Accurate Valuation and Effective Targeting of Low Response Programmes using Propensity Score Matching Techniques

Credit Scoring and Credit Control XII

Kevin Chisholm

August 2011

Capital One is a successful global bank, specialising in credit cards

Globally	UK
• Top 10 US bank • 40 million customers • Offices in US, UK & Canada	• 800 employees • Top 10 credit card issuer • Nottingham & London offices

In our short history, we have expanded from a small regional credit card provider to a fully diversified international bank

1988 – Book first customers

1994 – Capital One IPO

1996 – Expand into Canada and UK

1998 – Acquire auto-lending business

2005 – Acquire Hibernia Bank

2006 – Acquire North Fork Bank

2009 – Acquire Chevy Chase Bank

Today – 40m customers worldwide; top ten US Bank

Contents

Balance build programmes drive mutually beneficial long term customer relationships

Balance Build Programme Valuations

Super prime customers can choose between a number of cards in their wallet

If they lose sight of the benefits of our card, spend will reduce and they may become completely inactive

We mail low-APR balance transfer (BT) offers to valued customers

If they accept, they are likely to continue using our card beyond the end of the BT offer

For many customers, the proposition is not compelling and the response rate from mailing everyone would be low.

By intelligent targeting, we can reduce programme costs and reduce the number of unwanted mailshots.

Although there is positive financial benefit from these campaigns, this has been difficult to demonstrate in our monitoring

Balance Build Programme Valuations

Cumulative Profit per Original (Test-Control)

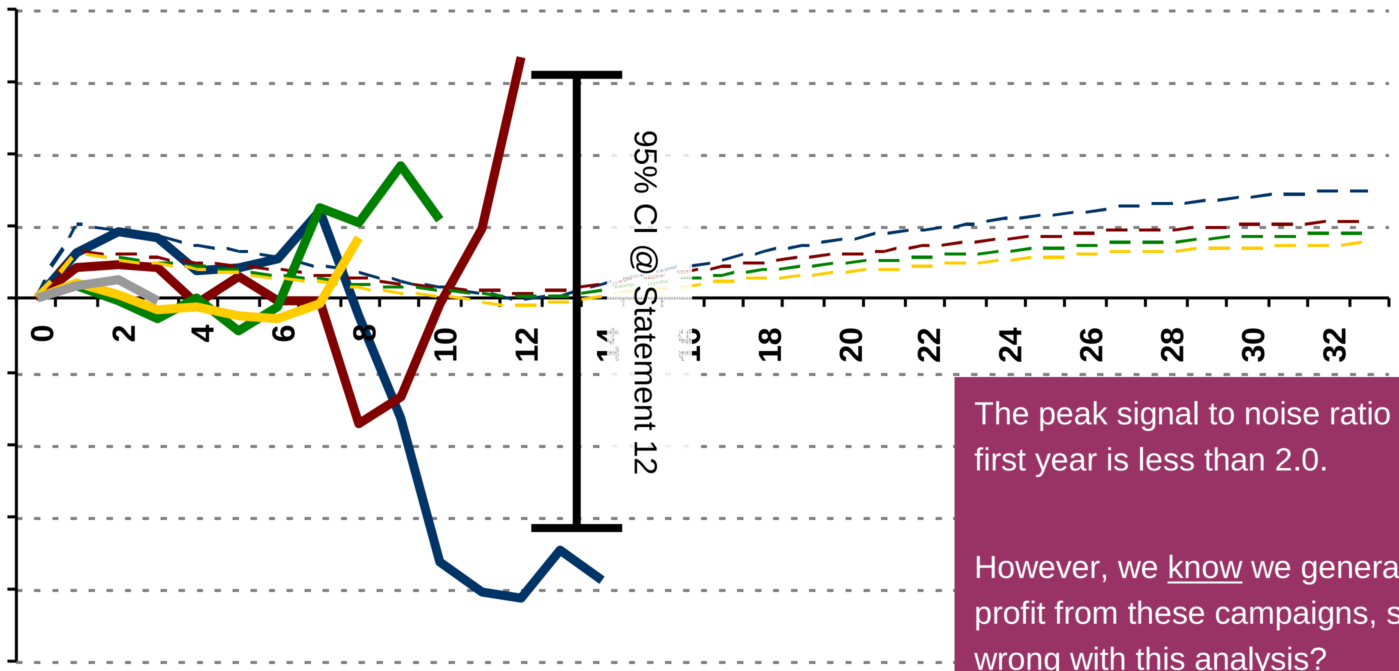

Statement

The peak signal to noise ratio (μ/σ) in the first year is less than 2.0.

However, we <u>know</u> we generate consistent profit from these campaigns, so what is wrong with this analysis?

Key:

Dashed line – Programme assumption

Dotted line – Programme actuals

With low response rates, the vast majority of the accounts we monitor are not directly contributing to the value of the campaign

By averaging across the whole population, the responders' high incremental value signal is diluted due to the low overall response probability.

$$E\{\Delta v\} = E\{\Delta v \mid R\}P_R + E\{\Delta v \mid R'\}(1 - P_R)$$

If we hypothesise that a non-responder who receives a mailing:

•does not stimulate their spend,

•does not reduce their spend, and

•does not modify their risk behaviour,

then non-responders contribute **zero** to incremental value.

Their variance does contribute to the standard error in calculating the mean campaign value!

$$\mathrm{var}\{\Delta v\} = \left(\mathrm{var}\{\Delta v \mid R\} + E\{\Delta v \mid R\}^2\right)P_R + \left(\mathrm{var}\{\Delta v \mid R'\} + E\{\Delta v \mid R'\}^2\right)(1 - P_R) - E\{\Delta v\}^2$$

We would also like to improve targeting so that exposure is better matched to customer needs

Balance Build

Programme Valuations

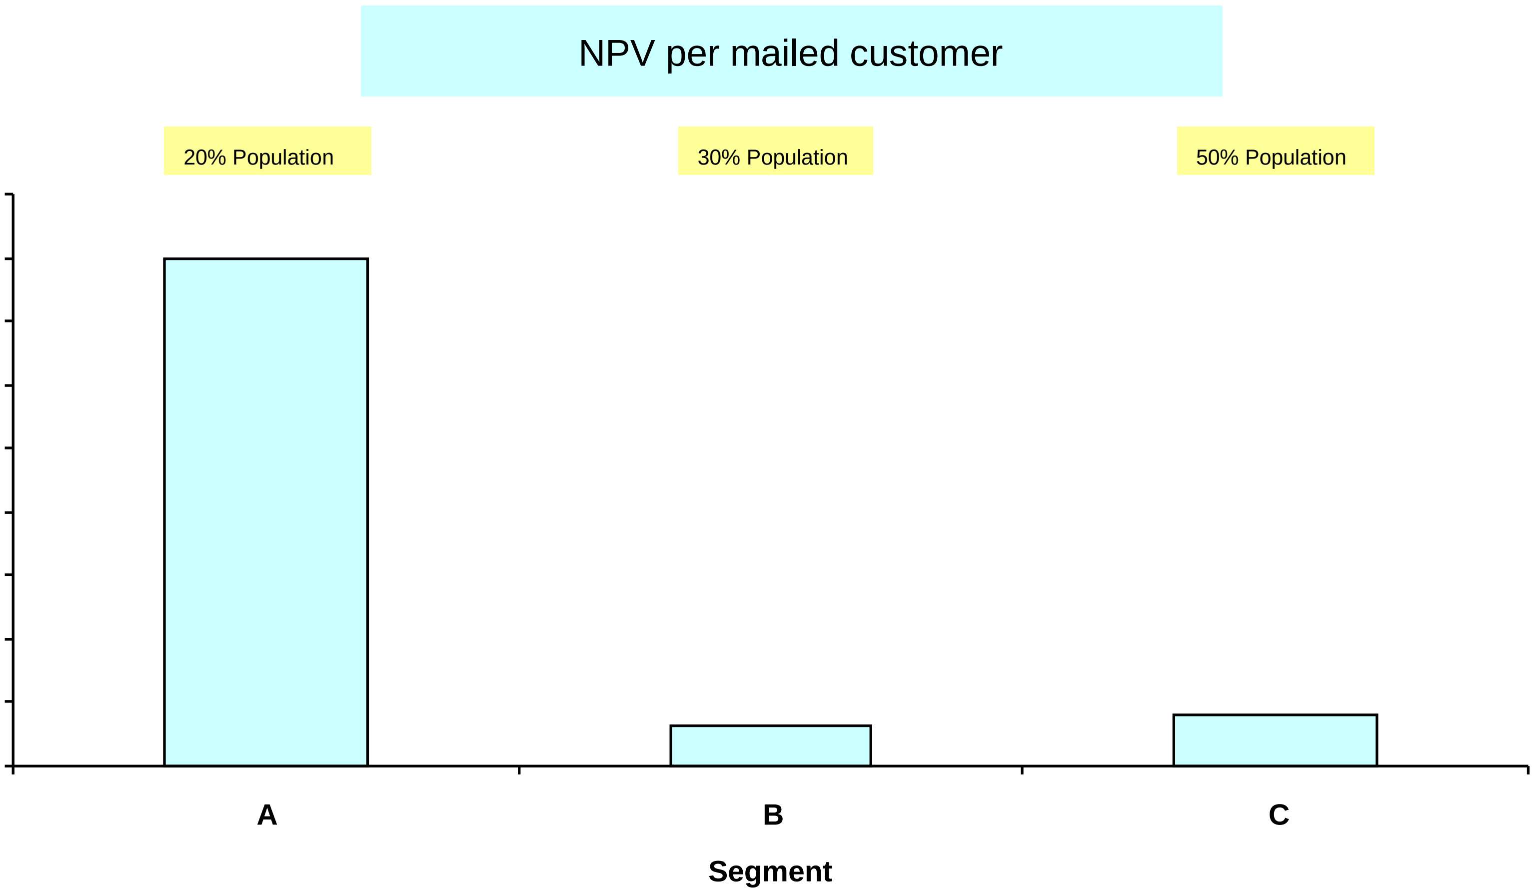

Contents

If we could focus on the responder population, we stand to win improved signal to noise

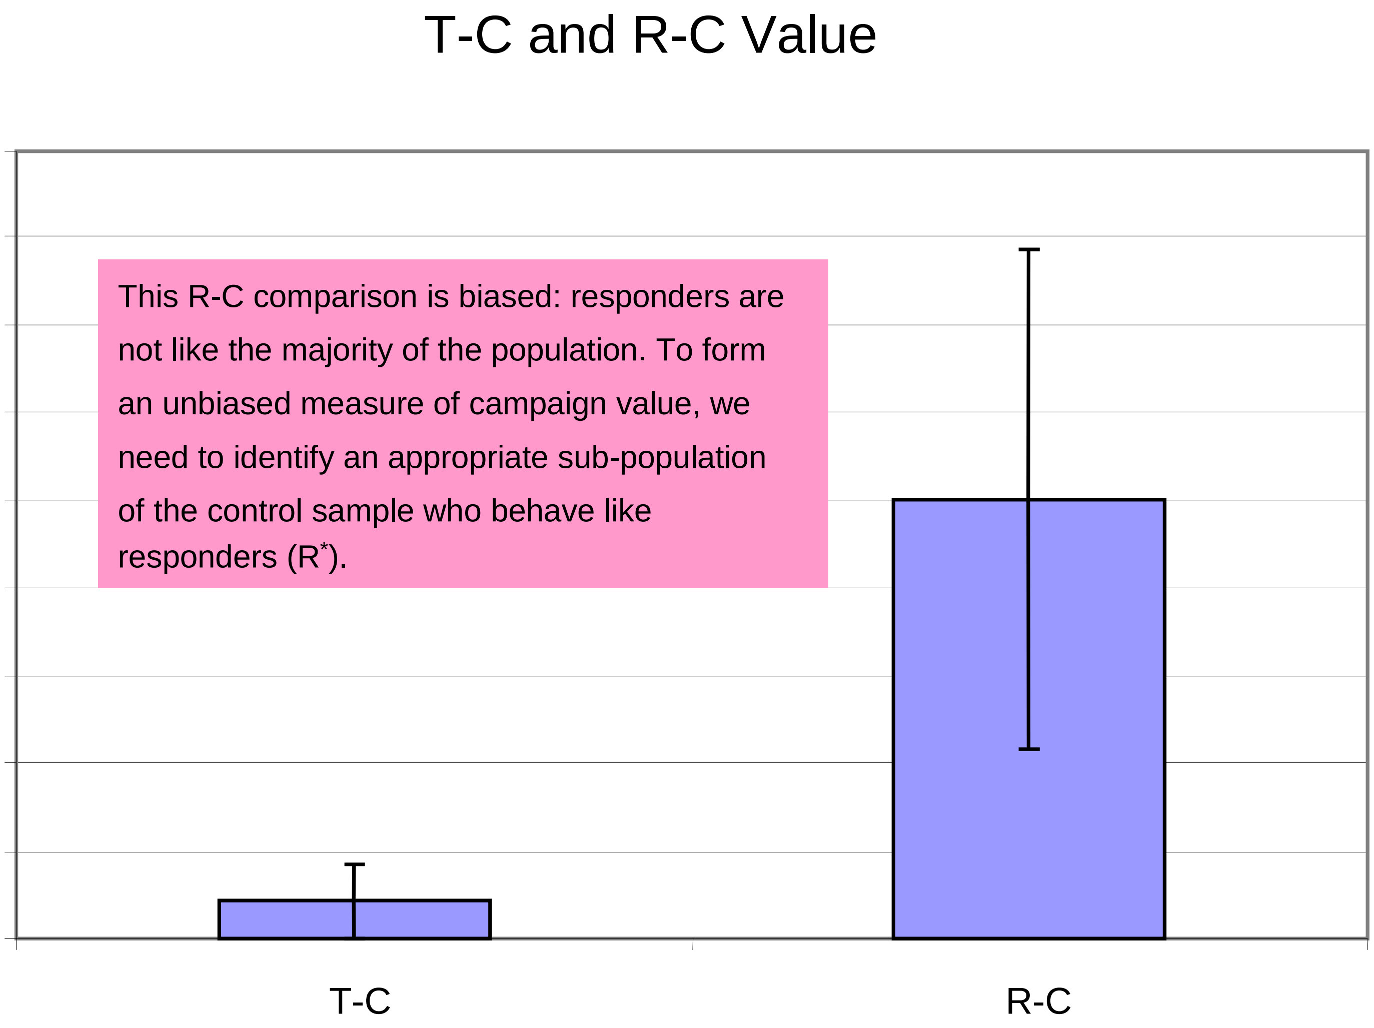

There is a strong analogy with problems related to non-random treatment assignment

Relevance of Propensity Score Matching

Non-random treatment assignment	Non-random response to mailing
Observational studies or imperfectly randomised tests have treatments whose assignment may be predicted by observed covariates.	A well-randomised treatment is assigned, but we are only interested in responder behaviour. Response may be predicted by observed covariates.
The predicted probability of treatment is known as a “propensity score”	We predict probability of response and refer to this as the “response score”
Matching propensity scores for treated and untreated units approximates a randomised test and **balances covariate distributions**	Matching response scores for mailed and unmailed customers identifies a “fair” control sample and **balances covariate distributions**

If there is good overlap and a suitably large control sample, we can obtain matched score distributions without exhausting the control sample first

Response Score Distributions

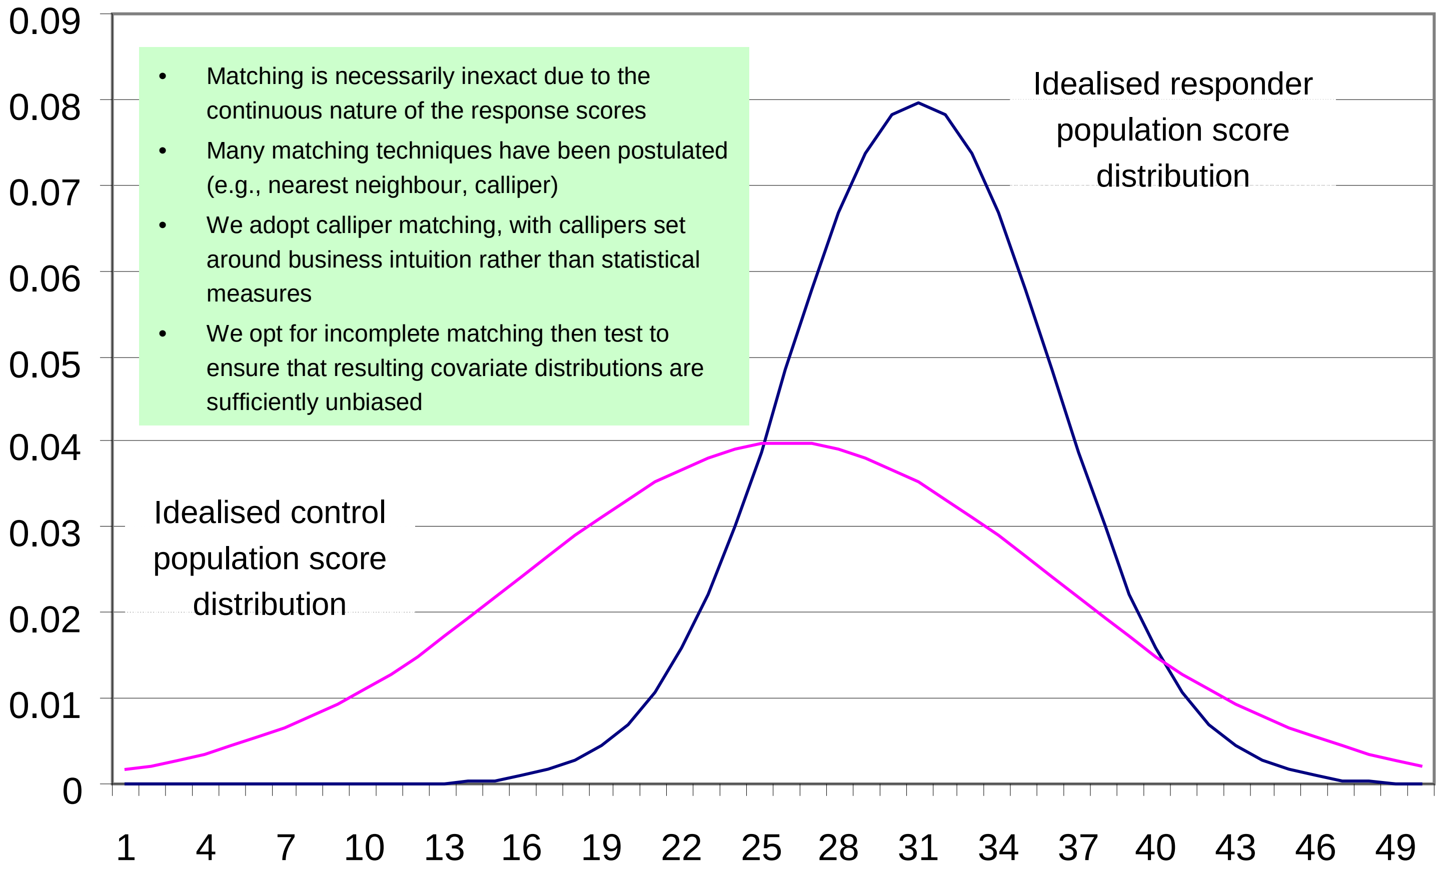

Contents

- **Balance Build Programme Valuations**
  - Overview
  - Challenges
- **Relevance of Propensity Score Matching**
  - Analogy with non-random treatment assignment
  - Practical implementation
- $\boxed{\textbf{Incremental Value Segmentation}}$
  - Pair-wise matching for synthetic account level valuation
  - Regression tree analysis
- **Conclusions**

To improve the campaign segmentation, we use decision trees to identify homogeneous value subpopulations

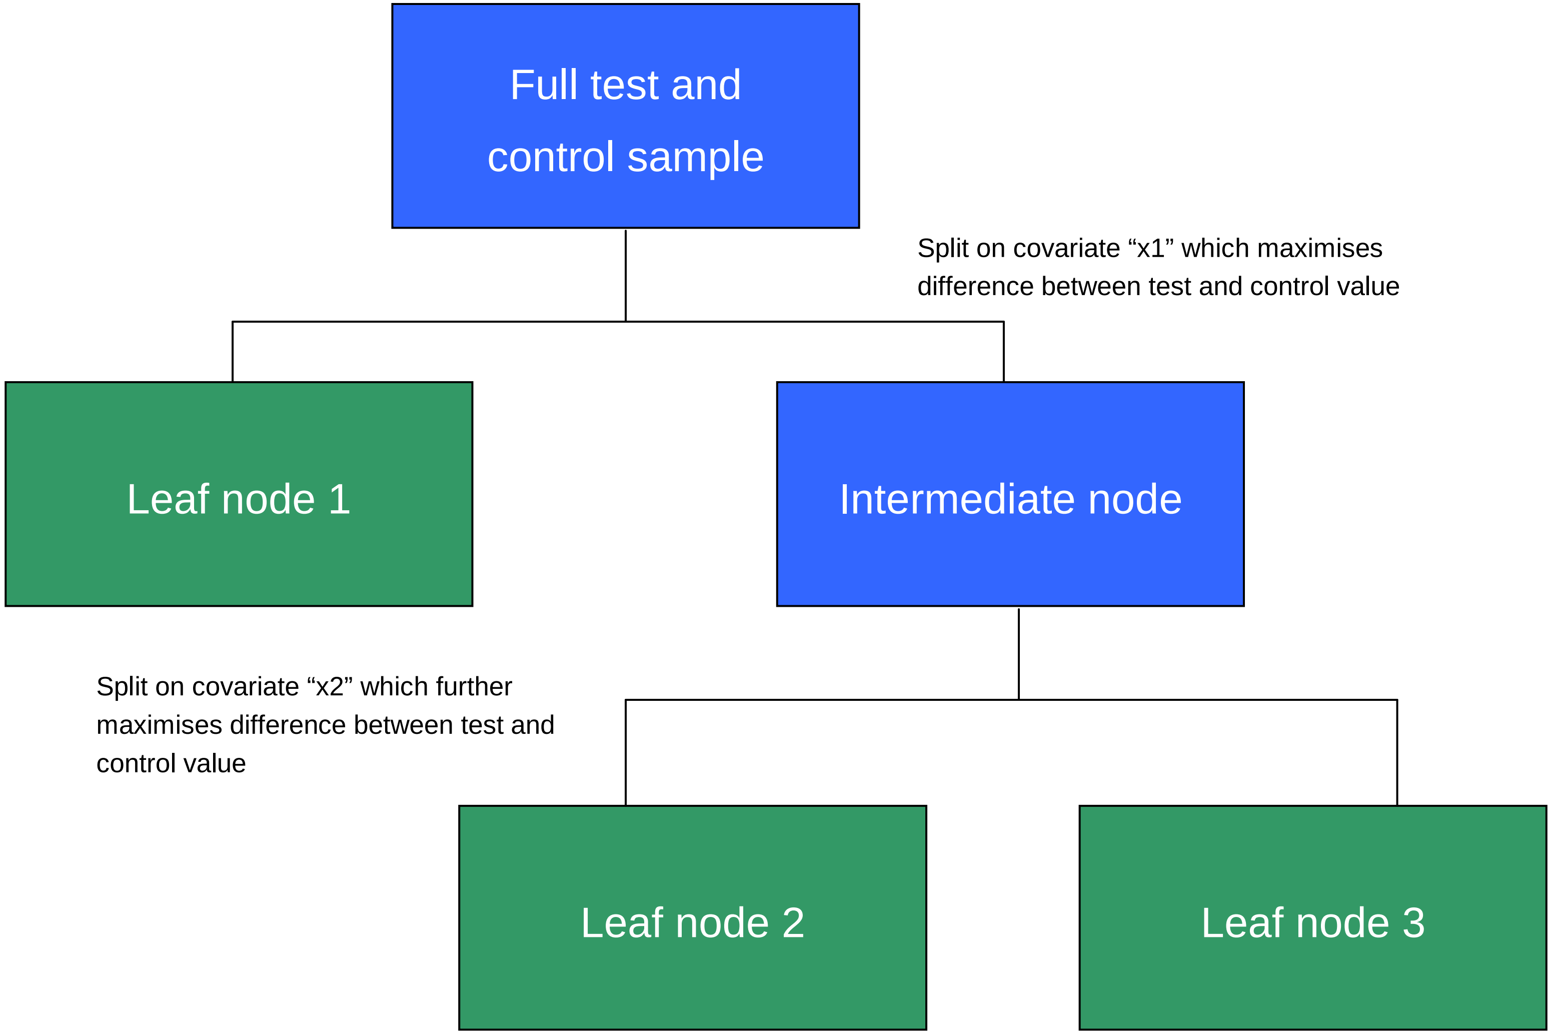

We identify matched test and control pairs and treat as synthetic individuals

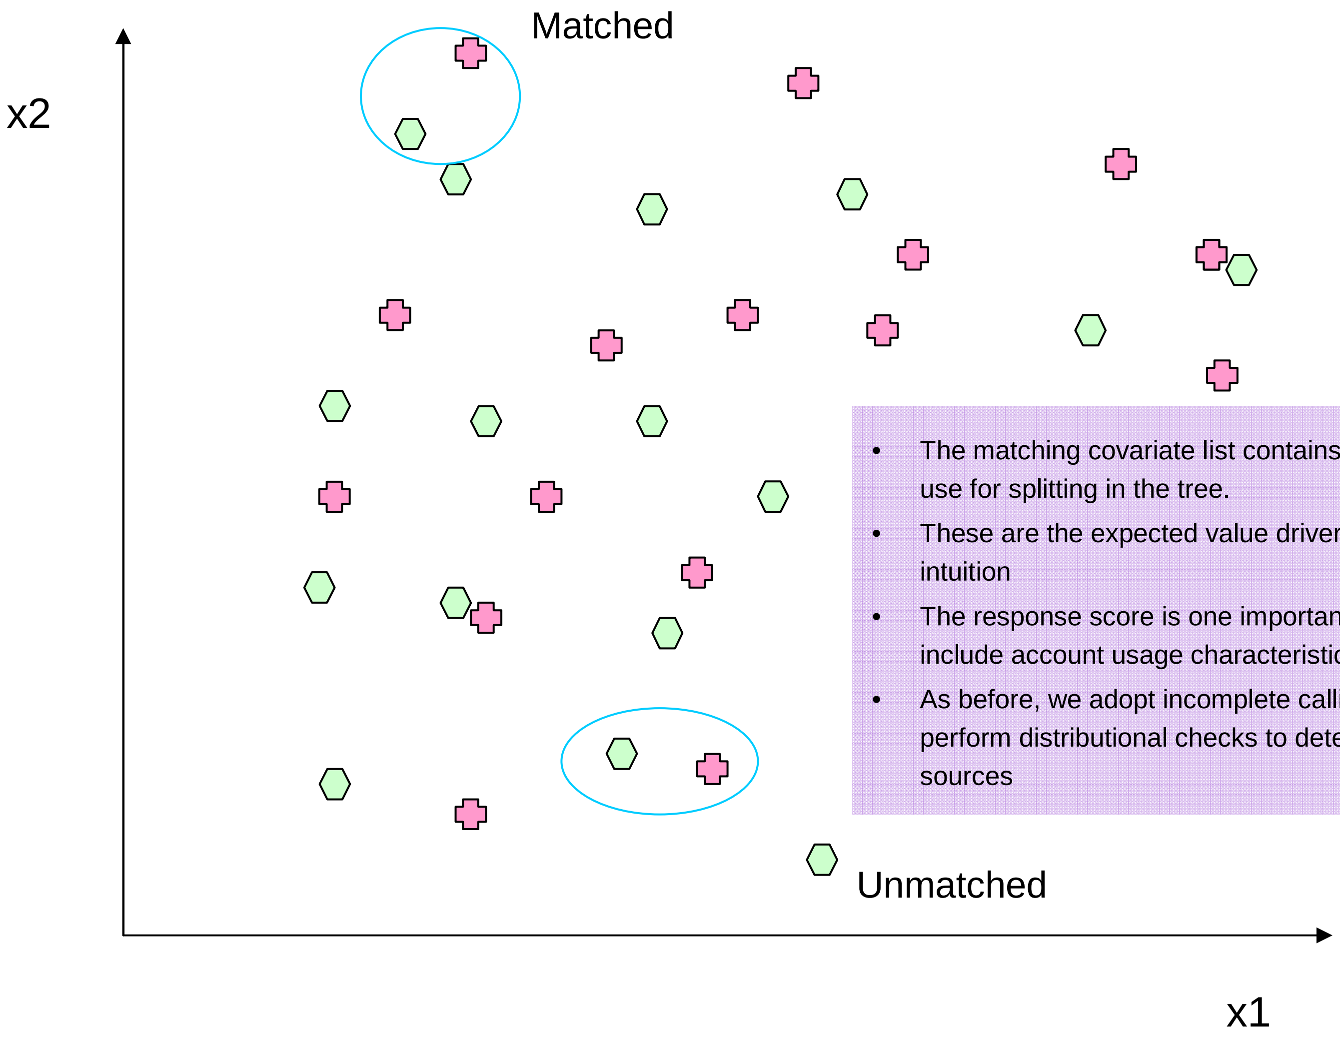

- The matching covariate list contains all variables we may use for splitting in the tree.
- These are the expected value drivers based on business intuition
- The response score is one important variable; others include account usage characteristics and bureau data
- As before, we adopt incomplete calliper matching and perform distributional checks to detect potential bias sources

All important value metrics are sloped effectively (and with statistical significance)

Incremental Value Segmentation

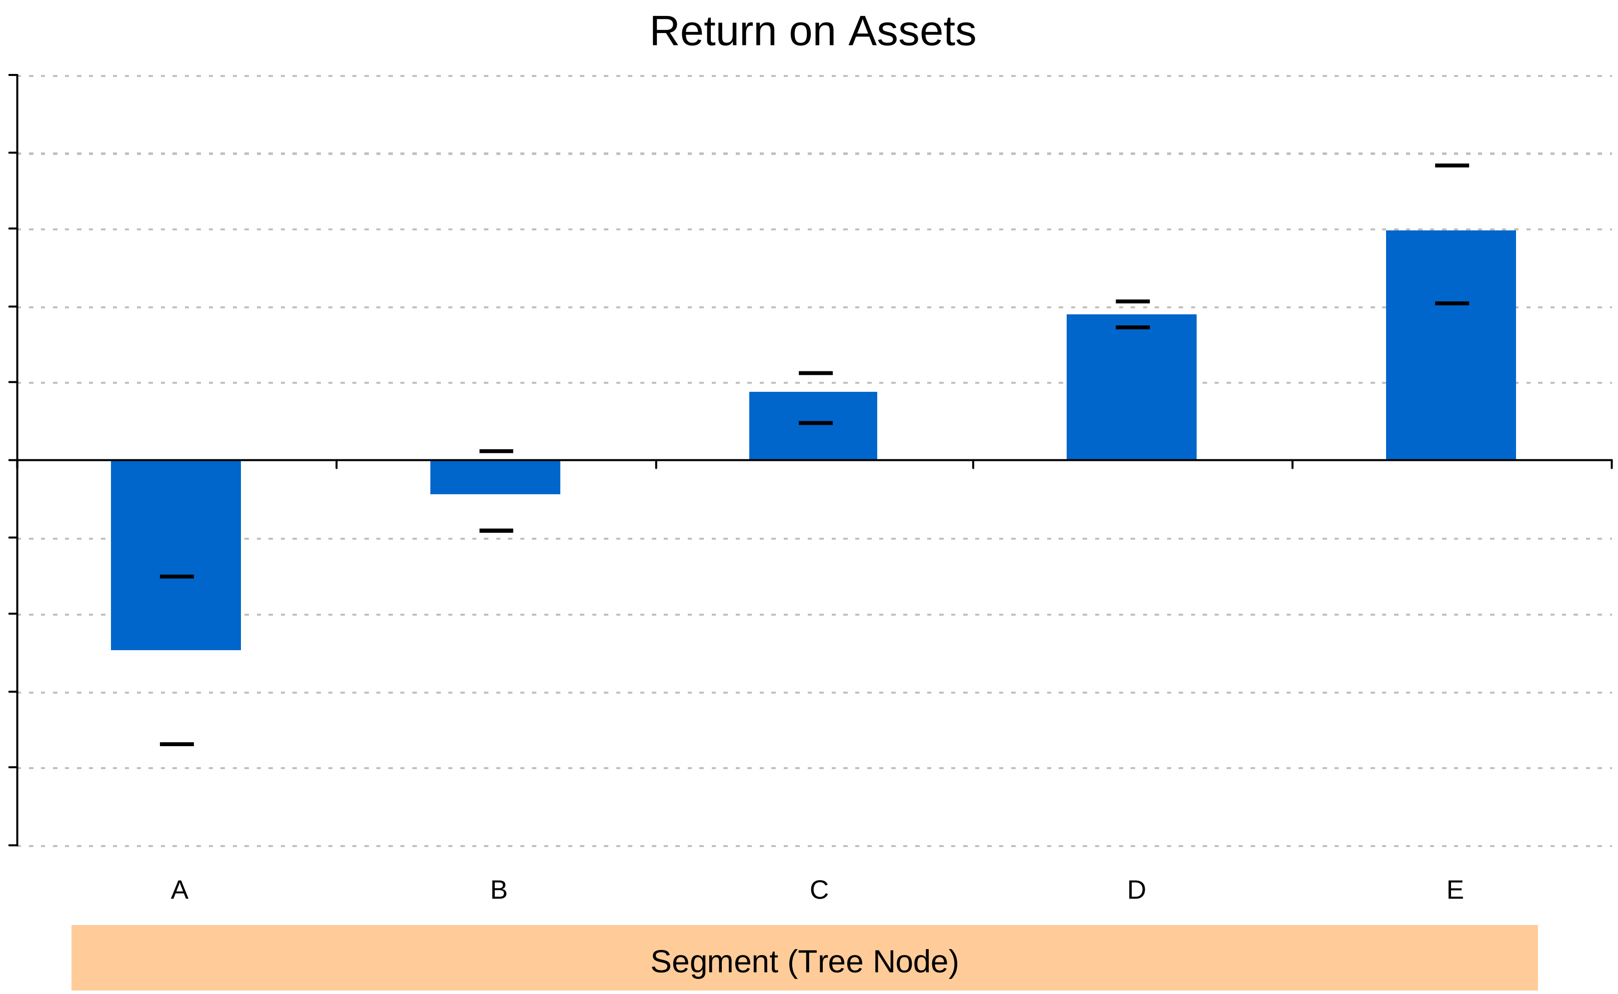

Contents

- **Balance Build Programme Valuations**
  - Overview
  - Challenges
- **Relevance of Propensity Score Matching**
  - Analogy with non-random treatment assignment
  - Practical implementation
- **Incremental Value Segmentation**
  - Pair-wise matching for synthetic account level valuation
  - Regression tree analysis
- **Conclusions**

We are excited about the value driven from the R-R* framework and are looking to apply it to other programmes

- **Improved signal to noise characteristics for monitoring low response rate programmes can be obtained by focusing on the responder population**

- **Response score matching is an effective technique for generating "fair" control samples**

- **Creating synthetic "T-C accounts" is an effective method allowing data mining techniques to be applied to a number of business problems**
  - We recognise that matching on too many dimensions will rapidly exhaust the data
  - Good business insight is required to manage the initial covariate list