Summary of Changes

Note: This document uses the quotation environment to reference content from `OGPM_diff.pdf`. Within the quotation environment, newly added or modified content is highlighted in blue. Cross-references in this document have been updated, so numbering may differ from `OGPM_diff.pdf`. References to specific locations in `OGPM_diff.pdf` are indicated in gray text, e.g. Page 6.

This document summarizes the changes made in the revised version. Section 1 details changes in the main text, while Section 2 covers changes in the appendix. The key updates are as follows:

1. **Softening the optimality claim for $m = 3$:** The claim that $m = 3$ is optimal (original Theorem 3.7) has been softened to a hypothesis. Monte Carlo results now are presented as support to this hypothesis. Correspondingly, theorems about optimal instantiations incorporate this hypothesis. To further substantiate the hypothesis, two directions for an analytical proof and associated challenges are outlined in the appendix. (Section 1.1 and Section 2.1)

2. **Clarifying OGPM's optimality conditions:** The conditions under which OGPM is optimal are explicitly stated in the main text. Additionally, the appendix discusses the generality of these conditions and related works on optimality. (Section 1.2 and Section 2.2)

3. **Expanded experimental comparisons:** Experiments and discussions have been added to compare OGPM with other mechanisms applicable to bounded numerical domains. These include: (i) The PM and SW mechanisms with post-processing by truncation; (ii) Variants of the Laplace mechanism, including the staircase mechanism, the Laplace mechanism with truncation, and the bounded Laplace mechanism; (iii) Mechanisms from the VMF paper. (Section 1.3 and Section 2.3)

4. **Clarifying no post-processing assumption:** In the section "Distribution and Mean Estimation," it is clarified that the data collector uses the reported values as observed, without applying post-processing. (Section 1.4)

5. **Additional error plots:** Whole-domain error plots for $\varepsilon = 0.4$ and $\varepsilon = 0.8$ in the classical domain have been added. (Section 1.5 and Section 2.4)

6. **Removal of inaccurate claim regarding truncation:** The footnote and appendix claiming that truncation of LDP mechanisms may harm the privacy level have been removed. (Section 1.6)

The remainder of this document provides details of the changes made to the main text and appendix, organized into two sections.

1 Changes in Main Text

This section details the six changes made to the main text of the paper.

1.1 Softening the Optimality Claim for $m = 3$

(Page 6) The claim that $m = 3$ is optimal has been softened to a hypothesis. Monte Carlo results are now presented as support for this hypothesis. The revised content is as follows:

> **Hypothesis 1.1.** *(original Theorem 3.7) For any domain $\mathcal{D} \to \mathcal{D}$, absolute error and square error metrics, the optimal piecewise-based mechanism falls into 3-GPM.*

Support. We validated this hypothesis for $\mathcal{D} = [0, 1)$ by performing Monte Carlo sampling on 10^4 random (ε, x) pairs. Theorem 3.6 at Page 5 then extends this optimality to any $\mathcal{D}$. Given the continuity of the objective functions w.r.t. ε and x, we posit that $m = 3$ is indeed the exact optimal. To further support this hypothesis, Appendix 2.1 outlines two directions for analytical proof and highlights the associated challenges.

Correspondingly, theorems about optimal instantiations will incorporate this hypothesis. The result for the classical domain is now stated as follows:

Theorem 1.1. *(Theorem 3.8 at Page 6) If hypothesis 1.1 holds, then GPM* $\mathcal{M} : [0, 1) \to [0, 1)$ *with the following closed-form instantiation*

$$pdf[\mathcal{M}(x) = y] = \begin{cases} p_\varepsilon & \text{if } y \in [l_{x,\varepsilon}, r_{x,\varepsilon}), \\ p_\varepsilon / \exp(\varepsilon) & y \in [0, 1) \setminus [l_{x,\varepsilon}, r_{x,\varepsilon}), \end{cases}$$

where $p_\varepsilon = \exp(\varepsilon/2)$,

$$[l_{x,\varepsilon}, r_{x,\varepsilon}) = \begin{cases} [0, 2C) & \text{if } x \in [0, C), \\ x + [-C, C) & \text{if } x \in [C, 1 - C), \\ [1 - 2C, 1) & \text{otherwise}, \end{cases}$$

with $C = (\exp(\varepsilon/2) - 1)/(2\exp(\varepsilon) - 2)$, *is optimal for* $[0, 1) \to [0, 1)$ *under the absolute error and square error metric.*

OGPM for the circular domain is updated accordingly (Theorem 4.2 at Page 7). Also, at Page 8, we updated the context and proof of the unbiased OGPM (Theorem 5.1). The current version of this part is as follows:

Unbiased mean estimation. Note that an unbiased mean estimator can be achieved by enlarging the output domain $\mathcal{D} \to \tilde{\mathcal{D}}_\varepsilon$. Mathematically, this involves incorporating the unbiasedness constraint $\mathrm{E}[\mathcal{M}(x)] = x$ into optimization problems for solving $\mathcal{M}$. Following the same optimization process as in the classical domain, we hypothesize that the 3-GPM remains optimal for domain $\tilde{\mathcal{D}}_\varepsilon$.

Hypothesis 1.2. *For any domain* $\mathcal{D} \to \tilde{\mathcal{D}}_\varepsilon$*, where* $\tilde{\mathcal{D}}_\varepsilon$ *is a variable w.r.t* ε*, and under absolute error and square error metrics, the optimal piecewise-based mechanism falls into 3-GPM.*

Under the 3-GPM, an unbiased mechanism $\mathcal{M}$ with a variable output domain $\tilde{\mathcal{D}}_\varepsilon$ can be analytically derived by incorporating the unbiasedness constraint. As a complement to Theorem 3.7, we provide Theorem 1.2 for mean estimation in the classical domain.

Theorem 1.2. *(Theorem 5.1) Denote* $\tilde{\mathcal{D}}_\varepsilon = [-C, C+1)$ *with* $C = (\exp(\varepsilon/2)+1)/(\exp(\varepsilon/2)-1)$. *If Hypothesis 1.2 holds, then among the unbiased GPM* $\mathcal{M} : [0, 1) \to \tilde{\mathcal{D}}_\varepsilon$ *(i.e.* $\mathrm{E}[\mathcal{M}(x)] = x$*), closed form*

$$pdf[\mathcal{M}(x) = y] = \begin{cases} p_\varepsilon & \text{if } y \in [l_{x,\varepsilon}, r_{x,\varepsilon}), \\ p_\varepsilon / \exp(\varepsilon) & y \in \tilde{\mathcal{D}}_\varepsilon \setminus [l_{x,\varepsilon}, r_{x,\varepsilon}), \end{cases}$$

where $p = \exp(\varepsilon/2)/(2C + 1)$,

$$l_{x,\varepsilon} = \frac{C+1}{2} \cdot x - \frac{(3C+1)(C-1)}{4C},$$
$$r_{x,\varepsilon} = \frac{C+1}{2} \cdot x + \frac{(C+1)(C-1)}{4C}.$$

is optimal for $[0, 1) \to \tilde{\mathcal{D}}_\varepsilon$ *and the square error metric.*

Proof. ~~The optimality of 3-GPM on $[0, 1) \to [-C, C + 1)$ can be proved as in Theorem 3.1, i.e. $m = 3$ is optimal.~~ Optimal C, p, $l_{x,\varepsilon}$, and $r_{x,\varepsilon}$ are derived by analytical deduction as in Theorem 3.7. Appendix A.6 proves the unbiasedness. □

Hypothesis 1.2 naturally extends Hypothesis 1.1, since $\tilde{\mathcal{D}}_\varepsilon$ becomes explicit once ε is specified. Finding the optimal GPM distribution over $\tilde{\mathcal{D}}_\varepsilon$ uses the same optimization approach as for $\mathcal{D}$. Therefore, this hypothesis is also expected to hold for $\tilde{\mathcal{D}}_\varepsilon$ as well.

1.2 Clarifying OGPM's Optimality Conditions

(Page 4) We added a paragraph to clarify the conditions under which OGPM is optimal. The content of the paragraph is as follows:

> **Conditions for optimality.** When discussing optimality, the following aspects should be specified: (i) the error metric, (ii) the data domain and family of mechanisms, (iii) the strength of the optimality, and (iv) whether post-processing is allowed. In this paper, the optimality of GPM is defined with respect to: (i) the worst-case L_p-similar error metric, (ii) bounded numerical domains $\mathcal{D} \to \tilde{\mathcal{D}}$ and mechanisms based on piecewise distributions, (iii) minimization of error value (not asymptotic or order-of-magnitude optimality), and (iv) without post-processing. These conditions are widely applicable in practice and literature. However, varying any of them may lead to different optimality results. Appendix 2.2 provides a detailed discussion of these conditions and related optimalities.

1.3 Expanded Experimental Comparisons

(Page 9-10) In the evaluation setup, we expanded the description to include additional mechanisms applicable to bounded numerical domains. The updated content is as follows:

> This section evaluates the theoretical and experimental data utility of our methods by comparing them with existing instantiations and their variants:
>
> - OGPM: closed-form optimal GPM;
> - OGPM-U: unbiased closed-form optimal GPM for mean estimation in the classical domain;
> - PM [8], SW [7], and their post-processed versions: PM is the first TPM designed for mean estimation, while SW is designed for distribution estimation. Both mechanisms output enlarged domains but can be post-processed by truncating outputs to the input domain. These post-processed versions are referred to as T-PM and T-SW for convenience.
> - PM-C and SW-C: the compressed versions of PM and SW for $\mathcal{D} \to \mathcal{D}$. For the best potential of PM and SW, we adapt them to $\mathcal{D} \to \mathcal{D}$ as PM-C and SW-C by linearly compressing their output domain $\tilde{\mathcal{D}}_\varepsilon$ to $\mathcal{D}$, i.e. transformation invariants, which maintains the privacy level.
>
> We also compare OGPM's expected error with non-piecewise-based mechanisms that can be applied to bounded domains:
>
> - Variants of the Laplace mechanism: including the staircase mechanism [3], Laplace mechanism with post-processing by truncation (T-Laplace), and the bounded Laplace mechanism (B-Laplace) [6], which designs a bounded Laplace-shape distribution.
> - Purkayastha mechanism [9]: a mechanism for directional data on spheres $\mathbb{S}^{n-1}$. When $n = 2$, it is a counterpart of OGPM in the circular domain.

(Page 11-12) We added experimental results and discussions comparing OGPM with the aforementioned new mechanisms. The content of this part is as follows:

> ### 1.3.1 Comparison with PM and SW
>
> Figure 1 presents the comparison of the whole-domain error in the classical domain for the original PM and SW mechanisms, along with their post-processed versions, T-PM and T-SW. For a fair comparison, OGPM is adapted to the domain $\mathcal{D} = [-1, 1)$ to match PM's design, while SW and OGPM remain consistent with $\mathcal{D} = [0, 1)$. The post-processing of PM and SW involves truncating their outputs in the enlarged domain to the input domain, i.e. applying $\mathcal{I} \circ \mathcal{M}(x)$, where $\mathcal{I} : \tilde{\mathcal{D}} \to \mathcal{D}$ is the truncation operator. We use the distance metric $\mathcal{L} = |y - x|^2$ and set $\varepsilon = 2$ for the comparison among these five mechanisms. It can be observed that OGPM consistently achieves the lowest error across all x values, with a more significant advantage compared to the comparison with PM-C and SW-C. This is because the original PM and SW output larger domains, resulting in higher errors. Meanwhile, T-PM reduces the error of PM more effectively than T-SW reduces the error of SW, as the original PM has a more enlarged

output domain than SW, making truncation more impactful. This comparison highlights OGPM's error advantage over the original PM, SW, and their post-processed versions when applied to their respective data domains.

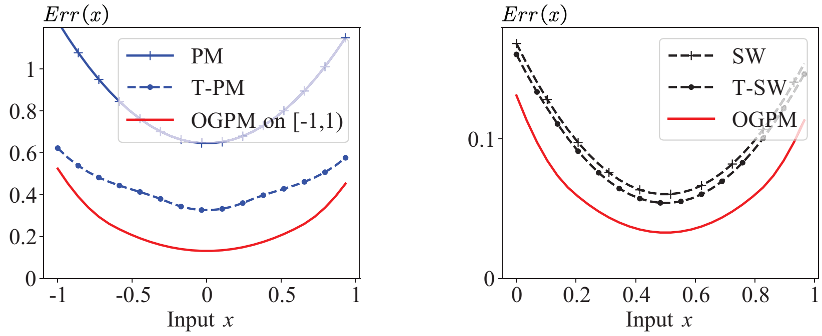

(a) Comparison with PM. (b) Comparison with SW.

Figure 1: Whole-domain error comparison with PM and SW on their data domains (i.e. $\mathcal{D} = [-1, 1)$ and $\mathcal{D} = [0, 1)$, respectively) when $\varepsilon = 2$.

1.3.2 Comparison with the Staircase Mechanism, T-Laplace, and B-Laplace

In addition to piecewise-based mechanisms, the Laplace mechanism and its variants can also be applied to the classical domain to achieve LDP. Among these, the staircase mechanism [3] claims to be optimal under certain assumptions. For the input domain $\mathcal{D} = [0, 1)$ (i.e. sensitivity $\Delta = 1$) and error metric $\mathcal{L} = |y - x|$, its expected error is given by Theorem 3 in [3]: $\exp(\varepsilon/2)/(\exp(\varepsilon) - 1)$. Another approach involves using the Laplace mechanism with truncation [6], referred to here as T-Laplace for convenience. T-Laplace preserves the privacy guarantees of the Laplace mechanism while reducing the expected error, particularly for data points near the endpoints or for small ε values. Additionally, the bounded Laplace mechanism (B-Laplace) [6] introduces a redesigned bounded Laplace-shaped distribution tailored for bounded domains.*

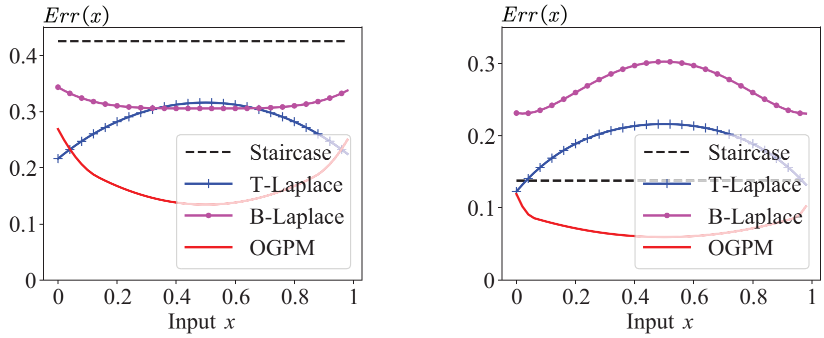

(a) Privacy parameter $\varepsilon = 2$. (b) Privacy parameter $\varepsilon = 4$.

Figure 2: Whole-domain error comparison with the staircase mechanism [3], T-Laplace and B-Laplace mechanisms [6] in the classical domain with error metric $\mathcal{L} = |y - x|$.

Figure 2 compares the whole-domain error in the classical domain $\mathcal{D} = [0, 1)$ for the staircase mechanism, T-Laplace, and B-Laplace. These mechanisms exhibit distinct error patterns across the domain. For the staircase mechanism, the error remains constant, as it is determined by a fixed staircase distribution and is independent of x. For T-Laplace, the error reaches its maximum at the midpoint and its minimum at the endpoints, as truncation favors the endpoints. For instance, when $x = 0$, it is error-free with a probability of $1/2$, due to the symmetry of the Laplace distribution around 0. For B-Laplace, the error trend varies with ε. When $\varepsilon = 2$, the error decreases with x and reaches its minimum at the midpoint, whereas

*Appendix 2.3 provides details on the expected error of B-Laplace.

for $\varepsilon = 4$, the error increases with x and peaks at the midpoint. Despite these differing error patterns, OGPM consistently achieves lower errors than the staircase mechanism and T-Laplace across the whole domain.

Figure 3 compares the worst-case error w.r.t. ε in the classical domain. OGPM consistently achieves the lowest worst-case error across all ε values. For small ε, T-Laplace and B-Laplace exhibit a significant advantage over the staircase mechanism; however, this advantage diminishes as ε increases. At larger ε values, the error of the staircase mechanism approaches that of OGPM.

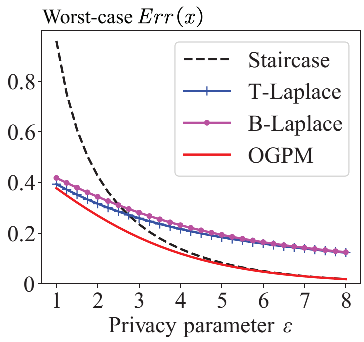

Figure 3: Worst-case error comparison (continued from Figure 2).

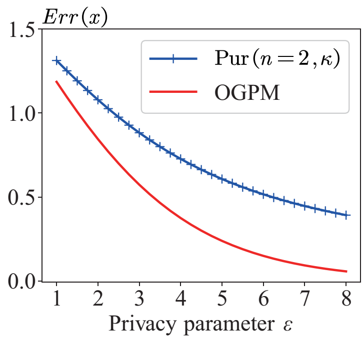

Figure 4: Comparison with the Purkayastha mechanism [9] for sphere $\mathbb{S}^{n-1}$.

1.3.3 Comparison with the Purkayastha Mechanism

The paper "Differential Privacy for Directional Data" [9] introduces two mechanisms for data on spheres $\mathbb{S}^{n-1}$: the VMF mechanism (ensuring indistinguishability of any two points with distance *through* the sphere) and the Purkayastha mechanism (ensuring indistinguishability of any two points with distance *along* the sphere). When $n = 2$, the sphere $\mathbb{S}^1$ corresponds to a circle, making the Purkayastha mechanism a counterpart of OGPM in the circular domain. Therefore, we compare them in the circular domain.[†]

Figure 4 presents the comparison of the expected error in the circular domain between OGPM and the Purkayastha mechanism. The expected error of the Purkayastha mechanism is derived using the closed-form expressions in Theorem 19 and 22 of [9], with $\kappa = \varepsilon/\Delta_{\measuredangle}$. Since the errors of both mechanisms are x-independent in the circular domain, it suffices to compare their worst-case errors. The results demonstrate that OGPM consistently outperforms the Purkayastha mechanism, achieving significantly lower errors.

(Page 6) We also weakened the claim regarding LDP mechanisms for the circular domain; it now states that none of the existing piecewise-based mechanisms address this type of domain. The first paragraph of the section "Optimal GPM for Circular Domain" is now as follows:

> This section presents the optimal GPM for the circular domain, another type of bounded domain. Circular domains are widely used in cyclic data such as time, angle, and compass direction. However, none of the existing piecewise-based mechanisms consider this type of domain, limiting their applicability.

1.4 Clarifying No Post-processing Assumption

(Page 8) In the section "Distribution and Mean Estimation," we clarified that the data collector uses the reported values as observed, without applying any post-processing. The revised second paragraph is as follows:

[†] We omit the comparison with the VMF mechanism also because (i) it has been shown that the Purkayastha mechanism outperforms the VMF mechanism (with the same sensitivity $\Delta_{\measuredangle} = \pi$ for sphere $\mathbb{S}^1$, e.g. Figure 5 and 10 in [9]), and (ii) the expected error of the VMF mechanism lacks a closed-form expression (Theorem 17 in [9]), making it complex to compute.

Assume a set of users with sensitive data $\mathcal{X} = \{x_1, x_2, \ldots, x_n\}$. They apply $\mathcal{M}$ to produce randomized outputs $\mathcal{Y} = \{y_1, y_2, \ldots, y_n\}$. The data collector then estimates the distribution and mean of $\mathcal{X}$ using $\mathcal{Y}$. Specifically, the collector uses the values in $\mathcal{Y}$ as it is, i.e. without knowing and applying any post-processing based on prior knowledge of $\mathcal{X}$.

1.5 Additional Error Plots

(Page 11) We added a sentence to reference the whole-domain error plots at $\varepsilon = 0.4$ and $\varepsilon = 0.8$ for the classical domain, which are provided in the appendix. The content of the paragraph is as follows:

> **Classical Domain.** Figure 8 shows ... More comparisons under smaller ε values are provided in Appendix 2.4.

1.6 Removal of Inaccurate Claim Regarding Truncation

(Page 3) We removed the inaccurate claim regarding truncation in the footnote. Specifically, the footnote is now as follows:

> While post-processing the output by truncating it to $\mathcal{D}$ is possible, this approach may still result in low data utility. Sections 1.3.1 and 1.3.2 provide comparisons with mechanisms that include truncation.

2 Changes in Appendix

This section details changes made to the appendix of the paper.

2.1 (Appendix B.3) Directions for Analytically Proving Optimal $m = 3$

(Page 18) This appendix outlines two potential directions for analytically proving that the optimal m is 3, along with the challenges associated with each approach. The content of the appendix is as follows:

> Mathematically, finding the optimal m-piecewise mechanism is equivalent to identifying the optimal m-piecewise distribution under an L_p-similar error metric. It is seemingly true that the optimal m is 3: if the optimal m-piecewise distribution is not 3 but 4 or more, we can always shift the probability mass from the two side intervals (i.e. other pieces) to the central interval, thereby reducing the error. At the very least, the following fact holds:
>
> **Fact 2.1.** *The optimal m-piecewise distribution has a strict staircase shape, i.e. the probability density of the central interval is greater than that of the two side intervals.*

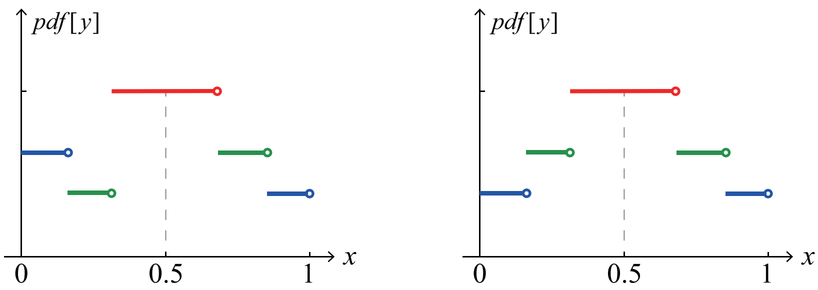

(a) Non-staircase distribution. (b) Staircase distribution.

Figure 5: A non-staircase distribution (left) can always be shifted into a staircase distribution (right) by moving some pieces closer to x, which reduces the error.

> Figure 5 illustrates this fact. Moving pieces while keeping their probabilities unchanged clearly maintains both the ε-LDP constraint and the probability normalization constraint. This observation reduces the problem to proving that a 3-staircase distribution can achieve the same optimal error as a 4-staircase distribution under the ε-LDP and probability normalization constraints.

Direction 1: If we can further move the green piece in Figure 5b "into" the red central piece while keeping the probabilities of the red and blue pieces unchanged, i.e. transform it into a 3-staircase distribution while ensuring a decrease in the error, then we can prove that the optimal m is 3. However, this is challenging, as it breaks the probability normalization constraint, also requiring adjustments to the probabilities of each piece to satisfy the ε-LDP constraint. The challenge arises here: it is hard to ensure that these adjustments will indeed decrease the error.

Direction 2: Another approach is to formulate the problems for 3-staircase and 4-staircase distributions as two constrained optimization problems. The goal would be to prove that the optimal error of the 3-staircase distribution is equivalent to that of the 4-staircase distribution. Ideally, these two multi-variable optimization problems could be solved analytically, resulting in two closed-form error expressions w.r.t. x and ε, thereby completing the proof for any x and ε. This direction aligns with our framework. However, the challenge lies in the complexity of solving such multi-variable optimization problems analytically. This is why we rely on an off-the-shelf optimization solver, which, while effective, only provides numerical solutions for specific x and ε values.

2.2 (Appendix B.2) Related Optimality

(Page 17-18) This section of the appendix has been rewritten to update the assumptions underlying OGPM's optimality and to provide a more detailed discussion of related optimality concepts. The revised content is as follows:

In this paper, the optimality of GPM is defined with respect to: (i) the worst-case L_p-similar error metric, (ii) bounded numerical domains $\mathcal{D} \to \tilde{\mathcal{D}}$ and mechanisms based on piecewise distributions, (iii) minimization of error value (not asymptotic or order-of-magnitude optimality), and (iv) without post-processing. L_p-similar error metrics are natural choices for evaluating data utility [5, 3, 8]. Bounded numerical domains are common in real-world applications. Focusing on error values allows for more precise comparisons between different mechanisms. By excluding post-processing, we can analyze the optimality of the mechanism itself, providing a more fundamental understanding than considering the mechanism combined with specific post-processing.

Other types of optimality have been explored in the literature, particularly for variants of Laplace mechanisms. The staircase mechanism [3] adopts the same utility model without prior knowledge or post-processing as this paper. It claims optimality under specific assumptions, one of which is that a staircase (piecewise) distribution can achieve the optimal error. The mechanism demonstrates better L_1-error performance than the Laplace mechanism on $\tilde{\mathcal{D}} = (-\infty, \infty)$, and its asymptotic optimality has been formally proven. *Universal optimality* is another type of optimality, defined from the perspective of a user's prior knowledge and post-processing ability [4]. In this utility model, the user observes the output of the mechanism and selects another value based on the output and their prior knowledge, i.e. under a Bayesian utility framework. Formally, if the user's prior is denoted as p_i on the data domain $i \in N$ (i.e. a discrete domain) and the user's post-processing is represented as a remap $z_{i,j}$ that reinterprets the output of the mechanism (on the sensitive value i) to j, then the utility model is

$$Err(i) = \sum_{i \in N} p_i \sum_{j \in N} z_{i,j} \cdot \mathcal{L}(i,j).$$

This utility model incorporates the user's prior knowledge and post-processing ability. A mechanism is called universally optimal if, for any prior p_i, there exists an optimal remap $z_{i,j}$. Under this utility model, it was proven that the truncated geometric mechanism (a discretized version of the Laplace mechanism) can achieve universal optimality for count queries[‡] and a legal error metric $\mathcal{L}(i,j)$. Such universal optimality was shown to be unachievable for more complex queries [1]. Under the same utility model, the universal optimality was extended to the truncated Laplace mechanism for a bounded numerical domain $\mathcal{D} = [0,1]$ by approximating the geometric mechanism with the Laplace mechanism and post-processing [2].

These optimality results do not hold in our utility model, i.e. utility model without prior and post-processing. Figure 2 has shown that OGPM generally has a smaller error than

[‡]This is in the centralized DP setting, where the data curator holds the dataset and uses *one* mechanism.

the truncated Laplace mechanism, indicating the sub-optimality of the truncated Laplace mechanism in the absence of prior and post-processing.

2.3 (Appendix B.8) Expected Error of the B-Laplace Mechanism

(Page 19-20) We added the computation of the expected error of the B-Laplace mechanism in the appendix. This computation, which is not included in the original B-Laplace paper [6], provides additional insights into its performance. For further details, please refer to `OGPM_diff.pdf`.

2.4 (Appendix B.7) Comparison under Small ε

(Page 19) We added the whole-domain error plots at $\varepsilon = 0.4$ and $\varepsilon = 0.8$ for the classical domain. The content of this appendix is as follows:

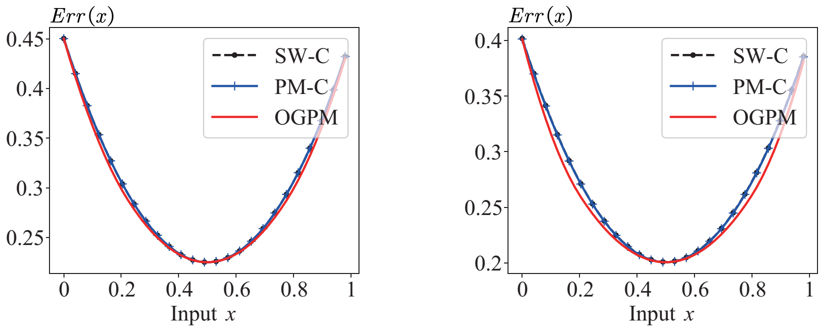

(a) Privacy parameter $\varepsilon = 0.4$.

(b) Privacy parameter $\varepsilon = 0.8$.

Figure 6: Whole-domain error comparison in the classical domain with error metric $\mathcal{L} = |y - x|$.

Figure 6 presents the whole-domain error comparison of OGPM, PM-C, and SW-C under smaller ε values, specifically $\varepsilon = 0.4$ and $\varepsilon = 0.8$. In these scenarios, all three mechanisms approach the uniform distribution more closely compared to cases with larger ε. Consequently, their errors are also more similar to each other. Statistically, when $\varepsilon = 0.4$, the error of OGPM is at most 0.008 smaller than that of PM-C and SW-C. For $\varepsilon = 0.8$, the error of OGPM is at most 0.015 smaller than that of PM-C and SW-C.

References

[1] Hai Brenner and Kobbi Nissim. Impossibility of differentially private universally optimal mechanisms. *SIAM J. Comput.*, 43(5):1513–1540, 2014.

[2] Natasha Fernandes, Annabelle McIver, and Carroll Morgan. The laplace mechanism has optimal utility for differential privacy over continuous queries. In *36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, Rome, Italy, June 29 - July 2, 2021*, pages 1–12. IEEE, 2021.

[3] Quan Geng and Pramod Viswanath. The optimal mechanism in differential privacy. In *2014 IEEE International Symposium on Information Theory, Honolulu, HI, USA, June 29 - July 4, 2014*, pages 2371–2375. IEEE, 2014.

[4] Arpita Ghosh, Tim Roughgarden, and Mukund Sundararajan. Universally utility-maximizing privacy mechanisms. In Michael Mitzenmacher, editor, *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 351–360. ACM, 2009.

[5] Moritz Hardt and Kunal Talwar. On the geometry of differential privacy. In Leonard J. Schulman, editor, *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 705–714. ACM, 2010.

[6] Naoise Holohan, Spiros Antonatos, Stefano Braghin, and Pól Mac Aonghusa. The bounded laplace mechanism in differential privacy. *J. Priv. Confidentiality*, 10(1), 2020.

[7] Zitao Li, Tianhao Wang, Milan Lopuhaä-Zwakenberg, Ninghui Li, and Boris Skoric. Estimating numerical distributions under local differential privacy. In David Maier, Rachel Pottinger, AnHai Doan, Wang-Chiew Tan, Abdussalam Alawini, and Hung Q. Ngo, editors, *Proceedings of the 2020 International Conference on Management of Data, SIGMOD Conference 2020, online conference [Portland, OR, USA], June 14-19, 2020*, pages 621–635. ACM, 2020.

[8] Ning Wang, Xiaokui Xiao, Yin Yang, Jun Zhao, Siu Cheung Hui, Hyejin Shin, Junbum Shin, and Ge Yu. Collecting and analyzing multidimensional data with local differential privacy. In *35th IEEE International Conference on Data Engineering, ICDE 2019, Macao, China, April 8-11, 2019*, pages 638–649. IEEE, 2019.

[9] Benjamin Weggenmann and Florian Kerschbaum. Differential privacy for directional data. In Yongdae Kim, Jong Kim, Giovanni Vigna, and Elaine Shi, editors, *CCS '21: 2021 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, Republic of Korea, November 15 - 19, 2021*, pages 1205–1222. ACM, 2021.