

Package & Premium Optimization (Direct Mail)

Outcomes we typically see*

-50%

Cost per \$ Raised

-62%

Cost per Piece

+9%

Net Revenue Lift

What this use case delivers

- Use model-guided offers and copy to lift response — creative without guesswork.
- Balance premium cost against incremental lift by cohort.
- Learn from historical data to narrow tests and move faster.
- Control frequency to avoid donor fatigue.

AI signals powering it

- Offer & premium fit: choose the incentive that pays back.
- Elasticity: estimate value changes with price or ask.
- Frequency guardrails: hit the sweet spot by segment.
- ROI forecast: incremental lift and net by package.

How it works

- Models score offers/premiums by segment and predict net revenue at package level.
- Suppression reduces cost by removing low-ROI packages before production.

Implementation (days, not months)

- Connect historical donor, campaign/appeal, cost, and channel data.
- AutoML trains models with validation; you get donor-level scores and drivers in plain language.
- AI agents push segments and next-best actions into your CRM, ESP, print, and ad tools.
- Launch fast with no-code setup and human-in-the-loop approvals where needed.

* Directional examples; actual results depend on list, offer, channel mix, and scale.

Package & Premium Optimization (Direct Mail) (cont.)

Data & integrations

- Donor master & gift history with unique IDs and dates.
- Campaign/appeal tables, list sources, package and premium costs.
- Channel touches (mail, email, SMS, web), and optional enrichment (co-ops, demographics).
- SFTP/API/cloud-DWH access; read-only connections to activation tools.

Governance

- Human-in-the-loop approvals for high-impact actions and content.
- Least-privilege data access, expirations, and audit logs for decisions and prompts.
- Bias monitoring and calibration checks; versioned models with rollback.

Measurement

- Track net revenue, CPDR, response rate, average gift, reactivation/upgrade rates, and time-to-first gift.
- Use randomized holdouts and pre-registered test plans; report lift with confidence intervals.
- Explainability: show top drivers and calibration so teams can answer the 'why'.

Next steps

- Send a recent drop; receive model-guided changes with projected lift and savings.
- Embed learning in future tests for faster cycles.