AI Olympiad, February 2026

1. During the development of a Large Language Model (LLM), a team implements "Differential Privacy" (DP) as an ethical design constraint to prevent the leakage of Personally Identifiable Information (PII). This process involves a privacy budget parameter ϵ.

   Which of the following best describes the resulting impact of this constraint on the system's performance?

   (a) The model's utility remains constant, but the training time increases exponentially to handle the privacy layers.

   (b) There is an inherent trade-off where increasing the privacy guarantee (decreasing ϵ) necessarily degrades the model's accuracy and convergence rate.

   (c) The constraint acts as a regularization term, improving generalization by preventing the model from over-fitting on specific data points.

   (d) Ethical constraints like Differential Privacy are purely architectural and do not affect the loss function or the final weights of the model.

   [3 Marks]

[Total 3 Marks]