International Journal of Forecasting 28 (2012) 248–260

Contents lists available at SciVerse ScienceDirect

International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecast

Estimating causal effects of credit decisions

Gerald Fahner *

FICO, 901 Marquette Avenue, Suite 3200, Minneapolis, MN 55402, USA

ARTICLE INFO

Keywords:
Causality
Data mining
Decision making
Econometric models
Finance
Model selection
Nonparametric methods
Regression

ABSTRACT

In principle, making credit decisions under uncertainty can be approached by estimating the potential future outcomes that will result from the various decision alternatives. In practice, estimation difficulties may arise as a result of selection bias and limited historic testing. We review some theoretical results and practical estimation tools from observation study design and causal modeling, and evaluate their relevance to credit decision problems. Building on these results and tools, we propose a novel approach for estimating potential outcomes for credit decisions with multiple alternatives based on matching on multiple propensity scores. We demonstrate the approach and discuss results for risk-based pricing and credit line increase problems. Among the strengths of our approach are its transparency about data support for the estimates and its ability to incorporate prior knowledge in the extrapolative inference of treatment-response curves.

1. Introduction

Making credit decisions poses a classical problem of decision making under uncertainty: treatments for individuals must be selected on the basis of estimates of potential future outcomes resulting from treatment alternatives. A growing number of organizations have developed or aspire to develop generalizations of traditional scoring models for delivering such estimates for business objectives such as response, revenue, profit, prepayment, attrition, default and loss. Such models have been referred to as profit scoring models, action-effect models, or, in a wider sense, decision models (Marshall & Oliver, 1995; Rosenberger & Nash, 2009, p. 138; Thomas, 2009, p. 204). The models are developed to predict individuals' potential future outcomes, based not only on characteristics describing individuals or accounts, as is the case with traditional scores, but also on the applicable treatments for decisions such as product offers, pricing, credit limits, authorizations and collection actions.

The main problem is to estimate likely future outcomes for individuals or groups as functions of alternative treatments, conditional on the individual or group being held fixed. This problem is shared across research areas ranging from the medical to social science fields. A significant body of statistical and econometric work has been dedicated to the characterization of the theoretical properties of the problem, and the demonstration of the conditions under which accurate estimates can be found. Based on these properties, powerful tools for tackling the estimation problems have been developed. One basic problem concerns the estimation of the average effect of a binary treatment alternative on a single outcome of interest. For example, does an alternative style of collection letter nudge a larger share of past-due customers to make their payments? The basic problem can be refined to estimating the treatment effects for multiple treatment alternatives, and conditional on groups or even individuals.

The plan of this paper is as follows: The next section discusses the potential outcomes framework which was first applied to the study of causation by Rubin (Holland, 1986), related practical estimation tools of propensity scoring and matching, and some extensions and applicability to credit operations. Section 3 outlines our approach to conceptually estimating the effects of credit decisions. Sections 4 and 5 give more details and present the case study results.

* Corresponding address: 4443 Stony Meadow Lane, Austin, TX 78731, USA.
E-mail address: geraldfahner@fico.com.

0169-2070/$ – see front matter
doi:10.1016/j.ijforecast.2010.10.002

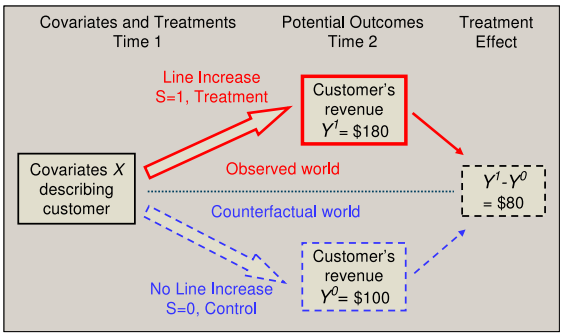

Fig. 1. A customer receives a credit line increase. Later, we observe the customer's account revenue. In a counterfactual world, the customer gets no increase. Boxes marked by dashed lines are unobserved. The diagram is not intended to imply that the treatment is selected based on the measured covariates only, although we will make this assumption later.

The style of our presentation is somewhat analogous to that of Rubin and Waterman (2006), while our work clearly extends beyond that of their article.

2. Problem formulation

2.1. Rubin causal model

The simplest formulation concerns dichotomous treatments, referred to in the following as 'control' and 'treatment'. We are given units (individuals, accounts) $i = 1, \ldots, N$. For each unit, we posit a pair of potential outcomes, Y_i^0, Y_i^1, under control and treatment, respectively. Each unit receives either control or treatment, as indicated by the treatment indicator ($S_i = 0$ if control, $S_i = 1$ if treated). Y_i is the observed outcome, such that $Y_i = Y_i^0$ if $S_i = 0$, and $Y_i = Y_i^1$ if $S_i = 1$. Only one potential outcome is observed: the other, called counterfactual, is unobserved. The units are characterized by their observed covariates X_i, which are variables which are not influenced by the treatment, and indeed are typically measured prior to the treatment. The unit-level treatment effect is defined as the difference between the potential outcomes, $Y^1 - Y^0$. This is an unobservable random variable, because we never observe Y^0 and Y^1 together for the same unit. Fig. 1 illustrates these definitions for a credit treatment.

Consider the problem of estimating the average (expected) treatment effect $\theta \equiv E\left[Y^1 - Y^0\right] = E\left[Y^1\right] - E\left[Y^0\right]$. This is an expectation over an unobserved variable, and thus cannot be estimated directly. We could estimate $E[Y|S = 1] - E[Y|S = 0]$, but this will not generally equal the average treatment effect, because the two treatment groups may not be comparable in their covariate distributions prior to the treatment, which can lead to a severe selection bias (Table 5 in Section 4 provides an example). However, it is possible to estimate the average treatment effect based on assumptions of unconfoundedness and common support.

- Unconfoundedness (sometimes called conditional independence or selection on observables) requires $\left(Y^0, Y^1\right) \perp S|X$; i.e., the potential outcomes and treatment are conditionally independent given X.
- Common support (sometimes called overlap) requires that $0 < p(X) \equiv \Pr\{S = 1|X\} < 1$; i.e., at each value of X, there is a nonzero probability of receiving each treatment. $p(X)$ is called the propensity score.

Unconfoundedness implies:

$$E\left[Y^k|X = x\right] = E\left[Y^k|S = k, X = x\right] = E[Y|S = k, X = x]; \quad k \in \{0, 1\}.$$

Define $\theta(x) \equiv E\left[Y^1 - Y^0|X = x\right]$, the average treatment effect for the subpopulation at $X = x$. From unconfoundedness, $\theta(x) = E[Y|S = 1, X = x] - E[Y|S = 0, X = x]$. This expression can be estimated for any value of x with a nonzero and below one probability of receiving each treatment, which is granted by the common support condition. Given a small region in the covariate space around $X = x$ with common support, we can, at least in principle, estimate the local average difference of the outcomes between the treated and very similar (matching) control units. As a consequence, $\theta = E_X[\theta(x)]$ can be estimated. Conditioning on $X = x$ removes the bias due to observable covariates. In the above, common support was assumed globally. A weaker condition, which is still useful for estimating the effects of treatments, is for common support to only hold for a subset of the covariate space.

2.2. Matching on the propensity score

For practical applications, we need to obtain estimates from finite samples. It can be difficult, if not impossible, to find units with very similar values for all covariates. Among the useful results derived by Rosenbaum and Rubin (1983) is the finding that, if unconfoundedness holds, then $\left(Y^0, Y^1\right) \perp S|p(X)$ also holds, i.e. the potential outcomes and treatment are conditionally independent, given the propensity score. Thus, biases due to observable covariates can be removed by conditioning on the propensity score alone:

$$E\left[Y^1 - Y^0|p(X = x)\right] = E[Y|S = 1, p(X = x)] - E[Y|S = 0, p(X = x)].$$

This reduces the dimension of the estimation problem from the high-dimensional covariate space to a single dimension, the propensity score. Given a small interval $[p, p + \Delta p]$ of propensity score values, we can determine the local average difference in outcomes between the treated and control units with propensity scores values in that interval. Matching is a method of sampling units to generate control and treatment groups which are comparable in terms of their covariate distributions. Matched sampling based on the propensity score alone ensures that the matched treated and controls have the same probability distributions for the observed covariates. This makes matched sampling a useful nonparametric tool for adjusting for treatment vs. control group differences in all observed covariates, thus removing any selection bias due to observables, before comparing the outcomes from treated and controls.

In practice, we seldom know the true propensity score; instead, we have to estimate it. Matching is then based on the estimated propensity score. The intuition