

ABSTRACT

- Applied **DP** to **both TabSyn stages** (VAE + latent diffusion) and benchmarked against FinDiff+DP on UCI Credit Default (**n=30,000**); 42 runs passed fail-closed DP auditing.
- At **ε≈10**, DP-TabSyn achieves **ROC-AUC 0.674** vs. **0.578** for FinDiff+DP — a **16.6% advantage** and the **strongest privacy-constrained** result among evaluated models.
- The **privacy–utility frontier** is **tunable**: AUC rises from **0.624** (ε≈2) to **0.715** (ε≈20), recovering up to **87.4%** of **non-private signal**.
- Deployment structure and data quality matter**: consortium fragmentation is associated with a **0.124 AUC gap** versus **row-matched central training**, while **missingness mechanism** materially changes utility at the **same missingness rate**.

I. DATASET & EVALUATION

- UCI Credit Default**: 30,000 clients · 23 predictors + binary default target · 22.1% default rate.
- Split**: 24,000 train | 3,000 validation | 3,000 held-out test (hash-locked).
- XGBoost on synthetic data**, tested on real held-out set; type-aware preprocessing; 36-config grid; validation-only selection.
- Fidelity**: KS, Wasserstein, JS divergence, correlation MAE.
- Subgroup disparities** evaluated across **SEX, EDUCATION** and **MARRIAGE**.

II. STATISTICAL FIDELITY

Metric	Real	TabSyn	DP-TabSyn	FinDiff+DP
Default rate	22.1%	21.7%	19.9%*	22.1%†
Mean KS ↓	—	0.066	0.198	0.145
Norm. Wasserstein ↓	—	0.182	0.805	3.122
Cat. JS Divergence ↓	—	0.002	0.009	0.088
Correlation MAE ↓	—	0.032	0.197	0.245
Target Prev. Error	—	0.004	0.023*	0.000†


† Limitation: DP-TabSyn target prevalence is seed-sensitive (10.2%–26.2% vs. 22.3% real), despite strong downstream utility.

† FinDiff target counts fixed by construction (exact real counts); not learned fidelity.

Wins on 3 of 4 Fidelity Dimensions

DP-TabSyn yields lower Wasserstein (0.81 vs 3.12), lower JS divergence & lower correlation error

III. FIDELITY PROFILE




- FinDiff+DP: extreme **tail inflation** (PAY_AMT2 99th pctl. ~940K vs 73K real)
- DP-TabSyn **avoids** this **distortion** — **critical for portfolio and regulatory use**

IV. MODELS


Model	Core Design	Primary Role	Privacy Style
TabSyn	VAE + latent diffusion	Utility-focused synthesis	Empirical (DCR)
DP-TabSyn	VAE + latent diffusion + FastDP (both stages)	Privacy-constrained utility synthesis	Formal ε-DP (ε≈10, δ=2e-5)
FinDiff Vanilla	Diffusion in input space	Non-private financial baseline	Empirical (DCR)
FinDiff+DP	Strict DP diffusion in input space	Strongest memorisation control	Formal ε-DP (Opacus ε=9.99)
FinDiff+DP + Mask	DP diffusion + attention masking	Privacy-oriented baseline	Formal ε-DP (Opacus ε=9.99)

V. ARCHITECTURE



- All features unified into **one continuous latent space**
- Both stages **privately optimised** via **FastDP** with automatic clipping

VI. MAIN RESULTS



Training Condition	ROC-AUC	Gini	PR-AUC	Wtd. F1
Real training data	0.7477	0.4955	0.5029	0.7964
TabSyn (non-private)	0.746	0.491	0.491	0.779
DP-TabSyn (ε≈10)	0.674	0.348	0.410	0.746
FinDiff Vanilla	0.711	0.422	0.447	0.767
FinDiff + DP	0.578	0.156	0.274	0.701
FinDiff + DP + Mask	0.561	0.122	0.255	0.696


± ranges (1σ across 3 seeds) available in supplementary table.

Highest Retention Private Generator

DP-TabSyn retains 70.8% of non-private above-chance AUC signal at ε≈10


(AUC 0.674 vs. non-private ceiling 0.746)

VII. PRIVACY–UTILITY FRONTIER




- Privacy has a quantifiable, tunable cost — **each ε increment** buys back **measurable utility**.
- ROC-AUC: 0.624 (ε≈2) → 0.674 (ε≈10) → 0.715 (ε≈20); signal **retention rises** from **50.6%** to **87.4%** of non-private TabSyn.
- Actionable: **select ε** to **target** a **specific signal-loss** threshold rather than treating privacy as a binary constraint.
- Outputs concatenated** → 24,000-row synthetic release; row-matched control separates aggregate row count from fragmented learning.

VIII. CONSORTIUM EXPERIMENT



- Centralised training dominates**: ROC-AUC 0.703 vs. consortium 0.580 (**Δ = 0.124**).
- The deficit is **not explained** by aggregate **row count**: **row-matched central training outperforms** the **consortium** in every seed (**0.073, 0.076, 0.222**).
- Recommendation**: centralise where **governance permits**; use consortium only when data cannot leave institutional boundaries.

X. MISSINGNESS ROBUSTNESS



- Rate alone** is **insufficient**: **mechanism and deployment conditions** determine how much **predictive information** is **lost**.
- MCAR/MAR/value-MNAR** retain **67–78% of signal** at **85%** predictor-cell corruption; **outcome-MNAR destroys signal** (–3.1% retention).
- Institutions** must **characterise** their data's **missingness mechanism** before trusting synthetic output quality.

XI. PRIVACY VERIFICATION & SCOPE

Experiment	Runs	ε/stg	VAE Cov.	Diff. Cov.
Principal DP-TabSyn	3	5	35/35 ✓	5/5 ✓
Privacy frontier	6	1, 10	35/35 ✓	5/5 ✓
10-inst. consortium	30	5	35/35 ✓	5/5 ✓
Row-matched control	3	5	35/35 ✓	5/5 ✓

- 42 verified runs**; fail-closed: mismatch → run rejected
- Fixed schedules**, predetermined checkpoints; no adaptive LR or early stopping
- Schema **projection** uses **public constraints only**

Condition	ε (seed)	δ	End-to-End DP?
DP-TabSyn	10.00	2e-5	Yes — nominal composed
FinDiff + DP	9.99	2e-5	No — generator only†
FinDiff + DP + Mask	9.99	2e-5	No — generator only†

† Opacus budget applies to generator only

- FinDiff target counts fixed outside model — end-to-end DP not claimed unless counts are public
- DP-TabSyn: fixed schedules, predetermined checkpoints, public-only preprocessing

XII. CONCLUSIONS

AUC 0.674 vs 0.578

DP-TabSyn outperforms next-best DP model by +16.6%

- DP-TabSyn is the preferred privacy-constrained generator** — AUC 0.674 vs. 0.578, with stronger fidelity across Wasserstein, JS divergence and correlation metrics.
- Privacy budget** is **actionable**: AUC rises from 0.624 (ε≈2) to 0.715 (ε≈20), recovering up to 87.4% of non-private signal.
- Fragmentation-associated penalty**: **consortium AUC 0.580** vs. **row-matched central 0.703**.
- Rate alone is insufficient**: at 85% corruption, MCAR/MAR/value-MNAR retain 67–78% of signal while outcome-MNAR falls below chance.
- Scope**: single credit-default dataset (**n=30,000**); **cross-domain** generalisation remains **to be validated**.

XIII. FUTURE WORK

- Rényi-DP accounting** for tighter composition; persist accountant state for realised ε recovery.
- Native mixed-type output heads** and **target-aware** conditional generation for prevalence stability.
- Multi-dataset validation** with **membership inference auditing**; parameter-sharing consortium protocols.

XIV. PRACTICAL GUIDANCE

- Among the **privacy-constrained** implementations evaluated, **DP-TabSyn** offers the **strongest overall balance** of predictive **utility** and statistical **fidelity**.
- ε≈10** provides a **useful middle operating point** in this study; deployment budgets should be selected from **application-specific privacy** and **utility requirements**.
- Centralise training** where governance permits — consortium fragments impose **Δ=0.124 AUC penalty**.
- Profile missingness mechanism before synthesis**: outcome-MNAR requires mechanism-aware preprocessing before any generator can produce trustworthy output.