protect consumer information while preserving useful personalization.

10. *Privacy-Preserving Personalized Recommender Systems* [Fu et al. 2026]

This paper studies personalized recommendation under differential privacy. It analyzes how a platform can use customer data to improve recommendations while limiting the disclosure of individual information. It also illustrates the broader role of differential privacy in platform decisions that depend on personalization.

REFERENCES

Chen, D. and Chua, G. A. 2026. An algorithmic approach to managing supply chain data security: The differentially private newsvendor. *Operations Research 74,* 2, 958–983.

Chen, X., Simchi-Levi, D., and Wang, Y. 2022. Privacy-preserving dynamic personalized pricing with demand learning. *Management Science 68,* 7, 4878–4898.

Cummings, R., Elzayn, H., Pountourakis, E., Gkatzelis, V., and Ziani, J. 2023. Optimal data acquisition with privacy-aware agents. In *2023 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*. IEEE, 210–224.

Dwork, C., McSherry, F., Nissim, K., and Smith, A. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*. Springer, 265–284.

Fallah, A., Makhdoumi, A., Malekian, A., and Ozdaglar, A. 2024. Optimal and differentially private data acquisition: Central and local mechanisms. *Operations Research 72,* 3, 1105–1123.

Fu, X., Chen, N., Gao, P., and Li, Y. 2026. Privacy-preserving personalized recommender systems. *Manufacturing & Service Operations Management 28,* 1, 271–289.

Ghosh, A. and Roth, A. 2015. Selling privacy at auction. *Games and Economic Behavior 91*, 334–346.

Lei, Y., Miao, S., and Momot, R. 2024. Privacy-preserving personalized revenue management. *Management Science 70,* 7, 4875–4892.

Liao, G., Su, Y., Ziani, J., Wierman, A., and Huang, J. 2024. The privacy paradox and optimal bias–variance trade-offs in data acquisition. *Mathematics of Operations Research 49,* 4, 2749–2767.

McSherry, F. and Talwar, K. 2007. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*. IEEE, 94–103.