
I am an applied public economist who engages with frontier theory and econometrics to analyze the welfare effects of government programs and how they can be improved. This research agenda has two main topical groupings: I examine how education policy should be informed by treatment effect heterogeneity and I explore how market imperfections affect government subsidies for alternative energy. My research is motivated jointly by a desire to define policies that make people better off and by a curiosity about the implications of incentives and decision making related to public programs.

Heterogeneity and Education Policy

This line of my research agenda uses new reduced form econometric methods on large administrative data sets to answer questions in education policy, especially K-12 education. Outside of the family, there is nothing more influential in children's lives than school, and I am excited to do research that informs which policies enhance learning and promote long-term growth.

My Job Market Paper, **Worth the Wait? Strategic Kindergarten Entry Improves Achievement But Reduces Equity** asks how parents strategically make kindergarten entry decisions and what implications those behaviors have for achievement and equity. My job market paper answers these questions I use Michigan policies that induce two birthday-based discontinuities to estimate marginal treatment effects among the universe of kindergartners from Michigan public schools. The results show that, counter to some of the received wisdom, families are selecting on gains into when children start kindergarten. This means that allowing strategic entry increases average scores. I also show that the existing entry patterns do increase racial- and class-based achievement gaps. Considering various policy levers by simulation, I conclude that rather than limiting strategic behavior (which would narrow achievement gaps but reduce overall achievement), it would likely be preferable to increase enrollment in public prekindergarten for low-income children (which would narrow achievement gaps but *increase* overall achievement).

How students sort into different educational decisions is also relative in my work on Career Technical Education (CTE) in high schools with Brian Jacob. Our working paper **The Demand for and Effects of Career Technical Education** estimates how CTE programs affect scholastic attainment. We estimate a discrete-choice demand model that allows us to selection correct graduation rates in each CTE program using within-school variation across cohorts in availability and travel time instruments. Preliminary results suggest that participating in professional CTE increases graduation rates—especially in years with less strict compulsory schooling requirements. It also seems to be the case that whereas the students who would benefit the most from participating in vocational programs tend to do so, the students who would benefit the most from participating in professional programs are the least likely to do so—perhaps because of social or other nonpecuniary costs of participating.

There may be other dimensions of heterogeneity besides sorting into investments that social planners care about. In my work in progress *From Value Added to Welfare Added: A So-*

cial Planner Approach to Education Policy and Statistics (with Tanner Eastmond, Nathan Mather, and Julian Betts) we examine the tension between mean-oriented value added measures and policy prioritization of low-achieving students. We propose a new method to estimate teacher value added heterogeneity over the achievement distribution based on endogenous stratification. Preliminary results from the San Diego Unified School District suggest that there is a great deal of both within-teacher and across-teacher heterogeneity and that existing within-school teacher allocations are more likely to give high-achieving kids high-quality matches, which is *not* efficient if the social planner's welfare criterion is concave.

In addition exploring unobserved heterogeneity, there are situations where heterogeneity on observables is the object of interest. For example, in the working paper **Do Perceptions of Public Goods Affect Support? Evidence from Higher Education Appropriations** (with Reuben Hurst and Andrew Simon) we explore the role of information in reducing polarization in the support for spending on public services. We survey 4,700 voting aged respondents about their perceptions of state spending at local public four-year universities and average graduation rates. After a random set of individuals received true information about the graduation rates, we asked all respondents for their perceptions of the quality of the public services and ideal levels of spending. We find that misperceptions about the level of public good provision predict ideal spending and account for some of the polarization across generations and political parties. On average correcting these misperceptions increases ideal spending by 5% and reduces the polarization between these groups—even conditional on priors.

Optimal Subsidy Policy for Energy

The final prong of my research agenda involves more traditional public finance topics—often exploring the compelling market imperfections present in the energy sector. Whereas the first two lines of inquiry have more direct implications for people's well being, I feel excited about my energy-related research for its potential to push policies that make the world a better place by solving externality problems.

For example, despite the theoretical superiority of Pigouvian output subsidies, governments often implement policies on multiple margins: output, investment, entry, etc.—especially in markets for clean energy. In the working paper **Optimal Energy Subsidies with Multiple Policy Instruments: The Case for Subsidies with Deadlines** (with Owen Kay) we explore these tradeoffs. We show that while an output subsidy with a deadline may be optimal when there are administrative costs, it is critical to pair this subsidy with a generous investment subsidy. This is not something present with existing output subsidies such as the ten-year wind power production tax credit. We also show that how responsive production is at the deadline is a sufficient statistic for the optimal time horizon before the deadline. We estimate this responsiveness using a regression discontinuity for wind facilities as they pass through the deadline and find that production decreases despite no changes to total capacity or wind speeds. If wind facilities manage to respond, we argue it is critical that policy makers consider the incentives created by deadlines, and (if permanent subsidies are costly) implement complementary investment subsidies as needed.

Another key concern about optimal subsidies is how subsidies distort other firm incentives, a question I am examining in my work in progress *Dynamic Inefficiencies of Production Subsidies: Technological Progress, Investment, and the Race for Wind Resource*. Firms often face simultaneous trade-offs when making entry and location decisions. For example, in the US wind industry, location decisions face a static tradeoff between wind resource and proximity to transmission, and entry decisions face a dynamic trade-off between early entry in preferred locations and waiting for technological progress. I show that market failures like inter-firm externalities from transmission investment can break Pigouvian subsidies calibrated only to pollution externalities, generating dynamic distortions by inducing entry at inefficient times and in inefficient places. The key parameters for this distortion are the effect of the market imperfection on entry timing and the degree of complementarity between productivity (in this case wind resource) and technology (feasible turbine rotor diameter). I will estimate these parameters and include them in a dynamic model of entry and investment to measure their effects and consider counterfactual policies.

Other Work

I also started two projects in grad school with friends from my time as an undergraduate. Although they are not currently tied into my research agenda, I am actively looking for avenues for arbitrage. In *Can Familiarity Breed Advocacy? Evidence from Volunteer Missionary Service for the Church of Jesus Christ of Latter-day Saints* (with Tanner Eastmond) we use the assignments of (Mormon) missionaries for the Church of Jesus Christ to see how prolonged exposure to minority-rich communities shapes long term racial attitudes and behaviors. Our pilot result suggest a compelling set of effects.

In the second, **Capital Utilization in Production Function Estimation** (with Michael Gmeiner), we extend existing IO production function estimation techniques to allow for variable capital utilization. Because the current draft of the paper relies on shape restrictions to predict utilization, we are applying for Census Research Data Center access to restricted use data with capital utilization surveys. This project will hopefully be the first in a set of RDC-related projects that are more aligned with my research agenda: considering the effect of tax and regulation changes on firms' total factor productivity and the extent to which research and development tax credits may increase market concentration.

Conclusion

I am excited to continue answering welfare-relevant questions at the heart of public economics using advances in applied econometrics (both structural and reduced form) to inform policy and make a difference in the world. While the applications of these interests are varied, my research agenda combines them in a way that makes important contributions, enjoys the intellectual arbitrage, and opens many, many more questions to explore.