Journal of Public Economics 124 (2015) 1–17

Contents lists available at ScienceDirect

Journal of Public Economics

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

Interpreting pre-trends as anticipation: Impact on estimated treatment effects from tort reform☆

Anup Malani, Julian Reif*

University of Chicago, United States
R\FF, United States
NBER, United States
University of Illinois at Urbana–Champaign, United States
IGPA, United States

ARTICLE INFO

Article history:
Received 17 June 2013
Received in revised form 21 January 2015
Accepted 23 January 2015
Available online 31 January 2015

JEL classification:
C50
K13
J20

Keywords:
Anticipation
Medical malpractice
Endogeneity
Tort reform

ABSTRACT

While conducting empirical work, researchers sometimes observe changes in outcomes before adoption of a new policy. The conventional diagnosis is that treatment is endogenous. This observation is also consistent, however, with anticipation effects that arise naturally out of many theoretical models. This paper illustrates that distinguishing endogeneity from anticipation matters greatly when estimating treatment effects. It provides a framework for comparing different methods for estimating anticipation effects and proposes a new set of instrumental variables to address the problem that subjects' expectations are unobservable. Finally, this paper examines a specific set of tort reforms that was not targeted at physicians but was likely anticipated by them. Interpreting pre-trends as evidence of anticipation increases the estimated effect of these reforms by a factor of two compared to a model that ignores anticipation.

1. Introduction

While conducting empirical work, researchers sometimes observe changes in outcomes before adoption of a new treatment program or policy. Fig. 1 provides an example from the medical malpractice liability context. It shows that equilibrium physician labor supply increased well before states adopted caps on punitive damages, which lower physician liability. The conventional diagnosis that researchers make upon observing such a pattern in the data is that the treatment was endogenous: states adopted these caps in response to the change in supply or for reasons correlated with supply (Angrist and Pischke, 2008, chapter 5).

Observing changes in outcomes prior to treatment is also consistent, however, with anticipation effects. Perhaps individuals began changing their behavior in response to an expectation that they would be treated in the future. Anticipation is a reasonable diagnosis if individuals are forward looking, have access to information on future treatment, and there is a benefit to acting before treatment is adopted.[1]

It is very difficult to rule out endogeneity as an explanation for the pre-trends such as those in Fig. 1, although previous studies have argued that the adoption of punitive damage caps was an exogenous event in the medical malpractice context (Avraham, 2007; Currie and MacLeod, 2008).[2] It may be equally difficult, however, to rule out anticipation as an explanation for pre-trends. For example, we present evidence that these reforms were discussed in newspapers years prior to their

☆ We thank Dan Black, Amitabh Chandra, Tatyana Deryugina, Steve Levitt, Jens Ludwig, Derek Neal, Seth Seabury, Heidi Williams, participants at workshops and conferences at IGPA, Harvard University, New York University, the Searle Center at Northwestern University, and the University of Chicago, and two anonymous referees for helpful comments. Anup Malani thanks the Samuel J. Kersten Faculty Fund at the University of Chicago for funding. Julian Reif thanks the National Science Foundation for financial support.

* Corresponding author.
E-mail address: jreif@illinois.edu (J. Reif).

[1] This paper is motivated by a case where there are differential pre-trends, i.e., different pre-trends in units that are treated and in units that are not. However, it is possible for anticipation to exist even when there are parallel pre-trends, i.e., identical pre-trends in both types of units. Suppose there are two states in which doctors have similar expectations about whether a cap will be adopted at some future date t, but at date t only one of the states actually adopts a cap. Pre-trends will be parallel and only post-treatment outcomes will diverge, yet there are anticipation effects by assumption.

[2] Although we selected our application because treatment is likely to be exogenous, it is possible that treatment is anticipated even when it is endogenous. Our methods may be useful even in that context, though care is required in interpreting results. For example, if it is known that endogeneity and anticipation work in opposite directions, our methods yield a bound.

http://dx.doi.org/10.1016/j.jpubeco.2015.01.001
0047-2727/

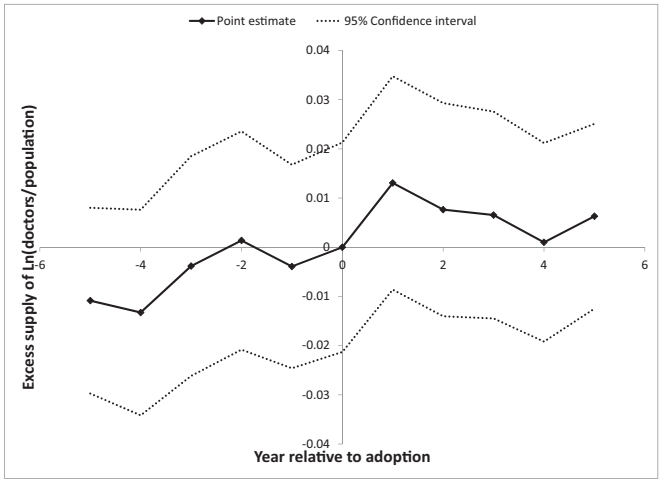

Fig. 1. Excess physician supply before and after punitive damage caps: annual leads and lags from 5 years before to 5 years after adoption. This figure plots the normalized coefficients λ_j from the following regression: $y_{ist} = \sum_{j=-5}^{5} \lambda_j D_{st+j} + \gamma_{is} + \gamma_{it} + u_{ist}$, where y_{ist} is the log of the physician count for specialty i in state s in year t, D_{st+j} is an indicator for whether punitive damage caps is adopted in period $t+j$, and γ_{is} and γ_{it} are the state-specialty and specialty-year fixed effects.

adoption and that there are economic reasons for doctors to change behavior prior to reforms.

Our point is not that the pre-trends in Fig. 1 must be anticipation rather than endogeneity. Rather, we argue that there is no good reason to estimate the treatment effect of punitive damage caps on physician supply assuming that pre-trends can only be evidence of endogeneity. They could also be due to anticipation. This matters because how one interprets these pre-trends has substantial implications for how one estimates treatment effects and how large those estimates are. For example, a researcher who does not account for anticipation effects with the same sign as post-adoption effects of a policy will underestimate the full treatment effect of that policy.

With this objective in mind, we organize the paper around two contributions. Our first contribution is to provide a framework for rigorously comparing and estimating the different models that may be employed to estimate anticipation effects. We start from the premise that there exists a wide array of applied economics topics in which a researcher may be confronted with forward-looking agents whose responses anticipate future treatment. Economic theory suggests, for example, that individuals are forward looking when purchasing durable goods such as cars or houses or making human capital investments, and that firms are forward looking when investing in physical capital or entering new markets.[3]

Two main difficulties arise when estimating models with anticipation effects. One is that researchers may not know how many periods in advance agents anticipate treatment. A common response in the empirical microeconomics literature is to estimate a "quasi-myopic" model that includes anticipation terms for only a finite number of periods.[4] Within these periods, however, anticipation effects are estimated in a non-parametric manner.

An alternative approach, common in the finance and macroeconomics literature, is to posit outcomes as a function of exponentially discounted expectations about future treatment (e.g., Chow, 1989). In this formulation treatment typically has a constant contemporaneous effect and an exponentially discounted anticipation effect. Exponential discounting has the useful feature that suitable differencing can eliminate nearly all anticipation terms.

We do not endorse any particular parameterization. The optimal approach will depend on the theory motivating the empirical analysis and on the limitations of the data. Instead, our framework advances the literature by highlighting the precise assumptions required to generate the regression models estimated in prior literature. It also provides a common benchmark for both the quasi-myopic and exponential discounting models that for the first time allows a comparison of the merits of each.

Another difficulty with estimating a model of anticipation effects is that expectations are generally unobserved. A common response is to examine shocks that alter expectations about treatment but do not actually administer a treatment. An example is a regulation that is enacted at time t but not implemented until time $t+k$ (e.g., Alpert, 2010; Blundell et al., 2010; Gruber and Koszegi, 2001; Lueck and Michael, 2003). Unless actual expectations are observed, however, the investigator can merely demonstrate that expectations affect outcomes. She cannot identify the precise slope of the relationship and thus cannot identify treatment effects that incorporate full anticipation effects.[5]

An alternative approach is to assume a model of belief formation, such as rational or adaptive expectations, in order to substitute observable variables for unobservable expectations of a variable. Unless the forecast error is orthogonal to the observable variables, however, the researcher will have to instrument for them. The traditional source for these instruments is a subset of the agent's information set, for instance, lags of the observable variable (McCallum, 1976). These lags influence the agent's unobservable forecast of a variable but do not directly influence the outcome variable.

A key technical innovation in this paper is our proposal of a novel, alternative set of instruments: leads of the observable outcome or treatment variable. In general, leads can complement lags as instruments for expectations in the forward-looking regression. We show that there are situations in which lags or leads are invalid, though leads are somewhat more robust.

Our second contribution is that we explore the practical implications of the foregoing analysis in an empirical application. Specifically, we estimate the effect of punitive damage caps on equilibrium physician supply and show that accounting for anticipation could increase their estimated effect by a factor of two or more compared to a model that ignores anticipation. We first estimate a model that ignores anticipation and thus corresponds to prior analyses of tort reform, e.g., Klick and Stratmann (2007). We find that caps on punitive damages have a positive treatment effect on physician supply of 1.1% after implementation of caps. Then, we interpret the pre-period trends visible in Fig. 1 as evidence of anticipation effects and estimate the different regression models discussed in our framework. We find that damage caps have a 1.5 to 2.6% post-implementation effect after accounting for all prior anticipation effects. In addition, we estimate that damage caps had a 0.9 to 1.9% effect in each of the two years immediately preceding reform. By contrast, prior models implicitly assume zero treatment effects prior to reform. Our results are robust to different models of anticipation, which suggests that the choice of how to parameterize anticipation

[3] Specific examples include R&D investment decisions (Acemoglu and Linn, 2004), present value asset pricing models (Chow, 1989), pricing of durable goods (Kahn, 1986), real estate pricing (Poterba, 1984), and investment in human capital (Ryoo and Rosen, 2004).

[4] A less than comprehensive list includes: Acemoglu and Linn (2004); Autor et al. (2006); Ayers et al. (2005); Bhattacharya and Vogt (2003); Finkelstein (2004); Gruber and Koszegi (2001); Lueck and Michael (2003) and Mertens and Ravn (2011).

[5] There is also a separate literature on Ashenfelter dips, in which an observed pretrend goes in the opposite direction as the post-implementation effects of treatment (Ashenfelter, 1978). The usual interpretation of such a dip is endogenous selection. A typical solution is to net out the dip by comparing post-implementation outcomes to pre-dip outcomes, in which case the slope of the dip does not matter. However, it is also possible for anticipation to cause opposite-signed pre-trends, in which case netting them out is inappropriate. For example, Lueck and Michael (2003) discuss a case where landowners were found to have killed endangered species on their land in anticipation of a law prohibiting development in areas inhabited by these species. This anticipation effect has the opposite sign of the post-implementation effect of the law, which preserves endangered species.