

USING INCENTIVES TO DRIVE PIPELINE

APPENDIX B: DETAILED EXPLANATIONS OF THE METHODS OF MEASUREMENT

Five Methods for Estimating Incrementality

Note: Creating true control groups would require deliberately withholding the incentive from some eligible partners, which would be difficult to defend to sales leadership. In practice, organizations often use participant vs. non-participant comparisons or trend projections.

Practical distinction: Methods A and B construct a comparison group (matched partners or regions). Method C compares opt-in participants vs. eligible non-participants (useful but vulnerable to self-selection). Method E uses random assignment and offers the strongest evidence, but is rarely appropriate unless relationship and leakage (control group learning about the incentives) risks are manageable.

Method A: Matched Control / Synthetic Control (Most Rigorous)

Create a control group of similar partners (matched on size, region, baseline growth, product focus) and compare performance before vs. during program implementation. This mirrors a traditional archival study design with a comparison group, analogous to propensity score matching in social science research (a technique used to create a fair comparison between two groups when you can't randomly assign people or partners to conditions). When selecting matching factors, consider running simple regressions or correlations in advance to identify which partner characteristics are most predictive of the outcome you care about. This is not causal proof, but it improves the quality of the match.

- Best When: You have sufficient partner volume and stable baseline data
- Calculation: $\text{Incremental Lift} = (\text{Program Group Growth \%} - \text{Control Group Growth \%}) \times \text{Program Group Baseline Revenue}$
- Example: Participants grew 18%, control grew 8%, baseline revenue \$10M. Therefore, Incremental lift = \$1M

A difference-in-differences (DiD) analysis is a useful extension of this approach, comparing the change in outcomes for program participants vs. non-participants before and after program launch, which helps control for pre-existing trends.

Method B: Geographic or Segment Holdout (Cleanest Operationally)

Pilot the program in Region A while holding out Region B, then roll forward and compare.

- **Best When:** Regions or segments are sufficiently comparable, and market conditions are similar
- **Advantage:** Clear before/after comparison without complex partner matching
- **Caution:** Ensure the holdout period accounts for seasonal variation

Method C: Participant vs. Non-Participant Comparison (Pragmatic)

When enrollment is opt-in, compare performance of participants versus eligible non-participants.

- **Best When:** Program has been running long enough to have meaningful participation rates
- **Caution:** Self-selection bias (motivated partners opt in) may overstate impact. Use with awareness of this limitation.
- **Example Finding:** “Program participants outperformed non-participants by approximately 25%” (Automotive case study in Appendix A)

Method D: Baseline Trend Projection (When Controls Are Unavailable)

Project what performance would have been based on historical growth trends, then compare actual results.

- **Best When:** No control group is available, but you have robust historical data
- **Method:** Calculate average growth rate for 3–5 years pre-program, project forward, compare to actual
- **Caution:** Market conditions or competitive dynamics can confound results. Document these factors explicitly.

Method E: Field Experiment / Randomized Assignment (Research Gold Standard — Rarely Practical in Channel Programs)

Randomly assign eligible partners to a program condition (treatment) or a business-as-usual condition (control). This is the only method that provides near-causal evidence of program impact.

- **Best When:** You need definitive proof of causation—for example, to justify a major scale-up or to defend a program under intense scrutiny
- **Practical Approach:** Randomly withhold the incentive from ~10–20% of eligible partners for the period, then report incremental sales or margin as the post-period difference between treatment and control, optionally adjusted for pre-period performance

- **Caution:** Requires confidence that the treatment is not diffused across partners (i.e., participants don't informally share program benefits with control partners). This is the most expensive and operationally complex method. While it provides the strongest evidence, it is not recommended for most channel programs due to relationship risks and operational complexity — reserve it only for high-stakes decisions where near-causal evidence is essential and the organizational relationship with partners can withstand it.
- **Caution:** Withholding meaningful incentives from a subset of eligible partners can create perceived unfairness, damage trust, and distort results if partners compare notes or shift business to competitors. If you pursue randomized testing, consider lower-stakes interventions (e.g., randomized communications, “encouragement” designs, or feature-level A/B tests) before experimenting with benefit denial, and involve legal/compliance and channel leadership up front.

Choosing the Right Method: A Practical Guide

The five methods vary significantly in rigor, cost, and organizational feasibility. The table below summarizes the key differences. Most programs should start with Method C or Method A, reserving Method E for situations where the stakes demand near-causal evidence.

Method	How the Comparison Group is Created	Key Risk	Best Starting Point?
A – Matched Control	You deliberately select similar non program partners at a comparison group before launch.	Matching is imperfect; Requires sufficient partner volume.	Yes: Most rigorous of the practical options.
B – Geographic Holdout	You pilot in one region and hold out in another.	Regions may not be comparable; Seasonal confounds.	Yes: operationally clean when regions are similar.
C – Participant vs. Non-Participant	You compare partners who opted in to eligible partners who did not..	Self-selection bias: motivated partners join, inflating apparent impact.	Yes: Most available in existing programs.
D – Baseline Trend Projection	You project historical growth forward and compare to actuals	Market shifts can confound results.	When a/ B/ C/ E are unavailable.
E – Randomized Assignment	You randomly withhold the incentive from some eligible partners	Partners who discover they are excluded may react negatively, shift business, or compare notes; Results are distorted if treatment diffuses.	High stakes decisions only.

The key distinction between Methods A, C, and E: All three compare participants to non-participants, but the mechanism differs critically.

Method C: the comparison group self-selects: partners who were eligible chose not to join. The program was available to everyone; some opted out. This is the easiest to run but introduces self-selection bias — partners who participate tend to be more motivated to begin with, which inflates the apparent effect.



Method A: you construct the comparison group deliberately: you identify non-program partners who closely resemble program participants in size, region, baseline growth, and product mix, and track them in parallel. This reduces self-selection bias because the comparison is based on deliberate matching, not voluntary choice.

Method E: you randomly withhold the program from some eligible partners. No one chooses whether to be in the comparison group — assignment is random. This is the only method that approaches true causal evidence, but it requires withholding a benefit from partners who would otherwise qualify, which can damage trust and relationships if partners become aware. It is expensive, operationally complex, and carries the risk of results contamination if treated and control partners interact.

In practice: start with Method C if the program is already running. Use Method A when launching a new program and you have sufficient partner volume to build matched pairs. Use Method B when you have geographically separable markets. Reserve Method E for high-stakes decisions where causal evidence is essential and the organizational relationship with partners can withstand the design.

One additional option worth noting for organizations with CRM or deal-registration workflows: lightweight “attribution tagging” – adding an “influence” field to deals (e.g., “Was this deal influenced by the program? Which activity mattered most?”) – provides qualitative directional data alongside the quantitative methods above. It is not a substitute for incrementality measurement but can surface patterns that strengthen the broader case. See Appendix B for more detailed descriptions of Methods A to E.