Social Science Research 52 (2015) 627–641

Contents lists available at ScienceDirect

Social Science Research

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

Curricular policy as a collective effects problem: A distributional approach

Andrew M. Penner [a,*], Thurston Domina [b], Emily K. Penner [c], AnneMarie Conley [b]

[a] Department of Sociology, University of California, Irvine, United States
[b] School of Education, University of California, Irvine, United States
[c] Center for Education Policy Analysis, Stanford University, United States

ARTICLE INFO

Article history:
Received 3 December 2013
Revised 7 March 2015
Accepted 20 March 2015
Available online 30 March 2015

Keywords:
Education
Inequality
Policy
Quantile regression

ABSTRACT

Current educational policies in the United States attempt to boost student achievement and promote equality by intensifying the curriculum and exposing students to more advanced coursework. This paper investigates the relationship between one such effort – California's push to enroll all 8th grade students in Algebra – and the distribution of student achievement. We suggest that this effort is an instance of a "collective effects" problem, where the population-level effects of a policy are different from its effects at the individual level. In such contexts, we argue that it is important to consider broader population effects as well as the difference between "treated" and "untreated" individuals. To do so, we present differences in inverse propensity score weighted distributions investigating how this curricular policy changed the distribution of student achievement. We find that California's attempt to intensify the curriculum did not raise test scores at the bottom of the distribution, but did lower scores at the top of the distribution. These results highlight the efficacy of inverse propensity score weighting approaches for examining distributional differences, and provide a cautionary tale for curricular intensification efforts and other policies with collective effects.

1. Introduction

In the effort to develop an empirical base for social policy-making, scholars often draw upon a medical research model to identify the anticipated effects of different policy interventions. In idealized form this model proceeds in three steps: (1) Based on basic research and observational data, policy-makers or other social actors develop an intervention to address a documented social problem; (2) Evaluators test this intervention on a small scale, typically by comparing outcomes for individuals who are exposed to the intervention ("treated") with those who are not ("control"); (3) Having demonstrated desirable effects in this experimental setting, policy-makers design policies to mandate or facilitate the intervention's adoption at scale. While this design and validation model holds great promise for improving the evidence base of social policy, several scholars have noted that the effects of social policies implemented at scale are often very different from the effects observed for the same interventions in small-scale demonstration projects (Dodge, 2011; Welsh et al., 2010).

In this paper, we consider one such example: Based on evidence indicating that students benefit when they take advanced courses (c.f. Domina, 2014; Heppen et al., 2012; Long et al., 2012), California public schools dramatically expanded 8th grade

* Corresponding author at: University of California, Irvine, Irvine, CA 92697-5100, United States.
E-mail address: andrew.penner@uci.edu (A.M. Penner).

http://dx.doi.org/10.1016/j.ssresearch.2015.03.008

Algebra enrollments between 2005 and 2010. Our analyses, reported here and elsewhere (Domina et al., 2014a,b), indicate this policy effort was counter-productive. We introduce the concept of "collective effects" in an attempt to explain this disconnect. We argue that most evaluation research that informs policy-making focuses on the effects of interventions on individuals. But most social policies affect not just individuals, but also schools, neighborhoods, and societies. Put simply, collective effects arise when the effect of a policy on a given individual diverges from the effects of that policy on the population at large.

A simple illustration encapsulates this insight: Standing up at a baseball game is likely to improve any given spectator's view. However, if every spectator in the stadium stands up at the same time, nobody's view is likely to improve appreciably. In other words, the observation that standing improves views at the individual level is insufficient for estimating the effects on a policy requiring all spectators to stand up. Analogous collective effects exist in many domains. Thus, while we often analyze social policies from a partial equilibrium perspective, holding everything in the model constant while shifting a single parameter, a general equilibrium model is likely to be more appropriate, since policies often lead to large-scale changes in the access to given interventions (cf. Lise et al., 2004). Put differently, while ceteris paribus is a helpful concept for understanding individual effects, when policies are put into place at the population level many things change.

More technically, one can view collective effects as suggesting that for many social policy interventions, the stable unit treatment value assumption (SUTVA) for causal inference is unlikely to be met unless assignment to treatment occurs at the population level (e.g. schools instead of students; communities instead of individuals) so that the effect of the treatment is not affected by whether others were treated. However, we believe it is more helpful to think about individual and population level effects as being fundamentally different questions, and to recognize that it is only under certain conditions that they have the same answer. Since the effects of many social interventions spill-over across individuals, we suggest that estimates of an intervention's effect derived from settings in which a limited number of individuals are treated may be of limited value for understanding the intervention operating at scale.

To return to the baseball analogy, a stadium designer likely cares less about the view from each particular seat than the broader distribution of views. Likewise, when designing social policy, we argue that it makes sense to think about effects on the population broadly. Most analysts would argue in favor of adopting a policy that has desirable effects when implemented at the population level, even if complying with the policy had undesirable effects on an individual who complies with the policy in isolation. For example, in a world in which few drivers comply with traffic regulations, compliance might arguably be dangerous for any given driver. However, near-universal compliance with traffic regulations undoubtedly improves safety for all drivers, including the few who do not comply.

We further argue that in considering the population perspective, it is often helpful to think beyond average differences and consider the broader distribution of outcomes. Once again the baseball stadium analogy is useful. If the people who are most likely to stand when others are sitting are the shortest (i.e., those who have the most to gain by standing), they are likely to lose the most if everyone stands, and their standing view may be substantially worse than if everyone (themselves included) were seated. However, one could imagine that the average view quality is the same regardless of whether people are sitting or standing, even though there is more inequality in views when people are standing. We thus argue that to understand the population level effects of a policy it is helpful to compare outcomes across the distribution, for example, by comparing each of the different percentiles of the relevant distributions.[1]

In this paper, we develop the notion of collective effects as we evaluate the distributional consequences of California's ambitious effort to improve high school mathematics achievement and narrow achievement inequalities by standardizing middle school mathematics curricula. Our analyses indicate that this policy environment is a clear example of an instance in which individual effects and collective effects diverge, both at the average and across the distribution. In the discussion, we build on this insight to provide a preliminary typology of collective effects in educational and social policy settings.

2. Collective effects and curricular intensification

California's effort to universalize 8th grade Algebra culminated in 2008, when the state attempted to require all 8th graders to enroll in Algebra. This push to intensify the mathematics curriculum entailed two major changes for schools. First, and most obviously, it involved exposing more students to relatively advanced Algebra concepts in the 8th grade. Second, the 8th grade Algebra push also precipitated important changes in the skills composition of 8th grade mathematics classrooms, moving low-performing students from pre-Algebra or other less advanced 8th grade math courses to 8th grade Algebra courses that were once reserved exclusively for relatively high-skilled students. In effect, therefore, this policy aimed to detrack mathematics instruction in California middle schools. To understand this change and its potential implications, it is therefore useful to review the literature related to course-taking patterns in secondary school as well as the broader literature on school tracking.

A great deal of research examines the consequences of course enrollment in middle and high school mathematics, where nearly all students are exposed to a subject-specific sequence of course offerings that begins with Algebra I and concludes

[1] We can also think about the distribution of individual effects. In this analogy, this would amount to examining how much each individual's view would change if only that particular spectator stood up, and examining the distribution of these changes (ignoring how this would affect the views of those seated behind the spectator). While distributional approaches can provide information about effects at the individual or population levels (like mean differences), we suggest that thinking about effects at the population level lends itself to thinking about how the broader distribution of outcomes changes.