Accident Analysis and Prevention 97 (2016) 153–161

Contents lists available at ScienceDirect

Accident Analysis and Prevention

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

Heterogeneous treatment effects of speed cameras on road safety

Haojie Li[a,*], Daniel J. Graham[b]

[a] School of Transportation, Southeast University, Nanjing 210096, China
[b] Centre for Transport Studies, Dept. of Civil and Environmental Engineering, Imperial College London, United Kingdom

ARTICLE INFO

Article history:
Received 19 January 2016
Received in revised form 20 July 2016
Accepted 8 September 2016

Keywords:
Speed camera
Heterogeneous treatment effect
Propensity score

ABSTRACT

This paper analyses how the effects of fixed speed cameras on road casualties vary across sites with different characteristics and evaluates the criteria for selecting camera sites. A total of 771 camera sites and 4787 potential control sites are observed for a period of 9 years across England. Site characteristics such as road class, crash history and site length are combined into a single index, referred to as a propensity score. We first estimate the average effect at each camera site using propensity score matching. The effects are then estimated as a function of propensity scores using local polynomial regression. The results show that the reduction in personal injury collisions ranges from 10% to 40% whilst the average effect is 25.9%, indicating that the effects of speed cameras are not uniform across camera sites and are dependent on site characteristics, as measured by propensity scores. We further evaluate the criteria for selecting camera sites in the UK by comparing the effects at camera sites meeting and not meeting the criteria. The results show that camera sites which meet the criteria perform better in reducing casualties, implying the current site selection criteria are rational.

1. Introduction

Speed limit enforcement cameras were first introduced in the UK in 1991 and were extended widely in the last decade. Numerous studies have been conducted to investigate the effect of safety cameras, and results show that the implementation of safety cameras has reduced vehicle speed and casualty numbers near camera sites (e.g. Mountain et al., 2005; Gains et al., 2004, 2005; Li et al., 2013). Despite the wealth of empirical evidence it remains unclear how such effects may vary across sites, referred to as heterogeneity of treatment effect (HTE). The hypothesis is that the variation in treatment effects is related to the differences in site characteristics, specifically the extent to which site characteristics meet treatment assignment criteria. The main objective of this study is to analyse how the site characteristics influence the effects of fixed speed cameras, and identify the locations which have benefited most from treatment.

Although the importance of HTE has been widely recognized in causal analysis, most previous studies on speed cameras usually report an average treatment effect (ATE), which neglects the fact that the effects of speed cameras may differ systematically by site characteristics. This is due in part to the fact that causal approaches for exploring HTE, used routinely in other areas of science such as medicine and epidemiology, have not yet been adopted in road safety studies. Understanding HTE has important implications for policy making. Treatments or trials, such as speed cameras, are usually costly. For example, the annual cost of safety cameras is around £100 million for 2003/04 in the UK (Gains et al., 2005). It is desirable that the treatment is operated in a way that maximises effectiveness with limited resources. By revealing patterns of HTE, policy makers can assign treatments to units most likely to benefit from the treatment, so as to improve the cost-effectiveness of intervention. In this paper we tackle this issue by applying and developing causal approaches for estimating heterogeneous treatment effects of speed cameras on road casualties.

The paper is organized as follows. The literature review is presented in Section 2. The method and data used in the analysis are described in Section 3 and Section 4. The results are presented and discussed in Section 5. The conclusions are given in the final section.

2. Literature review

Several studies have been conducted to analyse the effects of speed enforcement cameras on safety (Christie et al., 2003; Mountain et al., 2004; Cunningham et al., 2008; Shin et al., 2009; Jones et al., 2008). In general, these studies show that the implementation of speed cameras has significantly reduced vehicle speeds and the number of casualties near camera sites. There are two outstanding issues, however, which have yet to be fully

* Corresponding author.
E-mail addresses: h.li@seu.edu.cn, rockytj@126.com (H. Li).

http://dx.doi.org/10.1016/j.aap.2016.09.007
0001-4575/

addressed in the previous evaluations of the effects of speed cameras on road casualties.

The first issue regards the selection of the reference or control group. Most studies to date have used before-and-after methods with control groups (Gains et al., 2004; Christie et al., 2003; Cunningham et al., 2008; Jones et al., 2008). In these studies, a group of similar sites is usually selected as the control group in order to account for the general trend in casualties. However, this method is unable to control for effects of regression to mean (RTM), also known as selection bias, which is a type of bias due to a flaw in the sample selection process. The impact of the RTM is that it can make random variation appear as real change caused by treatments and therefore overestimate the effect of a safety treatment.

A reference or control group is usually required to estimate the counterfactual outcomes of the treatment group. However, treated and untreated units may differ in the absence of any treatment due to confounding characteristics, which affect both potential outcomes and treatment participation. In other words, confounding characteristics of units that are treated may differ in some systematic way from those that are not treated, and those characteristics also may have a bearing on the incidence of selection bias and the severity of its impact. This means that only untreated units with similar confounding characteristics to the treated can be used to approximate the counterfactual outcomes of the treatment group. However, in previous research, not only is there insufficient justification of the selection of control groups, how the treatment and control groups are matched is also unclear.

The propensity score matching (PSM) method is proposed by Rosenbaum and Rubin (1983) for selecting control groups and estimating causal effects. The PSM method has been widely used as a tool of evaluation in econometrics (Heckman et al., 1997; Hirano and Imbens, 2001; Dehejia, 2005; Dehejia and Wahba, 2002; Kurth et al., 2006; Lechner, 2001; Abadie and Imbens, 2004, 2009). Recently, this approach has been introduced and employed in evaluation studies of road safety measures (e.g. Li et al., 2013; Sasidharan and Donnell, 2013). We will discuss PSM in the next section.

The second issue arising from these studies is that only ATE is estimated, however, neglecting the fact that treatment effects can differ across the treated population. The ATE provides useful information, but policy makers also care about effects within specific subpopulations. Since road safety measures are usually costly, it is desirable that treatments are assigned to areas or units which are most likely to benefit from the treatment. A good knowledge of the pattern of treatment effects can help policy makers to make optimal decisions with limited resources. Most previous studies on the effect of speed cameras, however, focus on the average benefit, ignoring the fact that the impact may vary across sites with different characteristics.

Several approaches to estimating HTE based on the propensity scores have been proposed and applied in a few quantitative sociological studies. For example, Xie et al. (2012) discuss a practical approach to studying HTE as a function of treatment propensity under the unconfoundedness assumption. Three methods, one parametric and two non-parametric, are described for analysing interactions between treatment effects and the treatment propensity. They apply the three methods to estimate the effects of college attendance on women's fertility based on the work by Brand and Davis (2011). This study applies the approaches introduced by Xie et al. (2012) to estimate HTE of speed cameras on road casualties.

3. Methods

In this section, we first introduce the propensity score and the conditions under which it can be used to evaluate the effect of interventions. Then two approaches based on the propensity score are discussed for ATE and HTE estimation.

3.1. Propensity score matching

The treatment indicator is defined as $T_i = 1$ if unit i receives the treatment and $T_i = 0$ otherwise. $Y_i(T)$ denotes the potential outcome for unit i, where $i = 1,\ldots,N$ and N denote the total population. For instance, $E[Y(0)|T = 1]$ is the expected value of the outcome Y of treated units when not exposed to the treatment. The treatment effect for unit i can be described as:

$$\delta_i = Y_i^{(1)} - Y_i^{(0)} \text{(Individual Treatment Effect)}$$

The fundamental problem of causal inference is that it is impossible to observe the outcomes of the same unit i in both treatment conditions at the same time (Holland, 1986). In practice, control groups are usually selected from untreated units to construct counterfactual outcomes for treated units. However, since the treatment assignment is usually not random and affected by pre-treatment variables, there can be systematic differences between treated and untreated units, and they can affect the potential outcomes, Y.

The basic idea behind matching is to match each treated unit to untreated units with the same values on observed characteristics, such as a vector of control variables $\boldsymbol{X}$. The matching approach becomes more difficult to implement as the number of observed control variables used increases, however. This obstacle can be overcome by matching on a single index instead of multiple dimensions. The most well-known index is the propensity score, which is the probability that a unit is selected into the treatment group conditional on confounding variables. Conditional on the propensity score, differences in observed outcomes between the two groups can be solely attributed to the intervention impacts.

The validity of this approach rests on two assumptions, conditional independence assumption (CIA) and overlap assumption, which can be described as:

$$(Y(1), Y(0)) \perp T|P(\mathbf{X}), \forall \mathbf{X} \text{(Conditional independence assumption)}$$

$$0 < P(T = 1|X) < 1 \text{(Overlap assumption)}$$

For a full discussion of these assumptions please see Abadie and Imbens (2009). It is important to check the validity of the assumptions before estimating the treatment effects. There are several methods for checking these two assumptions and assessing the matching quality. We will discuss this in detail later.

Because linear probability models produce predictions outside the [0,1] bounds of probability, logit and probit models are usually used to estimate propensity scores. For binary treatment, logit and probit models usually yield similar results, hence the choice between them is not critical. Please refer to the paper by Smith (1997) for further discussion of this point. In this paper, a logit model is used:

$$P(T = 1|X) = \frac{\mathrm{EXP}\left(\alpha + \beta' X\right)}{1 + \mathrm{EXP}\left(\alpha + \beta' X\right)}$$

Where α is the intercept and β' is the vector of regression coefficients. The selection of control variables included in PSM will be discussed in Section 4.

3.2. Inferences on treatment effects

Here we discuss propensity score matching and regression methods for estimating ATE and HTE under the unconfoundedness assumption.