Expert Systems with Applications 41 (2014) 387–396

Contents lists available at ScienceDirect

Expert Systems with Applications

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

A causal inference approach to measure price elasticity in Automobile Insurance

Leo Guelman [a,*], Montserrat Guillén [b]

[a] Royal Bank of Canada, RBC Insurance, 6880 Financial Drive, Mississauga, Ontario L5N 7Y5, Canada
[b] Dept. Econometrics, Riskcenter, University of Barcelona, Diagonal 690, E-08034 Barcelona, Spain

ARTICLE INFO

Keywords:
Causal inference
Price elasticity
Price optimization
Insurance

ABSTRACT

Understanding the precise nature of price sensitivities at the individual policyholder level is extremely valuable for insurers. A rate increase has a direct impact on the premium customers are paying, but there is also the indirect impact as a result of the "causal" effect of the rate change on the customer's decision to renew the policy term. A rate increase may impair its intended impact on the overall profitability of the portfolio if it causes a large number of policyholders to lapse their policy and switch to an alternative insurer. The difficulty in measuring price elasticity from most insurance databases is that historical rate changes are reflective of a risk-based pricing exercise. As a result, the specific rate change at which a customer is exposed is a deterministic function of her observed covariates. The nature of the data is thus observational, rather than experimental. In this context, measuring the causal effect of a rate change on the policyholder's lapse outcome requires special modeling considerations. Conventional modeling approaches aimed to directly fit the lapse outcome as a function of the rate change and background covariates are likely to be inappropriate for the problem at hand. In this paper, we propose a causal inference framework to measure price elasticity in the context of Auto Insurance. One of the strengths of our approach is the transparency about the extent to which the database can support causal effects from rate changes. The model also allows us to more reliably estimate price-elasticity functions at the individual policyholder level. As the causal effect of a rate change varies across individuals, making an accurate rate change choice at the individual subject level is essential. The rate at which each subject is exposed could be optimized on the basis of the individual characteristics, for the purpose of maximizing the overall expected profitability of the portfolio.

1. Introduction

Cost-based pricing of individual risks is a fundamental concept in the actuarial ratemaking literature. The goal of ratemaking methodologies is to estimate the future costs related to the insurance coverage. The loss cost approach defines the price of an insurance policy as the ratio of the estimated costs of all expected future claims against the coverage provided by the policy to the risk exposure, plus expenses (Denuit et al., 2007). There is a wealth of actuarial literature regarding appropriate methodologies for using exposure and claims data in order to calculate indicated rates (Brown and Gottlieb, 2007; Finger, 2006).

A revised set of rates will impact the profitability of an insurance portfolio due to its direct impact on the premiums that policyholders are paying. However, there is also the indirect impact as a result of the policyholders' reaction to the rate change. As basic Auto Insurance is mandatory in many countries, a rate change exceeding a certain threshold will make a policyholder more likely to shop for an alternative insurer and potentially switch to another company. If the rate change causes a large number of customers to lapse their policy, the revised rates could impair its intended impact on the profitability of the insurance portfolio.

In recent years, insurers are switching from a pure cost-based pricing to a demand-based pricing. Price optimization strategies (Towers Perrin, 2007) aim to integrate the cost-based pricing with the customer's willingness to pay into an overall pricing framework. A key component of this framework involves predicting, to a high degree of accuracy, how customers will respond to alternative rate changes, conditional on the customer's characteristics being held fixed.[1]

If we let for a moment the rate change play the role of a treatment with varying 'dose' levels, the main problem involves the

* Corresponding author. Tel.: +1 905 606 1175; fax: +1 905 286 4756.
E-mail addresses: leo.guelman@rbc.com (L. Guelman), mguillen@ub.edu (M. Guillén).

[1] An additional issue is the reaction to new products or cross-selling (see, for instance Kaishev et al., 2012; Thuring et al., 2012).

0957-4174/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.eswa.2013.07.059

selection of optimal treatments for individuals on the basis of estimates of potential outcomes resulting from treatment alternatives. A similar kind of estimation problem is shared across many disciplines, ranging from economics to medicine. In this sense, the price elasticity problem can be conceived under a causal inference framework, which is typically interested in questions of the form "what would happen to a subject had she been exposed to treatment *B* instead of *A*?". The alternative choice *B* is a counterfactual with an associated potential outcome. Thus considerations about potential outcomes from alternative treatment choices seems inescapable from the price elasticity estimation problem.

A randomized controlled experiment is generally the best approach for drawing statistical inferences about effects caused by treatments. The most effective way to measure price elasticity at the portfolio level would be to randomize the allocation of policyholders to various treatment levels and then measure the impact on retention. However, in the most common situation, insurance databases contain historical price changes which are reflective of a risk-based pricing exercise. Under this situation, treatment assignment is a deterministic function of the policyholder's observed risk characteristics. The nature of the data is thus observational rather than experimental, as randomization is not used to assign treatments. In the absence of experimental design, causal inference is more difficult and requires appropriate modeling techniques.

The standard actuarial approach to measure price elasticity in insurance is to model the policyholder's lapse behavior as a function of the rate change and the policyholder's covariates (Anderson et al., 2007; Smith et al., 2000; Yeo et al., 2001). The key assumption is that the inclusion of those covariates will adjust for the potential exposure correlations between price elasticity and other explanatory variables. This approach will be unreliable for estimating causal effects from observational data due to masked extrapolation problems, and the sensitivity of the results to unwarranted assumptions about the form of the extrapolation (Berk, 2004, p. 115; Guo and Fraser, 2009, p. 82; Rubin, 1973, 1979; Morgan and Winship, 2007, p. 129). The problem is even worse when the number of explanatory variables is large, as groups may differ in a multivariate direction and so non-overlap problems are more difficult to detect (Rubin, 1997). Standard statistical software can be remarkably deceptive for this objective because regression diagnostics do not include careful analysis of the distribution of the predictors across treatment groups. When the overlap is too limited, the data cannot support any causal conclusions about the differential effects of treatments (Englund et al., 2008; Guelman et al., 2012; Guillén et al., 2012).

In this article, we propose a method for estimating price elasticity with roots in Rubin's causal model (Rosenbaum and Rubin, 1983, 1984; Rubin and Waterman, 2006). One of the strengths of our approach is the transparency about the data support for estimating the impact of rate changes on customer retention at the portfolio level. The model also allows us to more reliably estimate individual price-elasticity functions. As the causal effect of a rate change varies across individuals, an accurate choice of the treatment at the individual subject level is essential. Each subject's treatment could be optimized on the basis of individual characteristics, and thus maximize the overall positive impact of the rate change intervention.

This article is organized as follows. We first formalize the price elasticity estimation problem from a causal inference perspective. We follow with an overview of the key assumptions required to derive unbiased estimates of average causal effects caused by treatment interventions from observational data. Propensity scores and matching algorithms are discussed next. The second half of the paper presents a detailed application of our approach to price elasticity estimation in the context of Auto Insurance. Managerial implications and a conclusion are outlined at the end.

2. Price elasticity as a causal inference problem

We postulate the problem in the context of Rubin's model of causal inference. This model conceptualizes the causal inference problem in terms of potential outcomes under each treatment, only one of which is observed for each subject. In this paper, we draw on the terminology and framework of experiments, and use the words *treatment* and *rate change* interchangeably. The notation introduced below will be used throughout the paper.

The insurance portfolio is composed of L policyholders, $\ell = \{1, 2, \ldots, L\}$, characterized by a vector of pre-treatment covariates $\mathbf{x}_\ell$. We consider the case of T treatments (representing rate change levels), indexed by $t = \{1, 2, \ldots, T\}$. We let $Z_{\ell t}$ be a set of T binary treatment indicators, such that $Z_{\ell t} = 1$ if subject ℓ received treatment t, and $Z_{\ell t} = 0$ otherwise. We postulate the existence of *potential responses* $r_{\ell t}$ to denote the renewal outcome[2] that would be observed from policyholder ℓ if assigned to treatment t. The *observed response* for subject ℓ is $R_\ell = \sum_{t=1}^{T} Z_{\ell t} r_{\ell t}$.

Our interest lies in estimating price elasticity, defined here as the expected renewal outcomes that result and are caused by the rate change interventions. Here causation is in the sense of *ceteris paribus*, meaning that we hold all policyholder's covariates constant. Our aim is to obtain an estimate of the price-elasticity functions at the policyholder level, $\hat{r}_{\ell t}\ \forall t = \{1, \ldots, T\}$, and in particular in differences of the form $\hat{r}_{\ell j} - \hat{r}_{\ell k}$, the causal effect of exposing subject ℓ to treatment j rather than to treatment k (for any $j \neq k$). We then use these individual estimates to construct an aggregate price-elasticity function at the portfolio level, $\hat{\mu}(t) = (1/L)\sum_{\ell=1}^{L} \hat{r}_{\ell t}$. If the variability of the causal effect $\hat{r}_{\ell j} - \hat{r}_{\ell k}$ is large over L, then the average may not represent the causal effect on a specific policyholder ℓ. The assumption that the effect of t is the same on every subject is known as the *constant treatment effect*, and it is relaxed in this study.

In the context of observational data, policyholders exposed to different rate change levels are not directly comparable. As a result, price-elasticity estimation requires adjustment for differences in the pre-treatment covariates. As discussed above, when the number covariates is large and their distribution varies substantially among the different rate change levels, simple covariance adjustment methods are typically inadequate. In this paper, we propose using propensity scores (Rosenbaum and Rubin, 1983) and matching algorithms (Gu and Rosenbaum, 1993) as a method for removing all biases associated with differences in the pre-treatment variables. Our methodology offers a rigorous analysis of price-elasticity in the context of Auto Insurance based on causal inference foundations. The next section discusses the method in detail.

3. The method

Without loss of generality, in this section we will present the method in a simplified case. We will focus on the binary treatment case, with $t = \{0, 1\}$, and let $Z_\ell = 1$ if subject ℓ received the first treatment (the *treated* subjects), and $Z_\ell = 0$ if received the alternative treatment (the *control* subjects). In the context of this study, multi-valued treatments are handled by analyzing a set of binary treatment dichotomies. That is, given T treatments, we analyze the $T(T-1)/2$ unordered dichotomies.[3]

[2] We denote the renewal outcome equal to 1 if the policyholder lapses (does not renew), and 0 otherwise.

[3] For example, with three treatments ($T = 3$), there are $3 = T(T-1)/2$ unordered treatment dichotomies: $\{(1,2), (1,3), (2,3)\}$.