A Computer Simulation of School District Economics:

Modeling Allocation Effects of Choice Programs

A DISSERTATION

SUBMITTED TO THE FACULTY OF THE

UNIVERSITY OF MINNESOTA

BY

Peter Carl Kirwin

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

David Johnson, Adviser

FEBRUARY 2015

Table of Contents

List of Tables

List of Figures

CHAPTER 1: Problem Statement

Problem Statement

Evaluators can use cost-effectiveness analysis to help policy makers choose a course for improving student achievement. However, existing cost estimates of one such course, school choice programs, ignore significant non-linear allocation effects created by these policies. These effects are too complex for description by systems of equations; I propose that computer simulation will provide a more accurate and portable method of estimating these costs.

Significance

Dissatisfaction with American public schools arises often in both popular and academic writing. For example, some researchers argue that student achievement is too low based on comparisons with other countries, predicting that the global status enjoyed by the United States will not endure if our students do not keep pace with students in other countries on international tests (Hanushek & Lindseth, 2009). Others argue that the persistent racial achievement gap shows that American schools under-serve their students (Ladson-Billings, 2006).

A serious problem like low student achievement has elicited a great variety of responses from education reformers (see below). Each of these proposals, however, presents some sort of cost, and resource constraints require choosing among them instead of simply implementing a large number of ideas that demonstrate benefit to students. Public education has not recently enjoyed extravagant funding, and the national recession and increased focus on accountability testing have further squeezed education budgets to the point where class sizes are rising and programs are being cut (Caesar & Watanabe, 2011;-2

Beveridge, 2010). Budgets are especially tight for those students our education system should work hardest to help: funding levels and standardized test scores are positively correlated (Ladson-Billings, 2006), so the student populations whose achievement most needs raising also attend schools with relatively lower funding. According to Duncombe (2006), factors like high voter education level, a large proportion of voters over 65 years of age, and low average voter wealth increase pressure on districts, through election of school boards, to keep education spending as low as possible while still achieving satisfactory results. This pressure will prevent districts from spending sufficient resources to enact all beneficial educational interventions; districts must select policies that lead to desired educational outcomes at the lowest cost.

Policy makers responsible for implementing one of the many proposed methods of raising student achievement face a difficult choice. Yeh (2010) identified 22 proposed remedies for low student achievement, and these proposals differ remarkably from each other in terms of their targets and theories of operation. School choice programs—thei set of programs discussed primarily in this paper, including school vouchers, charter schools, and open enrollment—attempt to effect change in school districts that will improve schools through competition (Hoxby, 2003) and ultimately benefit all students. Some proposals, like comprehensive school reform, focus on changing schools directly (Rowan, Camburn, & Barnes, 2004). Others assert that student performance will rise as a result of changing requirements for teachers through, for example, requiring more intensive certification (Darlington-Hammond, Berry, Thoreson, 2001). Yet other proposals, like frequent formative assessment, act on students directly (Yeh, 2007). Because proposals for raising student achievement are so substantially diverse, meaningful and systematic comparison poses a challenge.

Despite the significant heterogeneity of these proposed solutions, limited resources makes choosing among them necessary. For policy makers faced with the task of improving education while keeping costs down, obtaining the biggest “bang for the buck” is essential. Some researchers and evaluators have adopted a formal approach, cost-effectiveness analysis, to find the most economically efficient policy for improving education with the hope of assisting policy makers with the tough choices they face. Economists have long employed this form of analysis to compare programs with different outcomes in search of the most efficient application of resources (Robinson, 1993). This method of analysis has effectively, although not broadly, been employed in education research (Levin & McEwan, 2001). Cost-effectiveness analysis measures all of the programs in terms of a standard metric of effectiveness—often student achievement on standardized test scores in math and reading for these kinds of programs—and the total economic cost of each program. Evaluators can then usefully compare these various programs by examining the cost per unit of effectiveness.

Meaningful cost-effectiveness analysis requires both accurate costs and quantities representing effectiveness for all options under comparison. A significant amount of research over many years has gone into assembling measures of effectiveness and estimates of costs that comparing the many existing proposals will require. Measures of the effectiveness of school choice programs exist in some depth and quantity in the literature, and will not receive extensive comment in this paper. Several school systems including Dayton, Ohio, New York City, and Washington DC (Howell, Wolf, Campbell, & Peterson, 2002) have attempted long-term, large scale voucher programs or charter school programs, and sufficient research has been undertaken on these projects to conclude that both voucher programs and charter schools have a positive overall effect on student achievement for those students using vouchers to choose a different school or attending a charter school (Howell, Wolf, Campbell, & Peterson, 2002; Barrow & Rouse, 2008).

The cost of school choice programs used in cost-effectiveness analysis, on the other hand, remains controversial and the subject of much debate, with statements of cost neutrality—nearly $0 per student—on the low end (Coulson, 1998) and costs as high as $9,646 per student on the high end (Yeh, 2010). If we base a cost-effectiveness analysis on a faulty, artificially low cost, the effectiveness-cost ratio will be unduly high, falsely indicating the desirability of the program and posing the risk of adopting a surprisingly expensive, inefficient policy. If we use an artificially high cost estimate, the program will appear less efficient than it is, and policy makers might pass up a potentially valuable opportunity to increase student achievement at a desirable price. Because determining efficient education policy to remedy the important problem of low student achievement requires meaningful cost-effectiveness analysis, and meaningful cost-effectiveness analysis requires accurate cost estimates, improving cost estimates of school choice programs is a vital step in addressing low student achievement.

Researchers have focused on many facets of the economic impact of school choice programs, and have used many tools to estimate those costs. A few examples of different foci while estimating costs of school choice programs include transportation costs and their associated environmental impact (Wilson, Wilson, & Krizek, 2007), the cost of providing information to parents and adjudicating disputes (Levin & Driver, 1997), and economic impacts on family mobility (Nechyba, 2000). The facet that I intend to focus on involves “allocation effects”: the impact on efficiency of moving students between schools due to school choice programs. The central issue of allocation effects concerns the different unit size between the primary revenue source for schools (i.e., students) and the primary cost (i.e., teachers). Yeh (2007) argued that if a school choice program draws a small number of students from a district, then a small number of students will be drawn from each grade in each school (between 0.05 and 3.27 students per grade level per school in the analyzed programs) and this small decrease in enrollment will not allow for a decrease in overhead costs or faculty costs. Most cost estimates of allocation effects of school choice programs (I will henceforth refer to this category of costs as allocation costs) take an aggregate perspective (Yeh, 2007; Yeh 2010; Heinesen, 2004), but the complexity of the school district as a system indicates that analytic descriptions of aggregate behavior may be intractable. Two characteristics indicate the appropriateness of a component- or agent-based approach (these approaches differ slightly in that agents can make decisions about their behavior, and components cannot). First, the relationship between the number of students removed from a school to participate in a school choice program and the reduction in cost to educate the remaining students is not linear, so analytically extending componential rules to aggregate behavior is likely not possible. Second, the components of the system—in this case the students and teachers—depend on each other to some degree, which further excludes an analytic solution. These non-linear behaviors and interdependent components are difficult to account for in aggregate analysis, but relatively easy to include in an agent-based model. The lack of agent-based cost estimation techniques for allocation costs in the literature indicates that our understanding of the system is less than what it could be, and that evaluators and policy makers are developing policies and programs based on unnecessarily imprecise information.

Purpose of the Study and Research Questions

This study has two purposes. Primarily, I hope to help numerical simulation gain a foothold in the field of cost estimations of education programs. To do this, I will explore a specific example of how cost estimations can be improved with numerical simulation. The specific goal of this research is to give evaluators and policy makers a tool for ex ante estimation of the allocation costs of a proposed school choice program for use in a practical, non-academic setting like a school district considering a policy change. School choice programs vary widely along many dimensions (Yeh, 2007), as do school districts, so policy makers cannot heavily rely on estimated costs in the literature when weighing costs and effects of particular options for policy implementation in a school district. Some existing data-based cost estimation techniques applied to school choice programs use rough assumptions and simple accounting (see Levin and Driver, 1997). This sort of estimation is useful as a way of abstractly describing the type of costs that might be generated by a school choice program, but these methods cannot account for subtle complexity, such as allocation effects, that would inform a particular policy discussion. Elasticity analysis, which relies on regression, offers another approach to estimating the economic impact of school choice programs, but this method often requires many years of quite specific data about enrollment and school organization—see Heinesen (2003) for an example of the quantity and specificity of data required to conduct a useful elasticity analysis on school choice programs (that study uses 12 years of data and very advanced statistical methods). This volume of data may not be available in all cases, and, if the policy shift under consideration is motivated by dramatic change in the school district, comparisons between the early data and the recent data may not be appropriate. In light of these needs, one aim of this study is to create a portable technique for estimating allocation costs for use in ground-level decision making. By ‘portable’ I mean that the technology, data, and expertise required to use it might reasonably be available to evaluators and policy makers considering a proposed school choice program.

Secondarily, this study intends to bring cost estimation in education one step closer to the state of the art proven in other economic fields by developing an agent-based method for estimating allocation costs. The perspectives and technology widely used in estimating allocation costs rely on aggregate behavior extrapolated by theorists from componential rules. This philosophy lags behind that adopted elsewhere in the economics of education and in other economic fields (for example, see Senge and Sterman, 1992). I will discuss the importance of this goal more in the Conceptual Framework section below.

In pursuit of these two goals, I will focus on the following questions:

(1) Can the resources (computing power, level of detail of data and information, etc.) available to practicing evaluators generate useful predictions of the costs of school choice programs?

(2) How sensitive are faculty costs to enrollment changes caused by school choice programs?

(3) How much will a proposed school choice program cost due to inefficiencies

created by allocation effects?

To answer these questions, I will select a school district to model with an agent-based computer simulation. Beginning with actual enrollment data and class size policy, I will construct a series of increasingly complex models ranging from models with no interaction between components (i.e., students) to more complex models that include both positive and negative feedback effects.

For each model, I will employ an approach pioneered in financial stress-testing (Sorge, 2004): I will subject the model of the district to counterfactual stress—in this case, the removal of a number of students for participation in a hypothetical school choice program—and use class size policy from that district to determine how many teaching positions could be eliminated. Many iterations of this process will allow me to create a probability distribution of faculty reductions for each scenario.

Changes in faculty size represents an intermediate step to discussing the cost of a school choice program. I will use the model of faculty size response to small enrollment changes described above and accounting data from the district (e.g., per pupil spending, average faculty salary, etc.) to estimate the financial impact on the students remaining in the district.

Two distinct methods to create and calibrate these models are available. Gilbert and Terna (2000) describe a hierarchy of methods for calibrating numerical simulation and assessing their accuracy. These methods aims to match the behavior of the model and its components as exactly as possible to the observed behavior of real entities, and to verify this similarity through the “quantitative agreement with empirical microstructures, as determined from cross-sectional and longitudinal analysis of the agent population" (p. 66). One method that I will pursue will follow this assessment and rely as much as possible on calibration to original research on actual parent and administrator decision making behavior in the district under study. This sort of specific information might offer a very accurate description of a particular district, but such specificity will make the results esoteric and less useful to evaluators in other districts. A more portable method for constructing this model will rely heavily on calibration to easily obtainable data, findings in the literature, and then choosing at random from a wide range of possible modes of operation, similar to the approach used by Nechyba (2000) in Tiebout modeling of the relationship between school choice programs and housing values; this method will allow evaluators to easily apply a computer simulation to new school districts. I can then conduct sensitivity analyses to determine how much specific information is necessary for evaluators and policy makers to arrive at a useful estimate of allocation costs.

I initially grounded this research as pertaining to the achievement gap, but that statement should be moderated for two reasons. General Creighton Abrams quipped that we should eat elephants one bite at a time. First, this work will not eliminate the achievement gap on its own, but I do hope that the ideas presented here will contribute to the improvement of education in the US. Second, although improving the cost estimates of school choice programs will (at least theoretically) contribute to the quality of the education system, using numerical simulation to tackle nonlinear systems has applications beyond the specific context described above. I hope that the general applicability of these methods will not get lost in the discussion of this specific project.

Key Terms and Concepts

Definition of some the frequently-used terms in this paper will enhance its clarity.

Componential and *aggregate*. *Componential* behavior refers to the rules that govern the behavior of a single unit. Taking an example from chemistry, a molecule of water has certain properties (mass, polarity, bond strength, etc.) that describe its action and interaction with other molecules. Because of these properties, a water molecule will react predictably in known environments. *Aggregate* behavior refers to the rules that govern a group of interacting units. To continue with the example from chemistry, a cup of water has its own set of properties (density, boiling point, etc.). In the case of water, some of the aggregate characteristics, like boiling point, can be thoroughly described in terms of some of the componential characteristics, like intermolecular forces. The difference, though, is that some of the aggregate properties do not have componential analogues: a single molecule of water does not have a boiling point; only a collection of interacting water molecules does.

Analytical and *Numerical.* Some aggregate behavior can be described *analytically*, i.e., an exact solution can be obtained using a manageable set of equations, and some cannot. Leombruni and Richiardi (2005) observed that “many systems are characterized by the fact that their aggregate properties [i.e., aggregate behavior] cannot be deduced simply by looking at how each component behaves [i.e., componential behavior]” (p. 1). The algebra required to solve the componential equations in these systems is either so vast that it is practically unsolvable with finite resources, or logically unsolvable. On the other hand, numerical methods do not make an attempt to deduce aggregate behavior from componential rules; *numerical* methods iteratively apply componential rules across time or agents within a system. Where the theorist converts the componential to the aggregate through brain power in analytical pursuits, in numerical methods a computer often performs the calculations that allow the system to evolve; the theorist then exams the resulting system and comments on evident aggregate behavior.

Conceptual Framework

This study is based on the premise that numerical simulation offers researchers and practitioners a powerful tool for solving mathematical problems, but also a new way to view questions and systems. In other words, numerical simulation offers a way to solve math problems in complex systems, but it also leads to a new vocabulary and a new set of concepts with which to define the world (Ostrom, 1988). In numerical simulation, exploring a system begins not with a search for analytic equations that describe aggregate behavior, but with simple componential rules that interact to describe the aggregate behavior.

The notion that social systems are too chaotic or complex (i.e., impossible to predict future states, sensitive to initial conditions, etc.) to describe with analytical models was borrowed from physics and mathematics beginning in the late 1980s (Gregersen & Sailer, 1993), so an example from physics may help clarify this perspective. The solutions to the two body problem (e.g., a comet orbiting the Sun) and the three body problem (e.g., two moons orbiting the same planet and gravitationally interacting with each other) in physics depict cases that show when analytic solutions work and when they do not. Before computers, physicists had to use increasingly sophisticated mathematics to describe observable behavior that follows simple rules.

Isaac Newton invented whole disciplines of mathematics in order to analytically describe the observable universe (Westfall, 1994). After Newton hypothesized that the only force at work in the heavens was gravity acting with an inverse square law, he used (and largely developed) calculus, then a nascent tool, to show that the orbits of comets and planets around the Sun described by Kepler could be explained solely by gravity. In this approach, the aggregate behavior of the comet can be explained to the extent that the theorist can conceptualize how the componential behavior of gravity and momentum leads to aggregate behavior of comet trajectory. For hundreds of years, the relationship between componential and aggregate proceeded in this fashion: the only way to use mathematics to understand the universe was to develop increasingly sophisticated equations that convert the componential to the aggregate through analytic solutions. But some questions lie beyond the capacity of sophisticated mathematics to find analytics solutions: Newton solved the two body problem analytically but never solved the three body problem, and neither has anyone else.

Newton could not solve the three body problem, but even inexpensive computers today can do so with ease (Appel, 1985). Computers accomplish this not by finding analytic solutions that humans cannot, but by applying the simple rules that govern componential behavior (the inverse square law of gravity and the definition of momentum in the case of the three body problem) many times for very short time intervals. The solutions are numerical, not analytic, in that they do not result in an equation that can be solved with arithmetic and algebra for where bodies will be, but it is possible to calculate the location of the three bodies at any time, and the accuracy of the position is limited only by the computational resources. Numerical solutions do not require that the theorist can mathematically conceptualize the relationship between the componential behavior and aggregate behavior; a computer applies the rules that govern the componential behavior, and the theorist observes and contextualizes the aggregate behavior that results (Leombruni & Richiardi, 2005).

Numerical solutions to the three body problem both allow scientists to determine if and when a particular asteroid will hit Earth, but also open the door to a new realm of questions to ask, questions that would have been beyond the horizon if only analytic solutions were available. Scientists can extend the three body problem to a system in which two galaxies, each made of millions of bodies, collide—all the math is the same (the inverse square law for gravity and the definition of momentum) but the number of bodies and the required computing power are greater by orders of magnitude. An analytical approach to describe colliding galaxies, called the Boltzmann equation, theoretically exists, "[u]nfortunately, we don't know how to obtain relevant analytics solutions to this equation except in the limit of weak interactions … or for rather limiting special systems …, so numerical methods must be used" (Barnes & Hernquist, 1992, p. 707-8). Numerical simulation, while certainly a tool to solve known problems, also focuses on observing aggregate behavior based on repeated application of componential rules, with little concern at the front end for how the former is caused by the latter. In this sense numerical simulation and searching for analytic solutions represent substantially different frameworks for describing aggregate behavior.

Physicists Stan Ulam and John Von Neumann developed the first numerical simulation in 1946 to solve problems involving the penetration of neutron radiation into various materials during radioactive decay (Eckhardt, 1987). Prior to numerical simulation, physicists could describe how one neutron would behave, but encountered insurmountable difficulty in developing equations to describe how many neutrons would behave. An analytic solution eluded them, so they invented a technique to use computers and random numbers to perform componential calculations repeatedly with the hope of understanding the aggregate behavior well enough to proceed with their research. Since that development, diverse fields such as statistics (Briggs, Wonderling, & Mooney, 1997), meteorology (Silver, 2012), epidemiology (Hethcote, 2000), asset risk analysis (Hertz, 1957), earthquake prediction (Bak, Tang, and Weisenfeld, 1988), and financial stress testing (Sorge, 2004) have shifted their focus from searching for analytic descriptions of aggregate behavior to numerical simulations based on componential rules.

Limitations and Boundaries

Numerical simulations iteratively apply simple rules to explore a complex system, and numerical simulation is itself an iterative process. The first models are relatively simplistic, and mostly intend to demonstrate that simulating the system is possible and a viable avenue for research. Once the original model has been created and accepted, future researchers add more sophisticated features over time (Gilbert & Terna, 2000). Because this study will contain the first attempt at modelling allocation effects of school choice programs in an agent-based model, it will expectedly be simple. Future iterations, though, will add more detailed componential rules to fill in and supplant the rough approximations in this model.

Organization of the Dissertation

This dissertation will be divided into five chapters. This chapter has described the problem, its importance, and given an overview of how I approach its solution. Chapter 2 will describe the literature about the cost of school choice programs and about numerical simulation in the economics of education. Chapter 3 will describe the selection of a school district to model and the construction of the series of computer models used to estimate the cost of a hypothetical school choice program. Chapter 4 will describe the output of the computer model constructed in chapter 3 and the meaning of these outputs as results of the study. Chapter 5 will discuss the cost of the hypothetical school choice program and the applicability of this technique for future use in ascertaining the cost of education programs for policy makers.

CHAPTER 2: Literature Review

Introduction

In this section I will provide an overview of the literature that pertains to the study that I propose. The first section describes the argument that underlies the rationale for school choice programs. In this first portion I will describe how the efficient-market hypothesis spawned the idea that schools need to operate in a competitive market in order to function efficiently. I will then show how this political philosophy became confused for a complete statement of the economic impact of school choice programs. The discussion of the broad body of literature will conclude with a more nuanced economic analysis of school choice programs that accounts for the impact of allocation effects on the economies of scale that schools rely on. The literature chosen for this section is not meant to exhaustively cover all of the studies that pertain to the economics of school choice programs, but rather to provide a representative sample rich enough to explain the argument that surrounds school choice programs.

I will also discuss the body of literature that lies nearest my study: the use of computer simulation in the analysis of education programs, with special focus on simulations that estimate economic impacts or school choice programs. The first area of literature discussed here, using simulations to explore Tiebout equilibriums under different voucher programs, offers a particularly good example of how simulations can assist in economic analysis and how simulations iteratively improve over time. But even more so, this literature shows how simulations shift the theorist's focus from trying to describe aggregate behavior to describing componential behavior and letting a computer simulate the aggregate behavior. This discussion will conclude with some examples of simulations used for other analyses in education that I have chosen because they demonstrate some important elements of approaching simulation and constructing a simulation well.

Economic Theory of School Choice Programs

The literature that explores the costs of school choice programs falls in two camps. One area of the literature focuses on the funding sources for schools. This body of literature exists more in the public policy arena than in peer-reviewed journals, but still represents an important perspective in the American education debate. The other area of literature focuses on producing cost estimates for use in cost-effectiveness and cost-feasibility analysis. This discussion represents the bulk of the academic literature, but is often too complicated to receive thorough treatment in the public policy arena.

Market Pressure in Education

To understand the argument about the cost of school choice programs, one must first understand the mechanism assumed to underlie their advantage. Proponents suppose school choice will improve the education system by subjecting all schools to more competitive pressure, which they contend will incent schools to use resources more efficiently and thereby improve educational outcomes. This competitive pressure comes from letting parents choose which institution receives the government money allotted for their children, a mechanism commonly referred to as “the money follows the child.” Some advocates erroneously assume that the mechanism of school choice completely describes the cost of the programs. This section will discuss the mechanism of school choice; an overview of attempts to describe the real cost of school choice programs will follow.

Early Rationales for School Choice

References to education systems that allow parents to direct government money for the education of their children can be traced back as far as Thomas Paine (Fishman, 1983). Paine's (1791) rationale for what would become known as school choice in "The Rights of Man" was principally egalitarian; he argued that quality education was out of reach for students in low economic classes, and a negative tax specifically created for the purchase of education services could help low income families increase their access to education. Educating poor children would benefit society in the long run: "the number of poor will hereafter become less, because their abilities, by the aid of education, will be greater" (Paine, 1791, chapter 5, part 5). The advent of free, compulsory education in the United States renders Paine's egalitarian argument largely moot, although access to education remains far from equal to students in different socioeconomic backgrounds (Biddle & Berliner, 2002).

The idea of school choice was reintroduced by the economist Milton Friedman as an abstract statement of economic theory and philosophy, more than education policy (Fox, 2011). Friedman (1955) proposed two reasons for enacting school choice. First, he believed market forces would allow the private sector to find ways to spend educational resources more efficiently than the government-administered public school system. Friedman's second reason for promoting school choice lies more in the realm of political philosophy than economic policy: the government takes on three roles in providing public education, only two of which are appropriate. He asserts that the government should (a) set standards and goals for education and (b) fund its provision. Friedman claims, however, that capitalist economic theory suggests that the third role, (c) actually administering the education system, should fall to the private sector instead of the government.

Creating a Marketplace of Education

Friedman viewed the benefit of market forces abstractly in terms of "more efficient utilization of their resources" (Friedman, 1955, p. 9) but offered no specific vision of how increased efficiency would benefit students. The proposition that an education system will respond positively to participating in a competitive market requires the explication of two underlying premises. First, monopolies allow inefficiencies because they have no incentive to remove them; existing inefficiencies allow the system to produce smaller outputs than could be produced by the inputs in a perfect system. Liebenstein (1966) first labeled this diminished output as "X-inefficiency" and defined it as the increase in price compared to outputs that result when firms operate in the absence of "competition and adversity [that] create some pressure for change" (p. 408). Second, public schools maintain a partial monopoly over consumers of education. Few communities have school choice programs of sufficient size to threaten public school systems with significant competitive pressure (Hoxby, 2003). Even in the absence of school choice programs, however, public schools do not hold a true monopoly in education: the presence of private schools allows some parents to exercise some choice. Additionally, the ability of families to relocate to a different neighborhood allows consumers to choose education other than that provided by their neighborhood school (Dougherty et al., 2009). Despite these options, most consumers use the "free" education provided by their neighborhood public school, and consequently these neighborhood schools do not compete in a truly open market to a great extent.

The literature offers direct evidence that the public education system operates below maximum efficiency (Hoxby, 2003; Eppel & Romano, 2008). Partially inefficient operation of public schools is so well accepted that it has become an accepted premise in other areas of education research as well. A recent debate has raged about the utility of cost functions in determining funding levels for education (see Costrell, Hanushek, & Loeb (2008) and Duncombe (2006) for examples of both sides of this argument). Despite differing perspectives, the two camps agree on the appropriateness of including a variable in the cost functions to approximate how efficiently schools use their resources—indicating wide agreement that public schools do not operate at maximum efficiency. The efficiency variable is included to account for the fact that public schools do not participate in a "competitive market," but rather in "the very different environment of public education" (Costrell, Hanushek, & Loeb, 2008, p. 207).

Effects of Market Pressure on Schools

Theoretical descriptions of the possible mechanisms for using choice to improve educational outcomes in public schools began to arise as early as the 1970s. Jencks (1970) posited that an education system can maintain teacher quality either through bureaucratic means such as training and academic requirements, or through the market by allowing well-informed parents to choose schools with better teachers. Schools providing services valued by the market will maintain enrollment levels because parents choose to send students to them; the schools not providing sufficient educational value to attract enough students will decline and the market will ultimately cull them. Jencks goes on to argue that education is nuanced, in contrast with heavy-handed bureaucratic requirements, so the market approach will more responsively ensure teacher quality. Modern analyses of the effect of competition due to school choice on the quality of public schools is mixed. Some research shows that the presence of a sufficiently large school choice program does produce measurable increases in the educational output of the public schools in the area (Hoxby, 2003), while Barrow & Rouse (2008) found no evidence of improvement in public schools as a result of voucher programs in Washington, D.C., Dayton, Ohio, and New York City. Such discussion of whether market pressure actually improves education is interesting and important, but falls outside the scope of this paper.

What the School Choice Mechanism Is and Isn't

The economic rationale for school choice described by economists primarily concerns reducing X-inefficiency by inducing the education system to function more like a firm in an open market. Some school choice advocates, however, have taken the idea that "the money follows the child" to mean that school choice programs are cost neutral. Consider the following quote taken from the Frequently Asked Questions section of a school choice advocacy group:

> The financial effects of vouchers, even vouchers for the full per-pupil expenditure of public schools, should be negligible. For every decrease in the amount of funds directed at public schools, there would be a commensurate reduction in the work load and hence costs of operating public schools. (Coulson, 1998)

This statement specifically references vouchers, but the underlying idea would apply equally to any school choice program. The school choice advocates who indicate that letting funding follow the child is a cost-neutral plan, as compared to a spending-neutral plan, mistake accounting costs for economic costs. Levin and McEwan (2001) explained that "the assumption that [budgets] will contain all the cost information that is needed is usually erroneous." (p. 45) Budgets serve useful functions, but they do not, and are not intended to, offer a complete picture of a program's costs.

Removing a Small Number of Students

School choice programs, at least as they have been proposed and implemented at present, often remove a small number of students from a relatively large number of classrooms. One impact of these so called "allocation effects" (Hoxby, 2003) is a disruption of an economy of scale that allows schools to operate efficiently. Other researchers have examined different aspects of estimating the costs of school choice programs. The final analysis of the cost of school choice programs will have to synthesize all of these findings, but that is not the task undertaken here. This paper aims to improve estimates of the costs associated with the impact of allocation effects on capacity utilization.

According to Goldhaber's (1999) analysis of proposed universal voucher programs, because some vouchers would go to parents who already choose to pay for private school tuition, "most of the benefits would go to those who are already in private schools. To offset this, vouchers could be made progressive by targeting the voucher to low-income families" (p. 23). Recent voucher programs seem to follow this course: all of the large-scale voucher programs implemented or ongoing in 2007 have been made available to only some students in the district. The selection process was based on socioeconomic targeting, special education status, quality of the student's original public school, lottery, or a combination of those factors (Barrow & Rouse, 2008).

Extant school choice programs impact only a small fraction of students. Large charter school programs have been in place for more than a decade, but even charter schools educate a small portion of this nation's students. Michigan's charter school program is regarded as one of the best and most expansive in the nation (Chakrabarti & Roy, 2011), yet only about 7% of Michigan's school-age children attend these schools (Michigan Department of Education, 2010). Minnesota, which also has a prominent charter school system, enrolls less than 5% of its students in charter schools (Minnesota Department of Education, 2011). These figures exceed the national average for charter school enrollment of 3.3% (National Center for Education Statistics, 2011).

Some question what types of students comprise the small numbers participating in school choice programs. While evidence that higher aptitude students self-select into school choice programs ("cream skimming") has not been found, as was at one time predicted (for example Jencks, 1970), studies in several fields have found that low-cost students—those students who require relatively few resources from their school—self-select into private schools and charter schools in school choice programs, likely because private and charter schools have little incentive to attract high-cost students (West, Ingram, & Hind, 2006; Biglaiser & Ma, 2003).

Inelasticity of Educational Costs for Small Enrollment Changes

Like predictions of cream skimming, claims of cost neutrality have also attracted scrutiny in the literature. Revenue will decrease in proportion to the decrease in enrollment, assuming vouchers equal to the full per pupil expenditure, but costs will not; the school will operate less efficiently and have fewer educational resources for the students who remain in the school. Scafidi (2012), a school choice advocate constructing the argument for later refutation, concisely described the position as follows:

> The claim is that when a child leaves [his or her former public school] via school choice that the public school retains significant fixed costs. A decrease in students means that there is less money to spend on these large fixed costs of operating a school. So, if students leave and these costs are truly fixed and must be paid in order for the school to operate, then the students who remain in public schools will have fewer resources devoted to their education. (p. 3)

The argument made by Yeh and acknowledged by Scafidi relies on the idea that some costs of operating a school are fixed, in that they do not depend on enrollment, and some costs vary directly with enrollment. A school might have a budget of *D* dollars that covers operating costs and an enrollment of *n* students; some school choice advocates (e.g., Coulson, 1998) apparently assume that each student costs *D/n* dollars to educate and removing a single student from the school would decrease both revenue and operating cost by *D/n* dollars. Assuming that all students require equivalent educational resources, the presence of fixed costs in operating a school requires that the marginal cost of educating one additional student is less than *D/n* dollars, so revenue would decrease by *D/n* dollars and operating costs would also decrease, but by a smaller amount.

The presence of fixed costs in operating a school has been widely, though perhaps not universally, accepted for decades, going back at least to the 1960s (Riew, 1966), and educational research in many areas relies on this assumption. In his seminal work on the optimal size of a high school, Kenny (1982) described “the level of ‘effective’ schooling input per pupil (v)” as

$$v = k (s^{\pi - 1}) \text{ (p. 3)} \qquad (2.1)$$

where k is the instructional input per student, and s is the number of students in the school, and π is the economy of scale coefficient. Kenny goes on to measure π empirically and finds that it takes on a value between 1.1 and 1.2. Any value of π greater than 1.0 indicates that larger high schools operate more efficiently. Recent updates to Kenny’s research by Ferris and West (2004) confirmed that an economy of scale exists in terms of educational productivity, but may be offset by increases in school violence with school size. This behavior agrees with the model of school finance that includes fixed costs: continuing the example from above, the marginal cost of an additional student is less than D/n dollars, but that student brings in D/n dollars in additional revenue. Looking at this example through the lens of Kenny’s findings, that surplus money can then go toward purchasing additional educational benefit for the students in the school which would increase the level of effective schooling as school size increases.

Even many school choice advocates acknowledge that schools have fixed costs. Consider the following quote from a policy document written for two voucher advocacy groups:

> School districts could be financially harmed if the dollar amount of vouchers exceeds the variable cost of educating each student. If this occurs, each departing school choice student will take some of the funds

> used to educated [*sic*] remaining students in the district. (Gottlob, 2004, p. 3)

The suggestion that vouchers would be set at or below the variable cost of educating a single student is a limited solution to the problems caused by allocation effects of school choice programs. This mechanism would not be available for other types of school choice programs like charter schools and open enrollment. Moreover, one might question if schools will still feel the competitive pressure that is assumed to result from school choice if the budgetary impact is designed not to financially harm schools as students transfer away as part of a school choice program.

Allocation effects of school choice programs diminish the economies of scale that exist in schools and create inefficient operation. These inefficiencies lead to costs in the form of reduced value provided to the students by the schools; but who bears these costs? Yeh (2007) described the result as follows: "[T]he reduction in revenue associated with declining enrollment is typically offset by reducing other services—art, music, extracurricular programs and other support services." Studies of how schools and districts respond to shrinking budgets confirm that districts will both cut teaching and aide positions, and they will eliminate programs such as fine arts and counseling (Jarman & Boyland, 2011). Thus the recipients of these services lose value otherwise provided by the school. The school's budget must remain balanced, and thus the value of these services lost by the students will not appear on accounting documents. Using Levin and McEwan's (2001) ingredients method of determining the cost of a program, these losses of value count as costs to participants, and so represent a real cost of the school choice program.

Sensitivity of Cost Categories

The costs of operating a school fall into three categories: fixed costs, variable costs, and step costs. Some costs of running a school are fixed with respect to enrollment in that they do not depend on the number of students at the school. Some costs vary with respect to enrollment in that they depend directly on the number of students at the school. Scafidi (2012) introduced a third category of costs into the description of school budgets from an accounting textbook by Horngren et al. (2009) referred to as ‘step costs’, which

> remain the same over various ranges of the level of activity, but the cost increases by discrete amounts—that is, increases in steps—as the level of the activity increases from one range to the next. (Horngren et al., 2009, p. 353 cited in Scafidi, 2012, p. 8)

At least one substantial cost in operating a school, faculty salaries, might behave like a step cost. Yeh (2007) argued that allocation effects caused by school choice programs “would not [allow] any of the public school building principals to eliminate teaching positions in response to the transfers” (p. 426). This statement assumes that the numbers of students that transfer are not large enough to allow faculty spending to move from its current level to a lower step. For example, suppose 80 students are taught in four classrooms by teachers with equal salaries, and a school choice program begins removing one student at a time. Because we will assume that 79 students will also require four classrooms and four teachers, removing one student will not allow the number of teachers to decrease. The total cost of the teachers’ salaries will remain constant as the students transfer out until enough students have left for the classes to be consolidated from four to three under the school's class size guidelines. At that point the faculty spending will be reduced to 75% of its initial value. The cost of faculty for this grade level thus exhibits step-cost behavior in that it remains fixed for some range of enrollment and, when the enrollment changes to a value outside of that range, the cost makes a large step to a new value that will remain fixed for a new range. See Figure 2.1 for a pictorial representation of this idea.

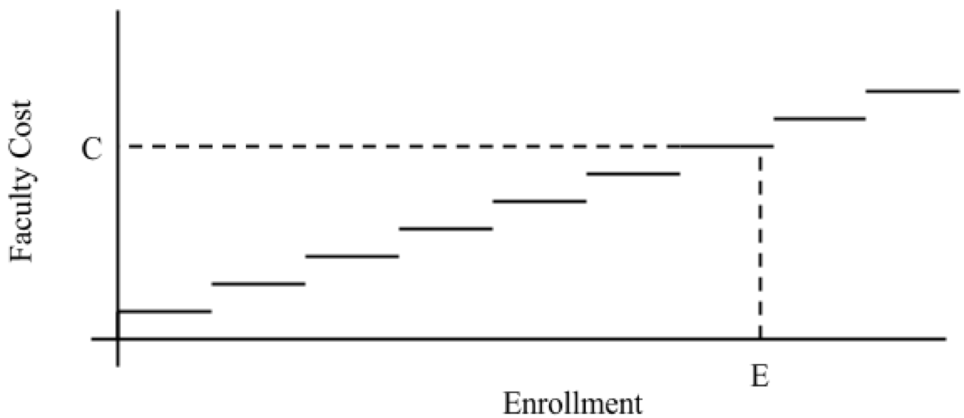

Figure 2.1. A demonstration of faculty cost as a function of enrollment when the faculty cost is conceptualized as a step cost. Note that several values of E will lead to the same value of C.

In addition to the intuitive description of faculty spending as step costs, empirical evidence in support of this comes from Heinesen's (2004) study of the relationship between enrollment and spending in Denmark. Heinesen's work confirmed that spending is partially inelastic with respect to enrollment, specifically "if the number of pupils falls by 1% from one year to the next, expenditure per pupil increases by 0.8%" (p. 443). In other words, a 1% decline in enrollment only leads to a 0.2% decline in expenditure. This study also finds that spending is less elastic with respect to enrollment decreases than enrollment increases due to "capacity utilization effects" (p.

443). (Classrooms and teachers make up capacity in a school; efficiently utilizing resources in education means keeping staffed classrooms as full as possible.) Heinesen's observation that enrollment increases lead to changes in spending more often than enrollment decreases can be understood in terms of step costs and capacity utilization: Because school districts use capacity efficiently, classrooms tend to be relatively full, and enrollment is thus closer to the high-enrollment end of the step; see Figure 2. Adding just a few more students will push the enrollment out of the step's range. To the extent that enrollment lies near the high end of the step, it will lie farther from the low end of the step, so more students would have to leave before enrollment would fall outside of the range of the current step and inside the range of the lower step.

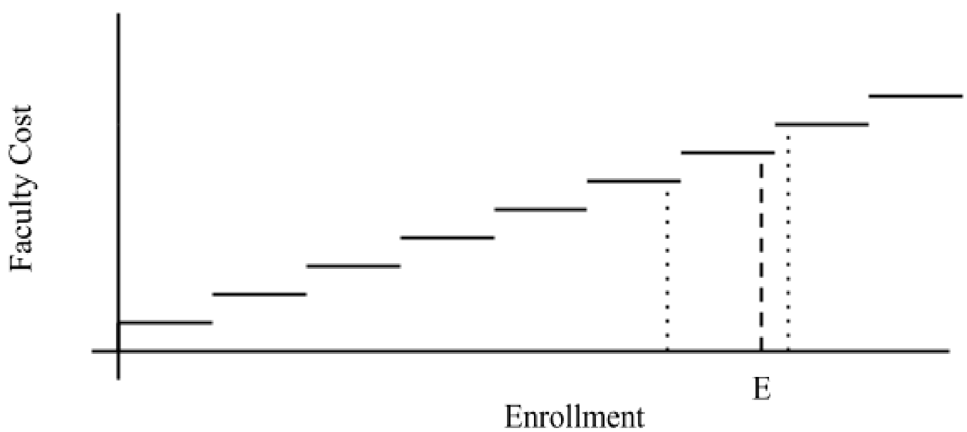

Figure 2.2 Similar to Figure 2.1, but also showing the smallest enrollment increase that would lead to a faculty cost increase, and the smallest enrollment decrease that would lead to a faculty cost savings. Notice that the enrollment (E) lies closer to the high-enrollment end of the step than to the low-enrollment end, this indicates that the school described with this figure uses its capacity (i.e., teachers and classrooms) efficiently. Because of this, the change in enrollment required to move to the higher faculty cost step is smaller than that required to move to the next lower step. Note that any decrease in enrollment that will not allow faculty spending to move to the next lower step will reduce the efficiency of capacity utilization and increase the per-student cost of operating the school.

The discussion so far has tacitly focused on the immediate cost of allocation effects, and has not addressed long term impacts. Little research has studied how long it takes a district to regain efficiency lost due to the allocation effects of school choice programs, or even to regain efficiency lost due to enrollment declines more generally. Cibulka (1983) analyzed ten urban districts that experienced enrollment decline during the late 1970s and found that efficiency is often not entirely regained, and to the extent that the system does find a new, efficient equilibrium it does so slowly and with considerable time lag. Unfortunately, Cibulka's findings offer no direct answers that apply to this discussion, but they do indicate that the answer will not be simple. The step-cost-based perspective described in this section could theoretically apply to longer time scales (e.g., by iteratively applying the model using the output from year $t - 1$ as the input for year t), but longitudinal analysis of allocation effects must necessarily begin with a valid model of immediate effects. Because a model of immediate effects must precede a model of effects over time, this discussion will continue with a focus on immediate effects.

Using Computer Simulation to Estimate Costs in Education

The advent of computers offered an economical way for researchers to approach complicated systems and diminished the need to consider mathematical feasibility when deciding what features to include in the model. This tradition of inquiry has made limited headway into the economics of education. I will discuss three categories of computer simulations used in modelling education programs: numerical solutions to complicated analytical equations, numerical solutions to equations developed with simulations in mind, and agent-based simulations. These three categories of simulation exist in two sub-bodies of literature that reside very near the study that I propose. The first set of articles represent the only literature I could find that uses simulations to explore vouchers. The second body of literature consists of studies that use simulation in other areas of education. I choose to include these as important for a variety of reasons (e.g., exemplars of a particular reason for using simulation, the use of cost-benefit analysis, etc).

Simulations that estimate the impact of vouchers

Epple and Romano (1998) developed an analytic general equilibrium model to show the impact of private schools in a generalized school district on public school quality and housing prices. Their model begins with an assumption that households have some amount of financial resources, and a child of some academic ability. Households weigh the utility they would obtain from each school in a district against how much they would have to pay—either in housing and taxes for public schools or in tuition for private schools. Households then choose their school so that they maximize their utility. As households relocate, though, housing prices and tax rates change, as do school qualities. An equilibrium—technically known as a Tiebout equilibrium (Tiebout, 1956)—exists when no household could find more utility by relocating or changing schools.

The analytic statement of the equilibrium developed in this work relies on a large number of assumptions that make the math possible. For example, the relationship between utility, U, that a household will find at a school, the household's after-tax income, y_t, and the school's average quality of students (θ), is assumed to satisfy the following relationship:

$$\frac{\partial}{\partial y_t}\frac{\frac{\partial U}{\partial \theta}}{\frac{\partial U}{\partial y_t}} > 0 \quad \text{(p. 37)} \qquad (2.2)$$

Simply stated, this condition indicates that the utility that a family obtains in a school is more sensitive to the quality of its student's peers for higher-income families. Or using Epple and Romano's words, "for students of the same ability, any indifference curve in the (θ, *p*)-plane of a higher-income household cuts any indifference curve of a lower-income household from below" (Epple & Romano, 1998, p. 37). The *p* in the previous explanation indicates an amount of private school tuition. Statements like this represent an attempt by the theorist to deduce aggregate behavior from componential rules by making powerful simplifications.

The analytic section of Epple and Romano's (1998) paper ends with the conclusion that a Tiebout equilibrium will exist in districts under reasonable assumptions and no vouchers. They then describe some general, provable conclusions about these equilibria in the absence of vouchers. For example, "the public sector [will have] the lowest-ability peer group" (Epple and Romano, 1998, pp. 40-41). Analytical proof of equilibria in the presence of vouchers cannot be deduced, however. The equilibrium equations cannot be solved analytically for voucher conditions, so the authors turn to a "computational model . . . to examine vouchers, and to explore issues for which comparative-static analysis may yield ambiguous results" (Epple and Roman, 1998, p. 45). The simulation described is simple compared to Nechyba (1999) and Nechyba (2000). The model ignores housing values in different neighborhoods, and instead treats all of the public schools in a district as one large school that contains all students not in a private school. Each household determines the utility of each school using the equation

$$U = (y_t - p)\theta^{0.06}b^{0.3} \text{ (p. 46)} \quad (2.3)$$

where U, y_t, p, and θ are as described above, and b is the ability of the child in the household. The cost function for schools is set by

$$C(k) = F + V(k) = 12 + 1{,}300k + 13{,}333k^2 \text{ (p. 46)} \quad (2.4)$$

where C is the cost of running a school and k is the percentage of students in the district who attend the school. Private schools are assumed to act as non-profit entities, so the cost of running the school is divided among the students attending that school and charged as tuition. A large number of households are created with values for y_t and b chosen from a distribution, with b for each household depending somewhat on y_t, indicating that household income and student ability are correlated. In each iteration, each household determines the utility it would derive from each existing school, and from a hypothetical new private school. If any household would derive more utility by enrolling in a different school, they do so. Equilibrium is reached when no household can increase its utility by transferring schools, and no new private school would attract enough students to survive if created.

The authors conduct a sensitivity analysis because “scant empirical evidence exists on some important parameters of the model” (Epple and Romano, 1998, p. 45). In the case of the correlation between household income and student ability, ρ, little direct evidence exists. The authors took some tangentially-related evidence and set the high end of ρ at 0.4, based on the argument that no evidence can be found to suggest that ρ could be any higher. They conducted a sensitivity analysis and found that “[i]ncreasing ρ from 0 to 0.4 reduces public-sector attendance [from 90%] to 88%” and they continue on to use $\rho = 0$ for the remainder of the analysis because the specific value of ρ has relatively little impact.

Because the simulation developed in this study is relatively juvenile, Epple and Romano candidly point out the aspects that would most benefit from more sophisticated development. The only factor considered in this analysis that captures the value of education is increased lifetime earnings for the students who attend better schools. The authors point out “longer-run externalities from education not considered by private schools, like reduced crime, may be present” (Epple & Romano, 1998, p. 44). Another area for improvement in this study concerns the impact of teachers. The authors point out that the model could “let teachers vary in skill with student achievement increasing in both the school’s teacher-student ratio and the mean skill of the teachers” and “[a]llowing a teachers utility to depend also on the skills of colleagues” (Epple & Romano, 1998, p. 55). The final area for acknowledged shortcoming of this simulation pertains to the public schools in the model. “The public sector is passive in this model for simplicity. Public-sector schools do not segment students by ability (track), increase educational inputs to compete more effectively with the private sector, or behave strategically in any way” (Epple & Romano, 1998, p. 38). In reality, though, public schools do track students as a way of catering to higher ability individuals (Epple, Figlio, & Romano, 2004), and this practice would influence the perceived peer quality in public schools.

Nechyba has published several studies of the impact of vouchers on Tiebout equilibria based on numerical simulation (for example, Nechyba, 1999; Nechyba, 2000). I have selected these two examples for their thoroughness and representativeness, and because they demonstrate a shift in the focus of the study by the author. Nechyba (1999) follows a pattern similar to Epple and Romano (1998): a thorough and detailed formal analysis of analytical equations seeking the aggregate behavior of schools in a district followed by a smaller section devoted to a numerical simulation. Unlike Epple and Romano (1998), Nechyba (1999) places households into different communities. Each community contains its own schools of some quality, its own tax rate, and its own housing prices. Agents in the simulation can choose a community and either enroll their child in the public school or pay tuition at a private school in such a way that their utility is maximized. Utility is defined in a way much like Epple and Romano (1998), except that now the households must also consider the quality of the house they live in.

Like Epple and Romano (1998), Nechyba (1999) turns to a numerical simulation because of "the complexity of the model described. Without making some specific assumptions … there seems little chance of deriving predictable implications." (p. 17) The rationale for Nechyba's use of a simulation in this article indicates a theoretical framework different than mine. "We therefore think of the work that follows as an extension of theoretical work with more specific functional form assumptions" (p. 17). Instead of using simulation to free the theorist from the boundaries of the analytical work as described by Ostrom (1988), Nechyba, like Epple and Romano (1998), uses it merely to solve intractable math problems at the end of an analytical inquiry.

The process of constructing the simulation used by Nechba (1999) begins with creating three communities and a set of house types that populate those communities. Households then populate the houses with randomly assigned incomes and child abilities; before any private schools or voucher programs are added, the households are allowed to iteratively relocate until the system reaches an equilibrium. Once the system reaches this first stage of equilibrium, tax policy is set in each community based on the desires of the median voter, and the system again relaxes into equilibrium. Once in this second equilibrium, the rules change by allowing private schools to exist and possibly introducing a voucher system, and the system evolves until it reaches a third equilibrium. Like Epple and Romano (1998), Nechyba (1999) reports few quantitative outcomes from these simulations. Instead, the results focus on qualitative descriptions, i.e., that "those choosing private schools in poor communities are not original . . . residents of those communities but rather are high income migrants" (p. 30). This process is then repeated with different funding laws—local financing, state funding, state foundation grants, and district power equalization. The analysis concludes with the outcomes of the third-level equilibrium under these different funding schemes; put briefly, the overall public school funding mechanism dramatically impacts the equity of voucher programs. District power equalization leads to the most equitable distribution of benefit from vouchers, and foundation grants lead to the greatest inequity.

In terms of one of the great themes of this study—that developing simulations is itself an iterative process—Nechyba (1999) shows substantial growth compared to Epple and Romano (1998). Much of this growth can be attributed to the function of real estate in the models. By not specifying neighborhoods and housing values, and by having only homogeneous public schools and a single tax rate, Epple and Romano tacitly assume that all households live in the same community. They reach some important conclusions about the manner of equilibrium reached with different kinds of vouchers by giving agents two choices: public or private. Nechyba (1999) includes neighborhoods with varying quality houses and tax laws, so that agents in the model have two more choices: public, private, relocate and choose a public school, and relocate for a private school. Nechyba (1999) does not invalidate the conclusions drawn by Epple and Romano (1998). Instead, the improvements made in this model enrich and expand upon the work that went before it.

The final equilibrium model that I will discuss also comes from Nechyba (2000). In one sense, the simulation described is simply “a richer and more realistic model than used in prior work” (p. 131). But the improvements made for this simulation are of key interest here. The most obvious change is that of focus. This simulation intends to examine the impact of vouchers under different targeting schemes instead of public school funding schemes, like in Nechyba (1999). The overall structure of the model remains the same, with several rounds of creating communities and letting them relax into equilibrium prior to the introduction of private schools or vouchers. Once in this initial equilibrium, a voucher program is introduced that targets certain demographics or neighborhoods. After several rounds of this process with different voucher targeting schemes, the resulting equilibria are compared to each other.

A second difference pertains to the magnitude of the simulation. The model from Nechyba (1999) contains three districts, which each contain five neighborhoods full of similarly-priced houses. Each household has one of ten income levels, and a child of one of five ability levels. The product of these variables leads to a total of 750 types of agents at the beginning of the model. In contrast, the simulation developed in Nechyba (1999) only has 45 types of agents. This difference in scale does not indicate that the earlier work lacked value; it merely represents an early stage in the development of the model.

A third difference concerns the degree of calibration. The housing prices used to initialize the model come not from national data, like in Nechyba (1999), but from specific neighborhoods in the suburbs of New York City. Using data specific to a certain location increases the realism of the simulation, but detracts from its generalizability. In response to this potential challenge Nechyba attempts to regain some wider applicability, "while the calibration is using data from New York suburbs, it should be noted that I have run similar simulations calibrated to data from New York City, which yield similar results" (2000, p. 133).

The most interesting difference from the perspective of the current study is the relative importance of the analytical section compared to the computational section. Epple and Romano (1998) develop a sophisticated analytical model for most of their discussion of equilibrium and turn to simulation to resolve some particulars when the math gets out of hand. Nechyba (1999) thoroughly pursued an analytical model, but paid approximately equal attention to the simulation, and even included several important features in the simulation (e.g., the public education funding model) that do not receive treatment in the analytical section. Nechyba (2000) focuses much more on the simulation than on the analytical model. The section that describes the model references prior work by the author (specifically, Nechyba (1999)), and discusses aggregate equilibrium in many of the same terms. Unlike Nechyba (1999), though, the description of the model also contains componential information like, "Parents . . . choose (i) where to live, (ii) whether to send their child to the local public or private school, and (iii) how to vote in local elections determining the level of public school spending." Although these ideas receive some analytical treatment, the explication of the model begins with a componential description, which indicates a shift away from the analytical mindset and toward the systems mindset that employs simulation as a primary method of inquiry.

Other simulations in education

Tengs, Osgood, and Chen (2001) used a simulation to conduct a cost-benefit analysis of a promising anti-tobacco program aimed at junior high school students. This article describes a high-quality simulation and provides insight into how to calibrate a simulation, how to use a sensitivity analysis when calibration is impossible, and thoroughly discusses the limits of the simulation and what improvements can be made in future iterations.

All simulations described in this review contain some treatment of calibration in which estimated quantities are stated and a reference is provided. Only Tengs, Osgood, and Chen (2001) indicate the importance of these assumptions by placing all estimated variables in a table along with their citations. The prominence of the calibration is further indicated by the fact that the table containing the calibrating estimates occurs on the second page of the article.

Some key parameters in the model received little direction from the literature or the literature indicated a wide range of estimated values. The literature suggested values for one such parameter, the effectiveness of anti-tobacco educational programs, that varied from 5% to 80% effective. Condensing that range down to a single-point estimate eliminates important variability in the system (Hertz, 1957). Instead of artificially collapsing the effectiveness estimate to a single point, the authors "decided to evaluate a wide range of possible scenarios" (Tengs, Osgood, & Chen, 2001). The anti-tobacco program would be administered to 7th and 8th graders; the value of the effectiveness assumed in the simulation is drawn from the range suggested by the literature, and the simulation uses a different value of effectiveness for each year that the program is administered. Allowing the effectiveness of the program to vary from one year to the next accounts for normal variation within the school system (e.g., from teacher turnover, new cohorts of students, etc.). Using a range of values allows the authors to discuss different outcomes under conservative and hopeful assumptions.

Another aspect of this study that demonstrates the quality of the simulation is the sensitivity analysis. The sensitivity analyses were used "to better understand the impact of uncertain parameters on cost-effectiveness" (Tengs, Osgood, and Chen, 2001, p. 564). Unlike parameters for which a wide range of possible values are available (e.g., anti-tobacco effectiveness), the literature offered little or no direction on some values, e.g., whether former smokers have health trajectories more like never smokers or current smokers. In this example, the authors parameterize health outcomes for former smokers to be some percent of the way from the health outcomes for never smokers toward the health outcomes for current smokers. To test the impact of this estimate on the final results, they "multiplied the differential between current and never smokers . . . with a constant α, to expand or shrink these differences" (Tengs, Osgood, and Chen, 2001, p. 564). The multiplier α was chosen at random from a normal distribution centered at 1 and with a standard deviation of 0.5 for each iteration; a value of 0 would totally eliminate the difference between never smokers and former smokers, and a value of 2 would move the health outcomes twice as far toward current smokers as the authors initially assumed. The authors repeated this process 5,000 times, and reported the range of outcomes for several confidence intervals.

The authors also discuss what specific aspects of the model they left out. Of particular theoretical importance, the simulation does not account for interaction between students. One might imagine that if a student's classmates do not begin smoking, then the student would be even less likely to smoke. Thus the anti-tobacco program would prevent some students from smoking, which would reduce the likelihood that other students would smoke and make them more susceptible to the message of the anti-tobacco program, preventing them from smoking. Complexity theorists refer to this type of behavior in social systems as positive feedback (Rickles, Haw, and Shiell, 2007). The feedback loop in this system is positive because the output of the system (i.e., students choosing not to smoke) becomes the input for the system (i.e., students' peers choosing not to smoke) in such a way as to amplify the behavior. Accounting for this feedback in analytical exploration would be challenging, but including the smoking behavior of a student's peers as an input in the simulated student's decision making would not likely pose an unreasonable challenge in a numerical simulation.

The final article that I will discuss does not involve vouchers or an agent-based simulation, but does demonstrate the difference between approaching research from an analytical mindset versus a simulation mindset. Marschke, Laursen, Nielsen, and Dunn-Rankin (2007) undertook a study to examine the impact of policies aimed at equalizing the gender makeup of university faculties. Agent-based simulations must, as their category implies, be based on the behavior of an agent (Gilbert and Terna, 2000), and by this description the simulation in Marschke, Laursen, Nielsen, and Dunn-Rankin is not agent-based. Instead, they adopt an aggregate approach, but from the beginning they assume that the equations they create cannot be solved longitudinally. The equations in this study use the value of a parameter (e.g., percentage of faculty that identifies as female during a certain year) to determine the change in that parameter (e.g., the percentage of new hires that identify as female in the next year). Equations that relate the change in a variable to its value are called differential equations, and very few types of differential equations are analytically tractable (Weidlich, 2002). The authors develop a system of differential equations that basically describes how the faculty in one year will become the faculty in the next year, but they never indicate that solving these equations for some future time is possible. Instead, they say that, "at each point in time, the faculty headcount is recalculated according to changes in the previous year that are themselves time-dependent" (Marschke, Laursen, Nielsen, and Dunn-Rankin, 2007, p. 9). A computer takes the state information for one year, and uses it to compute the state for the next year, then uses that information to compute the state for the year after, and so on.

The importance of this methodological approach and the theoretical perspective that justifies it is that it allows the model to include feedback loops, described by Rickles, Haw, and Shiell (2007) as important elements of complex systems. The authors identify an inertia in the gender composition of university faculties, "new hires are more likely to be women if the overall percentage of women is already relatively high" (Marschke, Laursen, Nielsen, and Dunn-Rankin, 2007, p. 3). The use of differential equations solved year to year by a computer allows the authors to explore theoretical features of their system that would be impossible from within the analytical mindset.

Summary

The literature that I have described here grounds the argument about school choice programs, what they cost, and why they might work. Some believe that schools must participate in competitive markets if they are to use their resources efficiently. Others mistakenly extend this argument and assume that if the money follows the child, the monetary inputs required by the education system will not increase or decrease, so the cost will remain the same. Many authors critique the assumption that school choice programs are cost neutral by pointing out, among other things, that allocation effects damage the economies of scale that schools rely on and reduce the value received by the students remaining in public schools.

I have also provided some examples of how simulations can help researchers explore the economics of education, including an area of the literature devoted to using simulation to explore Tiebout equilibrium in communities with both public and private schools. Comparing earlier models to later ones demonstrates how simulations iteratively develop over time. Once a model has been developed and shown useful, future authors can add richer behavior and more detailed calibration to improve the results of the simulations and allow it to enlighten a broader range of questions. Additionally, beginning a study with simulation in mind allows the theorist to consider a grander view of the behavior of the model, because only componential behavior needs to be described mathematically.

Conclusion

I searched Google Scholar using all combinations of the phrases (a) *education*, *school*, *school district* with the phrases (b) *simulation*, *numerical model*, *computational* *model* and the phrases (c) *school choice*, *voucher*, *charter school*, *allocation effects*. Many of the results I have described came from these searches, but I found no study within these searches that used numerical simulation to estimate the economic impact of allocation effects that arise due to school choice programs. Because some researchers conducting cost-effectiveness and cost-benefit analysis of school choice programs make assumptions about these economic impacts (e.g., Yeh, 2007; Yeh, 2010), finding better estimates of them using simulation will close an important gap in the literature.

CHAPTER 3: Methods

Introduction

I have described disagreement in the literature and in popular discourse about whether faculty costs should be viewed as fixed or variable in response to small enrollment changes. More technically, disagreement exists about how sensitively faculty costs, when viewed as step costs, can respond to small enrollment changes. As a result, I propose the following research questions:

(1) Can the resources (computing power, level of detail of data and information, etc.) available to practicing evaluators generate useful predictions of the costs of school choice programs?

(2) How sensitive are faculty costs to enrollment changes caused by school choice programs?

(3) How much will a proposed school choice program cost due to inefficiencies created by allocation effects?

Beyond the obvious value of answering these questions, answering them with numerical simulation will also help practitioners in the field of cost estimation in education become more comfortable with an approach and a tool set commonly used in other fields.

This chapter will describe the methods I will use to answer these research questions. I will begin by discussing a simplified version of the simulation that will help readers understand the process of conceptualizing, creating, and running a simulation. I will then describe both the general type of data that can create the initial state of a simulation like this, and the actual data I used in creating this simulation. Finally, I will cover the rules that govern the behavior of a simulation both generally and in detail about this research .

Study Design

Previous studies that have explored the relationship between enrollment and the cost of providing education are based on regression using many years of data (Chibulka, 1983; Heinesen, 2004). Using regression to analyze this data compares one year's enrollment and some metric (i.e., faculty size) to other years' enrollments and metrics. The trouble with this approach is that the signal of small enrollment declines due either to random variation or to school choice programs might get swamped by the noise of larger enrollment changes due to such factors as long-term demographic shifts, school openings or closings, changes in budgeting, school reform movements, etc. If we suppose that small enrollment declines function differently than large enrollment declines, then systems with a mix of large and small enrollment changes cannot help us understand the impact of small enrollment changes.

Ultimately, we wish to know how small changes in one year's enrollment will impact the faculty size *during that same year* in two parallel universes, one of which decided to implement a school choice program and the other did not. Because our interest lies in comparing a situation that exists to one that does not, building the entire system as a simulation offers the best method to address the research questions stated above.

The simulation that I will describe is based on the process of financial stress-testing employed by economists to measure the fiscal soundness of banks. The stress-testing models used in the financial world, as described by Sorge (2004), involve four primary steps: (1) represent the current state of the system, (2) define the rules that govern how the components interact, (3) apply a counterfactual change to the components—this is the stress to the system—and allow the components to interact and respond, (4) examine the resulting state of the system. I will follow this model closely for steps one through three, but my approach to step four will differ.

Sorge (2004) describes two methods of building a simulation of a financial system that economists might wish to stress-test. The first method involves constructing "hypothetical portfolios whose composition is intended to mimic the distribution of assets and risk exposures within a given financial system" (p. 4). The second method involves actual data, but this method is less common due to the unavailability of necessary information. Once the simulation has been created, the stress scenarios must be designed. Sorge (2004) offers the following summary,

> the design of any stress scenario includ[es] the choice of the type of risks to analyse (market, credit, interest rate, liquidity, etc.), whether single or multiple risk factors are to be shocked, what parameter(s) to shock (prices, volatility, correlations), by how much (based on historical or hypothetical scenarios) and over what time horizon. (p. 4)

The value of using a simulation comes from examining the impact of what Sorge (2004) calls "second-round effects" and what complexity theorists call "feedback" (Manson, 2001, p. 407). Once the shock has been applied to the system, the system will respond; the changes that occur during this response can themselves influence the system. These feedback effects are challenging to predict over short intervals, and essentially impossible to predict analytically over long intervals because complex systems exhibit final state sensitivity to initial conditions (Grebogi, McDonald, Ott, and Yorke, 1983). Imperceptibly small changes in the initial state can be amplified during feedback loops, leading to obvious and dramatic differences in the end result; only simulation can help researchers explore this behavior.

Much of my proposed simulation will follow this basic structure. Instead of initializing the simulation with financial data, I will use enrollment data for an entire district; instead of, say, making interest rates more volatile, I will remove a small number of students from the district. My approach to stress-testing a school district will follow closely the framework described by Sorge (2004) as described so far, but I will take a different course in the final step. The final step in financial stress testing involves defining a condition in the response variable in which the system is said to be in “crisis” (Sorge, 2004, p. 3). The financial system uses a binary outcome of the stress-test: in crisis versus not in crisis. My study, on the other hand, will calculate the elasticity (as a continuous variable) of faculty reduction that results from the enrollment decline.

A Simplistic Description of the Simulation

To elucidate my proposed method of estimating the costs of school choice programs, I will propose a toy example. Consider a fantasy school district with ten schools each containing one grade and taught by generalist (i.e., not content-specific) teachers; see Table 3.1. This model is clearly simplistic; a version of this model that is useful beyond demonstration will have to mimic or describe the actual state of the district in much more detail. The names of the schools, *School*, are randomly chosen from a list of words representing letters in the alphabet.

Table 3.1. The ten schools in the fantasy district.

School
Hartle
Grant
Uptown
Wessex
Pond
Zebra
Icelandic
Juniper
Verona
Open

I can add a random number of students, *Enroll*, to each school, uniformly chosen from [100, 200]; see Table 3.2. The decision to add a purely random number of students is an important assumption in the model—and one not supported by the literature. In practice, the number of students in each grade for each school is publicly available in many districts, and so is easily obtainable by evaluators.

Now that each school has a number of enrolled students, I can use some rule to determine *Faculty*, the number of teachers each school will require. For this toy example, I will set the maximum class size rigidly at 25 students and use that rule to determine a faculty size for each school; see Table 3.3.

Table 3.2. The school district with a random number of students in each school.

School	Enroll
Hartle	115
Grant	179
Uptown	124
Wessex	145
Pond	134
Zebra	135
Icelandic	196
Juniper	177
Verona	191
Open	177
Total	1573

The rule of using a rigid maximum class size to determine the number of faculty likely fails to capture the intricacies of education in the real world, but it will suffice for this toy example. See chapter four for a description of how I will handle decisions about faculty size in the actual model.

The model at this point describes the initial state of the fantasy district. Now we can borrow the method from financial stress-testing and subject the district to an imaginary stress. In this model, the imaginary stress takes the form of removing students from a school, *Students Removed,* to participate in a school choice program. The district contains a total of 1573 students, and I will subject the model to a school choice program that randomly removes a total of 25 students. The decision to remove these students purely randomly is an important, and perhaps inappropriate, assumption; see chapter four for coverage of how that assumption will be tested. Once the students are removed, I can calculate each school's new enrollment, *New Enroll*; see Table 3.4.

Table 3.3. The school district including faculty employed at each school.

School	Enroll	Faculty
Hartle	115	5
Grant	179	8
Uptown	124	5
Wessex	145	6
Pond	134	6
Zebra	135	6
Icelandic	196	8
Juniper	177	8
Verona	191	8
Open	177	8
Total	1573	68

We can then use the same rule that we used to calculate how many teachers each school employs to calculate how many teachers each school would need if the school choice program were to take these students, *New Faculty*. Comparing this faculty size to the actual faculty size allows us to predict how many positions could be eliminated, *Faculty Loss*, given our assumptions; see Table 3.5. Notice that the school choice program leads to the elimination of three teaching positions due to the random removal of this specific group of students.

Removing a different selection of students randomly from the district described in Table 3.4 will lead to a different result; see Table 3.6. Notice that removing the same number of students—but different specific students—leads to the elimination of one faculty positions; two fewer than in the previous example.

Table 3.4. The district after a counterfactual stress of removing students for a school choice program.

School	Enroll	Faculty	Students Removed	New Enroll
Hartle	115	5	2	113
Grant	179	8	4	175
Uptown	124	5	2	122
Wessex	145	6	3	142
Pond	134	6	3	131
Zebra	135	6	2	133
Icelandic	196	8	3	193
Juniper	177	8	3	174
Verona	191	8	1	190
Open	177	8	2	175
Total	1573	68	25	1548

Table 3.5. The results of the stress-test: how many faculty positions can be eliminated as a result of the school choice program.

School	Enroll	Faculty	Students Removed	New Enroll	New Faculty	Faculty Loss
Hartle	115	5	2	113	5	0
Grant	179	8	4	175	7	1
Uptown	124	5	2	122	5	0
Wessex	145	6	3	142	6	0
Pond	134	6	3	131	6	0
Zebra	135	6	2	133	6	0
Icelandic	196	8	3	193	8	0
Juniper	177	8	3	174	7	1
Verona	191	8	1	190	8	0
Open	177	8	2	175	7	1
Total	1573	68	25	1548	65	3

Table 3.6. A different impact on faculty of counterfactually removing 25 students from the district. The students removed in this trial differ from the students removed in the trial described in Table 3.5.

School	Enroll	Faculty	Students Removed	New Enroll	New Faculty	Faculty Loss
Hartle	115	5	2	113	5	0
Grant	179	8	3	176	8	0
Uptown	124	5	5	119	5	0
Wessex	145	6	1	144	6	0
Pond	134	6	0	134	6	0
Zebra	135	6	4	131	6	0
Icelandic	196	8	2	194	8	0
Juniper	177	8	1	176	8	0
Verona	191	8	4	187	8	0
Open	177	8	3	174	7	1
Total	1573	68	25	1548	67	1

Table 3.7. The scenario described in Table 3.6, but with cost information.

School	Fixed Costs	Enroll	Rev.[a]	Fac.	Salary[b]	Serv. per Stu.	Stu. Rmvd	New Enroll	New Rev.[a]	New Fac.	Fac. Loss	New Salary[b]	New SPS	Benefit/ Lost
Hartle	$100	115	$805	5	$250	$3.96	2	113	$791	5	0	$250	$3.90	$6.09
Grant	$100	179	$1,253	8	$400	$4.21	3	176	$1232	8	0	$400	$4.16	$8.38
Uptown	$100	124	$868	5	$250	$4.18	5	119	$833	5	0	$250	$4.06	$14.11
Wessex	$100	145	$1,015	6	$300	$4.24	1	144	$1008	6	0	$300	$4.22	$2.76
Pond	$100	134	$938	6	$300	$4.01	0	134	$938	6	0	$300	$4.01	$0.00
Zebra	$100	135	$945	6	$300	$4.04	4	131	$917	6	0	$300	$3.95	$11.85
Icelandic	$100	196	$1,372	8	$400	$4.45	2	194	$1,358	8	0	$400	$4.42	$5.10
Juniper	$100	177	$1,239	8	$400	$4.18	1	176	$1,232	8	0	$400	$4.16	$2.82
Verona	$100	191	$1,337	8	$400	$4.38	4	187	$1,309	8	0	$400	$4.33	$10.47
Open	$100	177	$1,239	8	$400	$4.18	3	174	$1,218	7	1	$350	$4.41	-$41.53
Total	$1,000	1573	$11,011	68	$3,400		25	1548	$10,836	67	1	$3,350		$19.93

Note. All cost figures in this table are displayed in thousands of dollars.
a Revenue columns assume a PPE of $7,000
b Salary columns assume that each teacher receives $50,000 in total compensation.

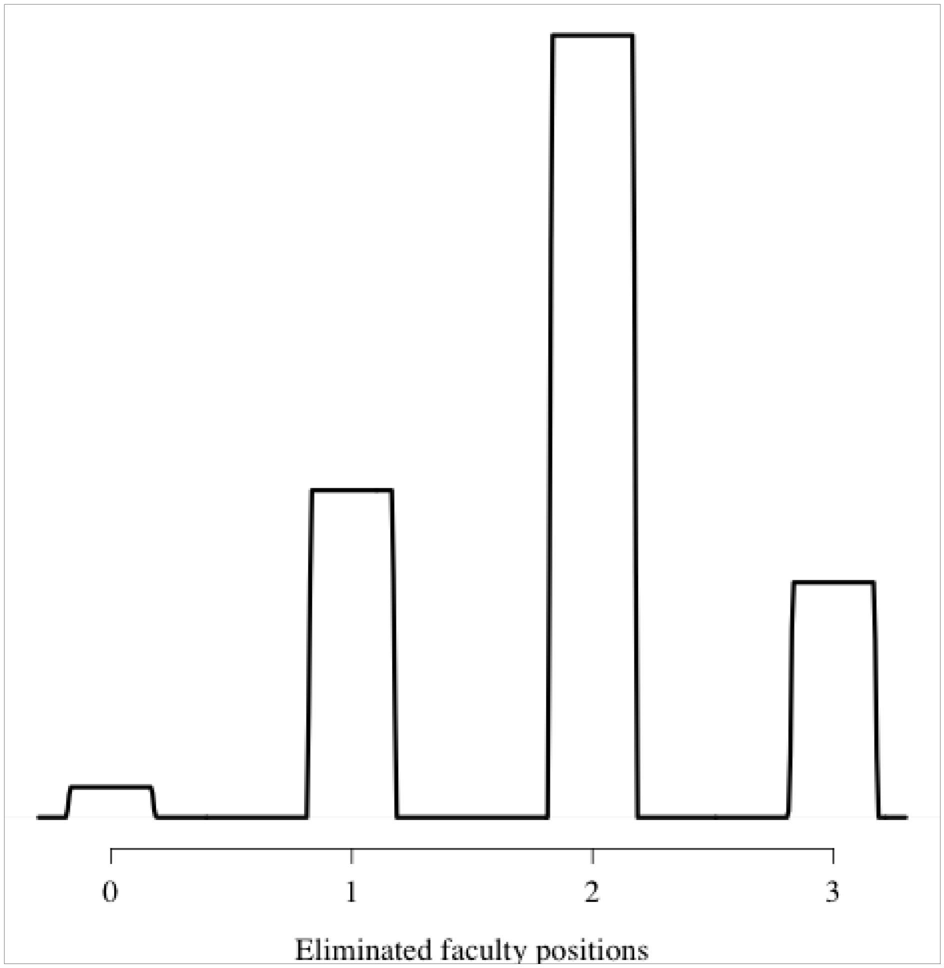

Figure 3.1. A histogram showing the relative frequency of eliminated faculty positions in 1000 trials of removing a random set of 25 students from the fantasy district using the enrollment numbers shown in Tables 3.2-3.6.

Repeating this process 1000 times will allow us to create a histogram showing the frequency of faculty position eliminations with different groups of students randomly removed; see Figure 3.1.

Table 3.5 shows that the size of the district's faculty will decrease by 4.4% (from 68 teachers to 65 teachers) when enrollment declines by 1.6% (from 1573 students to 1548 students), which indicates an elasticity ratio of 2.75. The 1.5% faculty reduction (from 68 teachers to 67 teachers) shown in Table 6 (as a result of an identical 1.6% enrollment decline) indicates an elasticity ratio of 0.92. Recall that the example here is meant only to demonstrate the function of the model, so these elasticity ratios will not help us make policy decisions in the real world, but calculating elasticity ratios for many trials using actual enrollment data will create a meaningful frequency distribution (like the one shown in Figure 3.1) that I can use to address the first research question.

Addressing the second research question will require estimating financial figures like per pupil expenditure (PPE), faculty salaries, and fixed costs for schools. The money left over after paying the fixed costs and faculty costs is available to provide other educational services to the students. The accounting spreadsheet will use these figures to calculate the amount of money available to purchase services per student—this figure is shown as Serv. per Stu. or SPS in Table 3.7. Comparing the sum of the services available per student before and after the counterfactual stress of removing students for the school choice program will show the economic impact of the allocation effects caused by the school choice program. As with the process described above, I will iterate this simulation and create a frequency distribution of the costs to students. Repeating this process for different combinations of districts and school choice programs will allow me to address the second research question.

The Role of Time

Before leaving this toy model, I will offer some clarification about the role of time in this simulation. This simulation attempts to help us understand the first year of a new school choice program. I use words like "before" and "after" to compare the numbers of students and teachers in the district without the stress and with the stress.

But these words can be difficult to interpret because we often view school years—in terms of a student attending a school or a school employing a teacher—as existing as a discrete object. It may at first seem odd to discuss “before” and “after” in an indivisible chunk of time. I will offer two functionally identical ways of viewing these temporal terms, and then an argument that which perspective we choose has little impact on our final conclusions.

The words “before” and “after” can refer to different moments in time. The “before” moment in time could either be when teaching contracts are signed prior to the beginning of the school year or when long term estimates are created about school enrollment within a district. If the decision to implement a school choice program and its implementation come after that moment in time, then all future enrollment figures and faculty decisions would be “after” that process. The two examples of before and after described here might lead to different decision making: perhaps fewer or no teaching positions would be eliminated if those eliminations entailed changing contracts. I do not favor the strict temporal interpretation of time in this study, so I will not attempt to address those differences here.

The words “before” and “after” can also refer to before and after the hypothetical creation of a parallel universe that includes a school choice program in the district. When faced with a decision about a contentious policy like a school choice program, what policy makers want from research is a glimpse into two alternate futures. They can then compare the alternate futures and choose on the course of action that will place them in the more desirable future. This simulation attempts to do exactly that: describe a school district as it currently is with no school choice program, and as it might be with one. I will take this perspctive in this paper, because it offers the most help to policy analysis.

However we view the words “before” and “after” in this study, they intend to describe the impact of implementing a school choice program. In this sense, the strict temporal perspective and the alternate future perspective would lead us to build largely identical simulations, so disagreement on this point should not change how we proceed.

Description of Data

The raw data for this simulation come from enrollment data and free/reduced lunch eligibility data from California in 2011. The enrollment data come from the Enrollment by School section of the California Department of Education website. The raw enrollment data contain a row for each race code in each grade of each public school in the state, 126,083 individual rows in total. These rows contain information about 1,026 districts, and 10,216 schools.

The free/reduced lunch eligibility data come from the Student Poverty - FRPM Data section of the California Department of Education website. The raw free/reduced lunch data contain one row per school, each with the number of students eligible for free lunch and the number of students eligible for reduced lunch in that school.

The districts range in size from 1 school (in 223 districts) to 916 schools in the Los Angeles Unified school district. Because only districts with more than one school will exhibit different behavior in the second level of iteration (described below and in chapter four), the simulation described in this work will only use the 535 districts with more than one elementary school. I have two reasons for only using grades one through three in this simulation. First, the class size rules are the same for grades one through three in California, but different in kindergarten and 4th grade (Gonzalez, 2013). Third, young elementary grades are more often taught by generalists instead of content area specialists. This means that elementary schools can theoretically more easily tailor their faculty size to enrollment, and thus have the greatest chance of responding to small enrollment declines. See the section on Generality in chapter five for a more thorough justification for decisions like this one.

Some explanations will benefit from using a single district as an example, in these cases I will use the Fallbrook Union Elementary school district in my examples. I chose this district because it has the smallest Mahalanobis distance to the centroid of the data (a Z-score of 0.000071) in terms of number of students and percent of students eligible for free lunch. In other words, this district is closest to average in terms of size and average socioeconomic status of all districts in California.

Anomalies and Data Cleaning Procedures

The data used in this simulation are not perfect. For example, 13.2% of districts have at least one school with more students eligible for free or reduced lunch than total students enrolled. The great majority of these anomalous schools fall into two categories. The first category involves schools that have, say 766 total students, but 767 students eligible for free or reduced lunch. The other type of anomalous datum is that of a category of student or a non-classroom education program, as opposed to an actual school. As stated above, the data for this study come from two sources, which alone may explain some of the discrepancy: the two data sources may have slightly different recording practices or definitions in the agencies that assemble the data. But it may also simply indicate the limit of precision of the individual measurements in these data. The final example of seemingly erroneous data, which are very uncommon, come from individual schools for one reason or another. One such example has only a single student enrolled based on the enrollment data, but has 30 student eligible for free or reduced lunch based on the free/reduced lunch data. The school in this example, a charter school, opened the year these data were gathered, 2011. The particular circumstance surrounding the launching of a charter school may explain why two agencies have different records for that school during that year. I can only observe discrepancies between these data sources; any errors in the data that do not conflict with the other data source are not detectable, so this is the only type of error I could observe.

No information was removed from the raw data. Instead, a function in the code tested whether the data for a district contained any conflicts prior to executing the simulation on that district. If the district was found to contain any anomalies, that school was removed. Many of the anomalous entries in the district were non-school education programs, which were removed because they contain too few students; see the description of thinning in the Varying Policy Parameters section of chapter four. Because these entities would be removed in the thinning process, a negligible number of anomalies remained in the data when the function testing for them ran.

Creating the Initial State of the Simulation

Calibrating this simulation with information from the real world will fall into two tasks: creating the initial state, and describing the componential behavior. Of these two tasks, creating the initial state will pose fewer challenges than describing the componential behavior. In this section I will first offer a very abstract description of the initial state, which might be of use to future researchers creating a similar but not identical simulation. I will then describe the data from California used in this simulation, and finally I will describe in detail how I turned those data into the initial state of this simulation.

A General Description of a Simulated District The fundamental object in this simulation is a district d containing a set of schools $\{s\}$, each of which contains its own set of 13 grades $\{g_j\}$, and a set of students $\{q\}$ who belong to grades in those schools, and a set of teachers $\{t\}$ who teach in those schools. A student q_{ijk} is the kth student in the jth grade in school s_i, and Q_{ij} represents the total size in the jth grade of school s_i. Similarly, T_i represents the total size of the teaching faculty at school s_i. Students and teachers in the district can be assigned various parameters as they become useful in the simulation; e.g., in this study I assigned each student a free lunch status, and each teacher a salary.

In practice, this simulated district could be created in a number of ways. I obtained publicly-available enrollment data from the California Department of Education that is aggregated at the school/grade level. Evaluators working with a school district could create the simulated district from actual student rosters. On the other end of the spectrum, researchers that are much more interested in exploring behavior than generating predictions about existing districts could create a simulated district by selecting district size and grade sizes from distributions.

Creating the Simulated District in this Study Before any simulated stress can be placed on a district, I must create an object that represents the district in the computer. I used the publicly available data about grade-level enrollment at schools in California and school-level free/reduced lunch eligibility discussed above to construct a list wherein each row represents a school. Each school in the simulated district contains the following: the school name, the district name, the enrollment in each grade from kindergarten through 12th grade, and the percentage of students eligible for free lunch. The raw data I obtained from the California Department of Education website has a row for each school/grade/race code that contains the enrollment for that group of students. I summed the enrollment in all the race codes for a grade in a school to determine the total enrollment in that grade in that school. I repeated that process for each grade in the school, and for each school in the district. With up to seven race codes and 13 grades, each row in the simulated district matrix that I create originally comes from up to 91 rows in the raw data.

I created a second set of grade-level enrollment figures in each row that show how many students in each grade are eligible for free lunch. Because grade-level free lunch statistics are not publicly available, I assumed that the percentage of students in each grade that are eligible for free lunch is equal to the percentage of students in the school eligible for free lunch; i.e., the students eligible for free lunch in the school are proportionately distributed among the grades.

The simulation executes two processes that are not based on actual data: assigning teachers to schools and counterfactually removing students as a part of the stress placed on the district. These processes are discussed in detail below, but the reader might find it helpful to know, from a data structure standpoint, that the number of teachers in the school before the stress, the counterfactual enrollments after the stress, and the number of teachers in the school after the stress are all added to each row in the simulated district matrix during the simulation.

For the financial simulation, I will also need several economic estimates for each state or district modeled by the simulation. In some cases I will obtain a point estimate for the figure, in other cases I will obtain a data-based set of values that represents the actual distribution, and in the remaining cases I will estimate a distribution from which to choose values (Tengs, Osgood, and Chen, 2001). The parameters I will need are (a) per pupil expenditure, PPE (b) faculty salaries, (c) the faculty benefit (i.e., health insurance, retirement contribution, etc.) ratio, (d) the fixed cost of operating the school. See Table 3.7 for the role of these estimates in the the financial simulation (the faculty benefit ratio is, in fact, absent from Table 3.7, but it will modify the salary in the actual model).

Modeling the Componential Behavior

The componential behavior of the agents and the environment in the simulation will consist of decisions that the computer makes, either pseudo-randomly or based on rules applied to some input. Transition from one state to the next in the simulation could involve some or all of the following: (a) selecting a student for participation in the school choice program, (b) determining where the selected students goes after selection, (c) applying the impact that a student's selection for participation has on other students in the district, (d) determining when to stop selecting students for participation, and (e) calculating the number of teachers each school requires.

The number and type of students participating in the school choice program—the shock to the system in the stress-test model—could be built to mimic existing school choice programs and proposals, or could span the range of parameters that a school choice program might take. The toy example described above shows one version of this stress: The scenarios shown in Tables 5 and 6 depict school choice programs that draw the same number of students (25, or 1.6%) and in the same way (purely randomly). Other stress situations will involve a greater or lesser enrollment decline and/or will target some demographic groups and/or will allow the decision of a student to impact the likelihood that other students in his or her class will also participate in the school choice program, as described in Tengs, Osgood, and Chen (2001).

How faculty size is determined based on enrollment will pose a greater challenge. States and school districts sometimes post class size targets, but individual schools do not always adhere to these goals (J. Franks, personal communication, 27 July 2012); sometimes principals allow class sizes to exceed the posted targets so that more students can attend their school. Because deciding on a number of teachers to employ is a rule—and not a value—using the technique of creating a distribution and selecting a number from that distribution (as used, for example, in Tengs, Osgood, and Chen, 2001) is not available in this case. Instead, I will develop a set of possible rules that will span the set of ways to select a number of teachers based on a number of students, and select a rule from this set. Allowing the hiring and firing rules to vary in different runs of the simulation will allow me to examine the impact of those rules on the outcome of the model, as in Nechyba (1999).

Iteration

With the initial state and componential behavior in place, I can conduct a single fundamental run of the simulation. Like other Monte Carlo simulations, this simulation will consist of iterating a simple step, recalculating the state of the system, and repeating either for some duration or until some terminating event occurs (Eckhardt, 1987). The simple step that will be repeated is removing a single student from the matrix; recalculating the state of the system will involve recalculating the number of teachers in each school. This sequence of steps will continue until the desired number of students have been removed. To conclude a single fundamental run of the simulation, the percentage of teaching positions eliminated and the economic impact will be recorded in a data file.

For the second stage in the nested iterations I will repeat the process described above with identical initial states and componential rules many times. The percentage of teaching positions eliminated for all of these fundamental runs will be collected into a frequency distribution (as in Hertz, 1957). This frequency distribution will allow me to generate point estimates (e.g., the ratio of percentage of teaching positions eliminated to percent enrollment decline) but also to discuss descriptive statistics (e.g., standard deviation and skew). This same approach, but with financial information, will allow me to create and analyze a frequency distribution for economic impact as a proportion of per pupil expenditure.

The third stage in the nested iteration will involve varying the parameters used to create the initial state of the simulation (except for the district parameters) and the rules that govern agent behavior. The simulation will run the two lower iterations described above and generate point estimates, which will be stored in a data file along with the values for the parameters used in that run. I will conduct a sensitivity analysis using ANCOVA to determine how the parameters influence the faculty responsiveness and economic impact.

The fourth and final level of iteration will involve applying the process described above to districts with different demographics. I will use 2-level hierarchical linear modeling (HLM) to assess the impact of demographic factors on the aggregate behavior described in the third level of iteration (Willms, 1999). The results of the HLM will allow me to answer the first and second research questions described above.

Conclusion

In this chapter I have presented a simplified model of the simulation to elucidate some important aspects of its function. I have also described the data I will use in this research, explained the componential rules that the simulation will use, and discussed the levels of iteration.

The reader will notice that I have not offered an extremely thorough plan for how the simulation will come together. Gilbert & Terna (1999) argue that creating a simulation can teach us more about a system than the results of that simulation alone. For this reason, I feel it inappropriate to describe the fully formed simulation in the methods section of this work: what I have to say about the structure and features of this simulation is as much a result of this research as the statistics and graphs.

CHAPTER 4: Results

Introduction and Structure

Simulation is both a tool and a process (Gilbert & Terna, 1999). I have decided to save the detailed descriptions of how the simulation was created until this results chapter because much of what we can learn about systems from simulation comes from what we will notice while creating them.

This chapter contains a description of the research done in this study. I will begin by describing the four levels of iteration in the simulation, and draw some conclusions when describing the fourth level of iteration. Finally, I will apply the simulation to a financial model of a school district and draw the remainder of conclusions presented in the work.

Levels of Iteration

Basic Step: Removing a Single Student

The fundamental operation in this model involves removing a single student from the simulated district. This action simulates a family deciding to have their student participate in the school choice program by enrolling in a charter school or a private school on a voucher (I will henceforth refer to charter schools and private schools as *non-district schools* to differentiate them from open-enrollment-like school choice programs). Because the simulated district only contains the public schools under the authority of the district, the students participating in the charter school or voucher-based school choice program are removed from the list of enrolled students, and not replaced. A more complete model of this system would then add that student to the list of students enrolled at the non-district school of the family's choice. If the school choice program in the simulation were an open enrollment model, the student would not be removed from the list of students enrolled in the district, rather the student's designated school would be changed to a different school in the district. Most of the controversy surrounding school choice programs specifically involves voucher programs (Editorial Projects in Education Research Center, 2004). For this reason, I will focus primarily on sending students to non-district schools.

The process of selecting a student for participation begins by generating a pool of students from all schools in the district for a particular grade using school-level enrollment figures. Each student in the pool has a weight which represents the likelihood that the student will be selected compared to other students in that grade in the district. How these weights are determined (e.g., based on free/reduced lunch status or special education program) and used in the selecting of students is described in the Varying Policy Parameters section below.

I use a two-step pseudo-random process for deciding which students to remove. First, the total weight is calculated for each school by summing the weights of all the students in that school. The first pseudo-random number then determines which school the student will come from. Once the school is chosen, a second pseudo-random number selects the student from that school based on the students' weights. This process is mathematically equivalent to selecting a student from the pool of all students in that grade in the district—that is, a particular student is equally likely to be chosen for participation in either pseudo-random selection scheme—but the two-step process runs much faster on a computer: In the largest district in California, Los Angeles Unified, removing students is 40 times faster using the two-step process than by removing students from the pool of all students.

Level 1: Removing Students from a District

The first level of iteration involves repeating the basic step until a pre-determined number of students have been removed, and the number of teachers in the district is determined both before and after the student removal process. The principal non-financial measurement in this analysis is the sensitivity of faculty size to enrollment, which I will refer to as the faculty/enrollment sensitivity, β; that is, the percent change in teachers p_T per percent change in students p_N, or

$$\beta = \frac{p_T}{p_N} = \frac{\Delta T}{T}\frac{N}{\Delta N} \tag{4.1}$$

Some policies only impact some grades, though. In this analysis T is the number of teachers teaching grades affected by the school choice program, and N is the number of students in the grades affected by the school choice program. Decisions about how many students to remove and how to determine the number of teachers needed to teach the students in a district will be discussed in the Varying Policy Parameters section below.

Level 2: Monte Carlo Iteration of Level 1

The impact of removing a small fraction of the students in a district will depend on which particular students are removed in addition to depending on the policy parameters and district characteristics. Because of the pseudo-random nature of Level 1, two runs on the same district with the same policies could have different outcomes in terms of the faculty/enrollment sensitivity. Table 4.1 shows actual data from the Fallbrook Union Elementary school district in California from 2011, along with three runs of Level 1.

Table 4.1 demonstrates that the size, or in some cases even existence, of faculty reductions depends on exactly which students decide to participate in the school choice program, even when the program parameters are the same. The faculty/enrollment sensitivities shown in this example suggest very different conclusions about the cost of the school choice program. Trial 1, with a sensitivity of 0.789, demonstrates that sometimes relatively little faculty reduction is possible, and the district will indeed lose funding, but not realize commensurate savings in faculty costs to offset them. In the case of Trial 1, money would need to be cut from elsewhere in the school's budget, likely leading to a reduction in value to the students who do not participate in the school choice program; see the Monetary Accounting Model section below for a more complete discussion of the implications of faculty/enrollment sensitivity on budgets.

Trial 2 in Table 4.1, in contrast, shows a faculty/enrollment sensitivity of 1.578, meaning that if 5% of students left the district, 7.9% of the teachers could be let go. Depending on how much of the district's spending is truly fixed, this outcome might free up more resources per student and be an overall financial victory for the district and a boon to its students.

Trial 3 falls in between: faculty size can decrease by 105% of the enrollment decline. This number is relatively close to a sensitivity of 1, and depending on how much of the district's spending is fixed, might not lead to a significant decline or increase in resources per student who remain in the district.

Table 4.1

An example of three runs of removing 5% of students from grades one through three in the Fallbrook Union Elementary school district

Trial	School	Grade 1	Grade 2	Grade 3	Teachers	Sensitivity
Original	San Onofre	106	102	99	12	
	Iowa Street	45	45	48	6	
	Frazier	97	95	87	10	
	Fallbrook	118	106	91	11	
	Live Oak	98	94	101	11	
	Pendleton	159	145	108	14	
	La Paloma	94	114	99	11	
Trial 1	San Onofre	100	99	98	12	0.789
	Iowa Street	43	43	45	6	
	Frazier	93	88	82	9	
	Fallbrook	113	100	86	11	
	Live Oak	92	87	97	10	
	Pendleton	150	140	101	14	
	La Paloma	90	108	92	10	
Trial 2	San Onofre	102	95	92	10	1.578
	Iowa Street	44	44	47	6	
	Frazier	92	88	84	9	
	Fallbrook	109	102	87	11	
	Live Oak	91	90	95	9	
	Pendleton	152	141	100	14	
	La Paloma	91	105	96	10	
Trial 3	San Onofre	101	99	92	11	1.052
	Iowa Street	42	44	46	6	
	Frazier	91	86	79	9	
	Fallbrook	114	100	87	11	
	Live Oak	92	90	98	10	
	Pendleton	151	142	103	14	
	La Paloma	90	104	96	10	

Note. This table uses a district ceiling rule; see below for details about this rule. This high average sensitivity is not representative; see the A Variety of Districts section below.

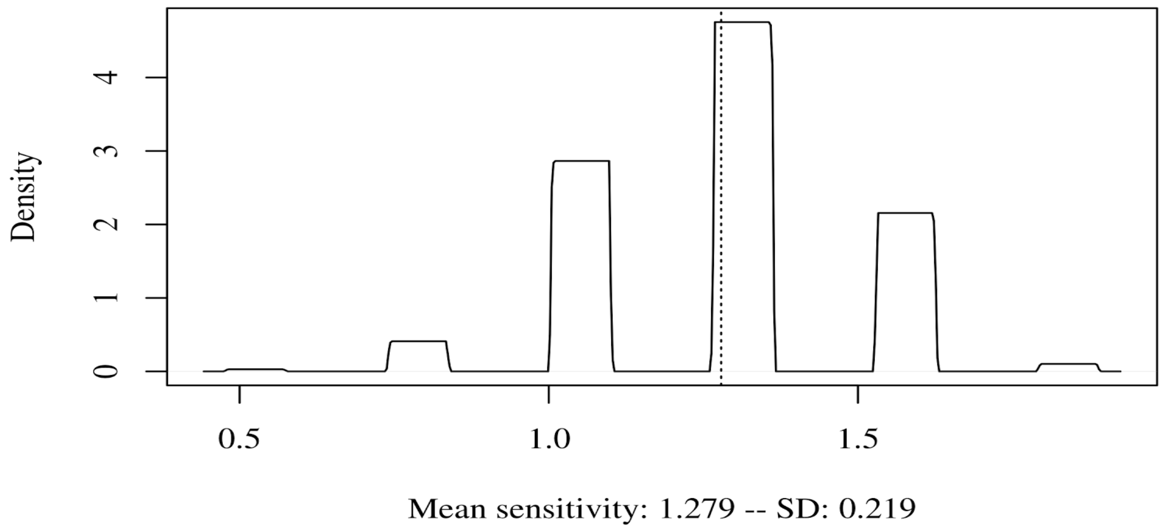

Figure 4.1. A histogram showing the frequencies of three faculty/enrollment sensitivity outcomes from removing 5% of first through third graders from the Fallbrook Union Elementary school district. This plot represents 10,000 trials of removing 5% of students and using a district ceiling rule to determine the number of teachers. The dotted vertical line (on the left side of the center bar) represents the mean faculty/enrollment sensitivity.

By running many trials, we can generate a list of faculty/enrollment sensitivities. We can then create a point estimate with an uncertainty estimate—for this district the point estimate, the mean faculty/enrollment sensitivity, is 1.279 and the uncertainty estimate, the standard deviation, is 0.219. On the other hand, Hertz (1957) would likely suggest that if we knew the relative frequencies of these outcomes we could make a better decision about the financial result of this school choice program than simply knowing a point value. Figure 4.1 shows the probability density of the outcomes of removing 5% of students in grades one through three from the Fallbrook Union Elementary school district. As seen in this plot, the cost saving 1.315 sensitivity outcome is most likely by a good deal. The relatively expensive 0.789 sensitivity outcome is quite unlikely, and the potentially cost-neutral 1.052 sensitivity is slightly more common than the even greater cost saving 1.578 sensitivity outcome. As with any single piece of information, I would not recommend that this plot be used alone to make the decision about whether or not to implement this school choice program, but those making that decision would likely find this information helpful.

Level 3: Varying Policy Parameters

Teacher assignment rules. ***Thinning districts.*** Several aspects about the assignment of teachers can be varied in this analysis: the softness of the class size ceiling, the level of district coordination, and which teachers teach classes that are impacted by the school choice program.

Most entities in a district are schools, but some are other educational programs. These other educational programs include support centers and alternative learning academies, which exist within another school. Others, like the Nonpublic, Nonsectarian designation, refer to students enrolled in a range of education programs outside the district, but paid for by the district (California Education Code, 2006). Because the students in these grades are unlikely to be taught as a group by a single teacher, it may be reasonable to remove these entities from the simulated district. Ignoring these education programs in the model, a process I will refer to as thinning, is accomplished by replacing all school/grade enrollment figures below a certain threshold, N_{min}, with zero. Figure 4.2 shows a kernel density plot of sensitivities determined using a soft district rule (see Ceiling Softness section below for a description of the rules that are used for determining the number of teachers in a district) for determining faculty size when the Pomona

Unified[1] district is thinned and when it is not. We can conclude that thinning the district does not seem to have an significant effect on the mean faculty/enrollment sensitivity ($t(1998) = 1.006$, $p = 0.315$). The consistency of this insignificance is discussed in the A Variety of Districts section below.

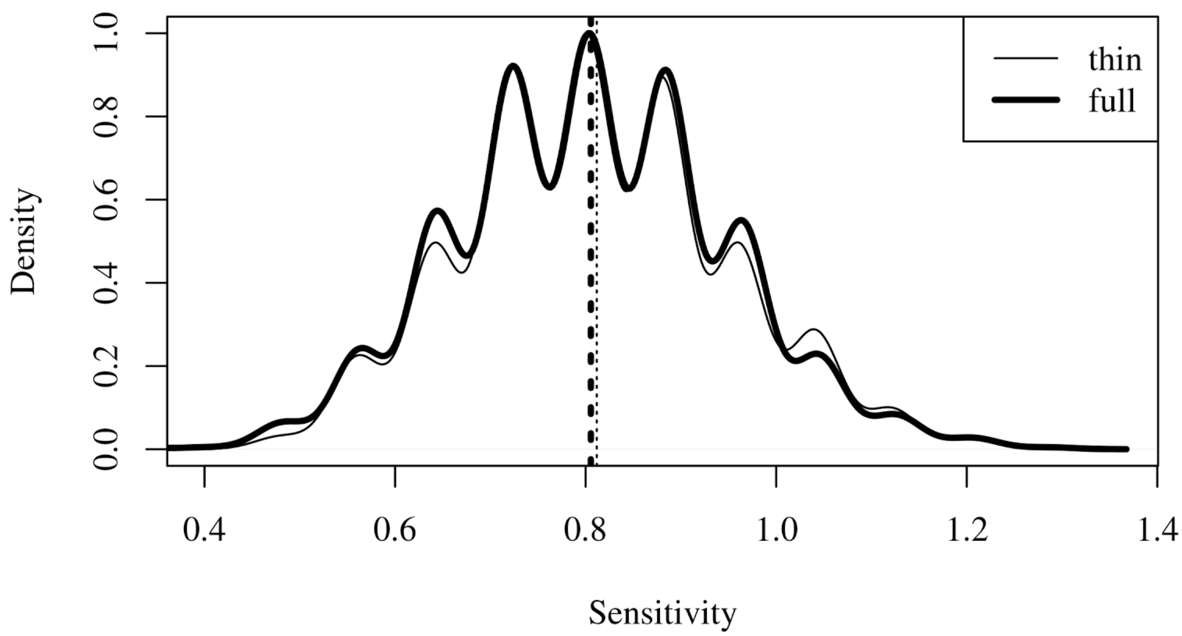

Figure 4.2. A frequency distribution of sensitivities in the Pomona Unified school district using a soft district ceiling before thinning the district of grades with less than 20 students enrolled (Full, in the legend) and after thinning it (Thin, in the legend). The mean faculty/enrollment sensitivities are not significantly different.

Ceiling Softness. California law suggests three ways of using class size to determine the number of teachers required to teach a grade in a school in this simulation,

1 I used the Pomona Unified school district instead of the Fallbrook Unified district because the Fallbrook Unified district had no grades in any schools with small enough enrollment to be removed in the thinning process.

a hard ceiling, a soft ceiling at the district level, and a soft ceiling at the school level. The "hard ceiling" rule uses the stated district average class size targets shown in Table 4.2 as the absolute maximum class size for each classroom. The same rule is used when determining the number of teachers in a district both before and after the students are removed to participate in the school choice program.

Table 4.2
Class size targets for California public schools, grades K-8

Grade Range	District Average Maximum	Individual Classroom Maximum
Kindergarten	31	33
Grades 1 - 3	30	32
Grades 4 - 8	30	30

Note. Class size figures stated in sections 41376 and 41378 of the California Education Code (2006)

The "soft ceiling" rules suggested by California's class size policy allows some classes to exceed the class size targets by a small number of students (in this case, two students per class in grades K-3) so long as the average class size for the district remains below the target. The following excerpt from Section 41367(a) of the California Education Code (2006) describes the method of calculating the excess number of students in a district:

> For grades 1 to 3, inclusive, [the Superintendent of Public Instruction] shall determine the number of classes, the number of pupils enrolled in each class, the total enrollment in all such classes, the average number of pupils enrolled per class, and the total of the numbers of pupils which are in excess of thirty (30) in each class.
>
> For those districts which do not have any classes with an enrollment in

excess of 32 and whose average size for all the classes is 30.0 or less, there shall be no excess declared. For those districts which have one or more classes in excess of an enrollment of 32 or whose average size for all the classes is more than 30, the excess shall be the total of the number of pupils which are in excess of 30 in each class having an enrollment of more than 30.

Schools that exceed either the average class size of 30 or the individual class size of 32 will lose a portion of the funding normally allocated to those students, as stated in subsections (c)-(e) of Section 41367 of the California Education Code (2006). I will express the statements in these subsections with the following equation,

$$\Delta S = 0.97 \times S_{ex} \frac{S_{CA,i}}{S_{CA,i-1}} \frac{S_{D,i-1}}{S_{D,i}} \tag{4.2}$$

where ΔS is the number of students deducted from the district's average daily attendance; a reduction in this number decreases the funding the district receives. S_x is the number of excess students in a district as calculated based on subsection 41367(a) above, S_{CA} and S_D are the enrollments in California and the district, respectively. S_i and S_{i-1} are the enrollments for the current year and previous year, respectively.

The California education code does not prescribe that schools should have individual classes containing 31 or 32 students, rather it describes a penalty for exceeding this number. I have created these soft ceiling rules to simulate how a district could maximize its spending ability for a given number of students.

I refer to the maximum number of students that can enroll in a class beyond the stated classroom-level target as the "softness," *m*, of the ceiling. We can interpret the soft ceiling in two ways. First, the soft ceiling can be applied using a hard ceiling at $n + m$

students plus enough teachers to push the average for the district below *n* students per class, and applying this rule both before and after the removal of students. This requires substantial coordination at the district level, so I will refer to it as the "district ceiling" rule. As with the hard ceiling rule, the district ceiling rule will be applied identically before and after the students are removed to participate in the school choice program; note that the hard ceiling rule is simply a soft ceiling rule where $m = 0$.

The second version of the soft ceiling rule, the "school ceiling" rule, does not require as much coordination at the district level. After the students decide to leave the district, schools might not, for whatever reason, make dramatic adjustments to the faculty unless enrollment falls to such a level that those adjustments could be made and still keep enrollment under the individual class size target, i.e., less than *n* students per class calculated at the school level. The number of teachers established before students are removed for the school choice program follows the same method as the district ceiling: the number of teachers is determined assuming a hard ceiling at $n + m$ students per class with enough additional teachers to keep the average class size below *n* students per class within the school. After the students are removed, however, the number of teachers is determined using a hard ceiling at *n* students. If the hard ceiling at *n* students requires more teachers than initially placed in any individual school, the original, smaller number will be used. Table 4.3 summarizes these rules.

Table 4.3

The rules used for deciding the number of teachers, T, in the simulation for a class size target of n and a softness parameter of m in a school with N_s students in the grade.

Rule	Initial Determination	Final Determination
Hard Ceiling	$T_i = \sum_s \left\lceil \frac{N_s}{n} \right\rceil$	$T_f = \sum_s \left\lceil \frac{N_s}{n} \right\rceil$
District Ceiling	$T_i = \sum_s \left\lceil \frac{N_s}{n+m} \right\rceil \bigvee \frac{\sum_s N_s}{n}$	$T_f = \sum_s \left\lceil \frac{N_s}{n+m} \right\rceil \bigvee \frac{\sum_s N_s}{n}$
School Ceiling	$T_i = \sum_s \left\lceil \frac{N_s}{n+m} \right\rceil \bigvee \frac{\sum_s N_s}{n}$	$T_f = \sum_s \left\lceil \frac{N_s}{n} \right\rceil \bigwedge T_i$

As asserted above, districts that operate in a fashion accurately modeled by the district ceiling rule would require more coordination than districts that operate according to the school ceiling rule. Suppose a small number of students leave a particular public school to participate in a non-district school choice program, and the decrease in enrollment would allow that school to reduce the number of teachers teaching second grade by one while keeping class sizes above *n* students, but below *n* + *m* students. The school could only enact this strategy if those making the decision could be certain that enough other classrooms in the district are sufficiently below the target class size to keep the district average class size below *n* students. If those making staffing decisions have no precise information about second grade class sizes in all schools in the district, the only way to certainly avoid the penalty associated with exceeding class sizes is to retain the full teaching faculty unless class sizes at that school would not exceed *n* students if a teaching position is eliminated. One can imagine that factors like the culture of intradistrict communication, reporting latency, building-level discretion, and the timing of contracts would all affect the level of coordination in a district; the specificity of these factors fall outside the scope of this research, so I will treat the district ceiling rule and the school ceiling rule as bookends on the actual behavior of districts.

The softness of the rule used to determine the number of teachers is an important determinant in the sensitivity of faculty size to a small reduction in enrollment. California law indicates we should expect that district officials used a softness of two (i.e., $m = 2$) when making staffing decisions for grades K-3. Logic would lead us to predict that softer ceilings would lead to a sensitivity nearer one using the district ceiling rule: large softness values would increase the likelihood that a district will be more constrained by the average class size of n students part of the rule than by the $n + m$ rule for individual classrooms. In a district where the number of teachers is only determined based on the average class size, then it does not matter which students leave the district to participate in a school choice program—only the number of students will matter. Put another way, if only the average class size matters, then removing s_R students from a district of s students means that $t_R = \lfloor s_R / n \rfloor$ teaching positions can be eliminated, regardless of exactly which students leave.

On the other hand, we would predict that a softer ceiling will lead to a sensitivity nearer zero for a school ceiling. Larger softness values would mean that a school with high capacity utilization (i.e., a school with nearly $n + m$ students in each classroom; see the A Variety of Districts section below for a thorough discussion of capacity utilization) with c classrooms would have to lose close to $c \times m$ students in order for class sizes in that school to fall to near n students, and then would have to lose n students beyond that before the faculty reductions would actually take place. Because many students would have to leave an efficiently utilized school before any changes in faculty size could take place, the sensitivity of such a district would be small for large values of m.

The simulation generally supports the logic described above for the school ceiling rule. For the district ceiling rule, however, the sensitivity increases with softness up to the legally allowed value of $m = 2$; after that it declines contrary to the logic described above. Because the data analyzed here are based on actual districts that make decisions based on a class size of 30 and a softness of two, analyzing these data with different softnesses (i.e., softnesses different than those actually used by the districts) may not produce useful insights in any event. To the extent, however, that the grade sizes have a somewhat random component, the different softness levels may indicate some patterns that would arise if the softness rules were different. See the A Variety of Districts section below for a discussion of the consistency of these findings. Figure 4.3 shows box plots representing the range of sensitivities measured in the Fallbrook Union Elementary district when 5% of students have been removed and the faculty sizes are calculated based on the district ceiling rule and the school ceiling rule, respectively. Note again that a softnes of zero equates to a hard ceiling.

Because these two soft ceiling rules create such different behavior in the simulation, I will conduct all of the following analysis with both rules. Whenever the analysis calls for presenting the findings in graphical form I will show plots using both rules, with the plot created with district ceiling rule on top and the plot created with the school ceiling rule on the bottom.

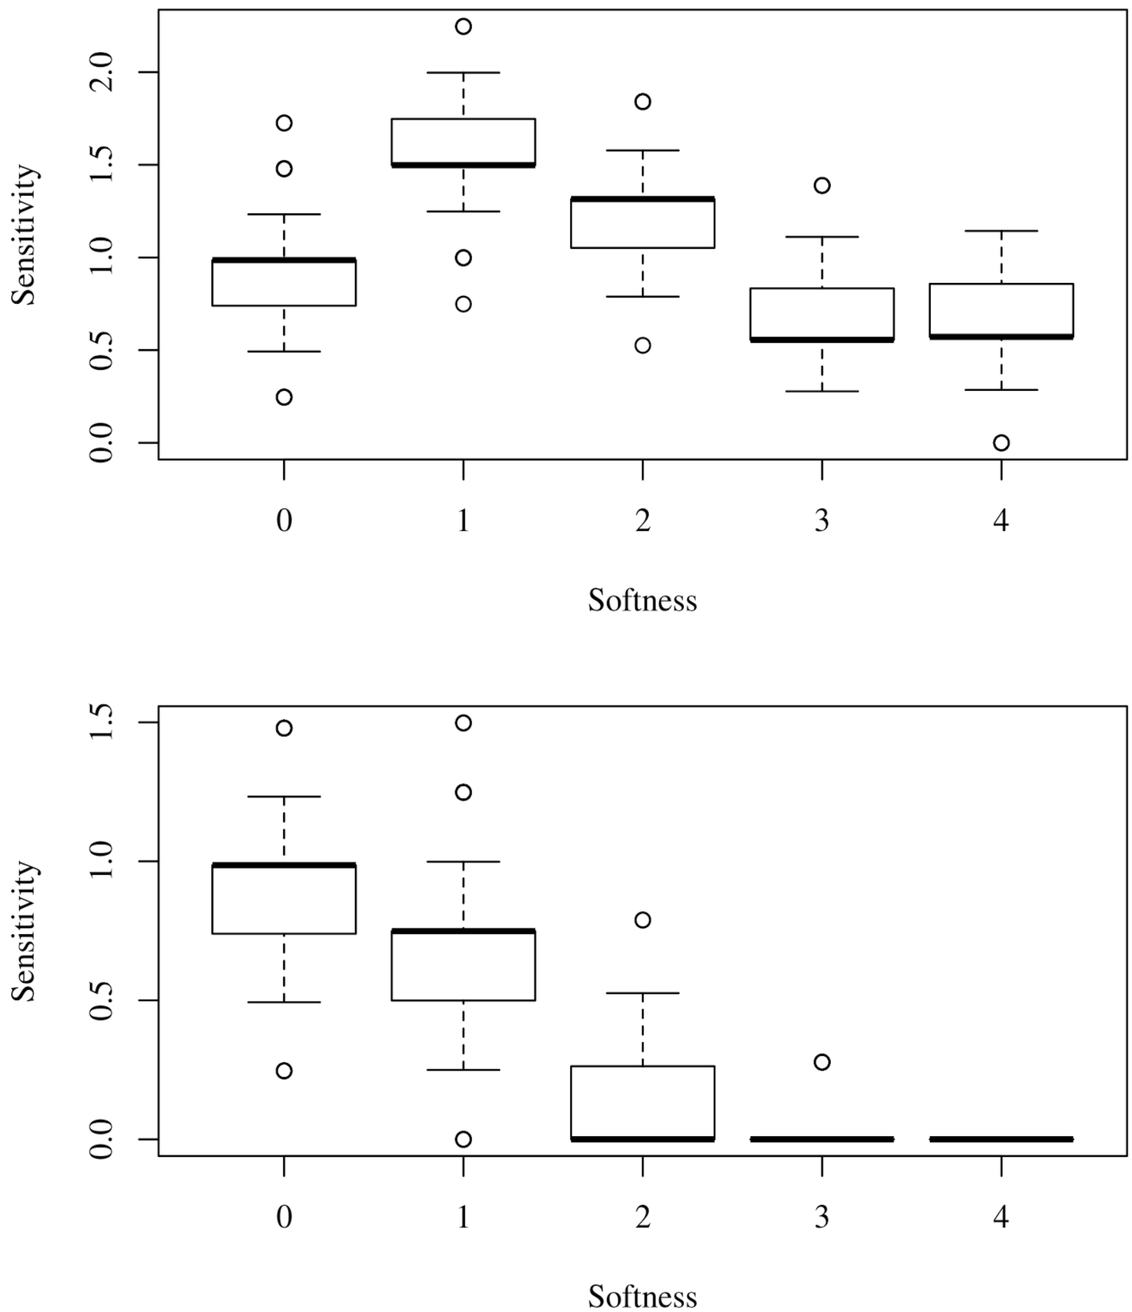

Figure 4.3. Box and whisker plots of the sensitivity of the Fallbrook Union Elementary school district with varying softnesses. The heavy line indicates the mean, the box encloses the interquartile range (IQR), and the dots represent observations beyond 1.5 times the IQR from the mean. The top plot shows the district ceiling; the bottom plot shows the school ceiling. Note that the case where $m = 0$, which equates to a hard ceiling, is the same for both rules.

School choice program parameters. ***Percent of students leaving the district.*** The central claim explored in this analysis is that faculty sizes cannot respond to small changes in enrollment, but what constitutes a "small" change in enrollment? We might guess that a district would have a better chance of eliminating teaching positions if 50% of students leave than if a single student—the smallest percentage of students possible—leaves. We would face a difficult task, though, if we tried to logically estimate with precision how removal percentages between 50% and one student would behave. Figure 4.4 shows the relationship between percentage of the student population removed and faculty/enrollment sensitivity. Surprisingly, though, for the district ceiling rule, some districts show a large faculty/enrollment sensitivity to very small removal percentages; this result is not indicative of a large pattern and will be further discussed in the A Variety of Districts section below.

Targeting student populations. Some school choice programs target students from certain demographic groups. The simulation can include this feature by giving students unequal probabilities of selection, or weights, when removing students for participation in the school choice program. Two methods of assigning these weights are compared in this analysis. The simplest method for removing students considers all students to be equally likely to participate in the school choice program—i.e., each student has a weight of 1. Some school choice programs are only available to low-income students; this can be represented in the simulation by giving students receiving free lunch a weight of 1 and all other students a weight of 0. The impact of only analyzing two methods of comparing weights is discussed in the Generality section of chapter five.

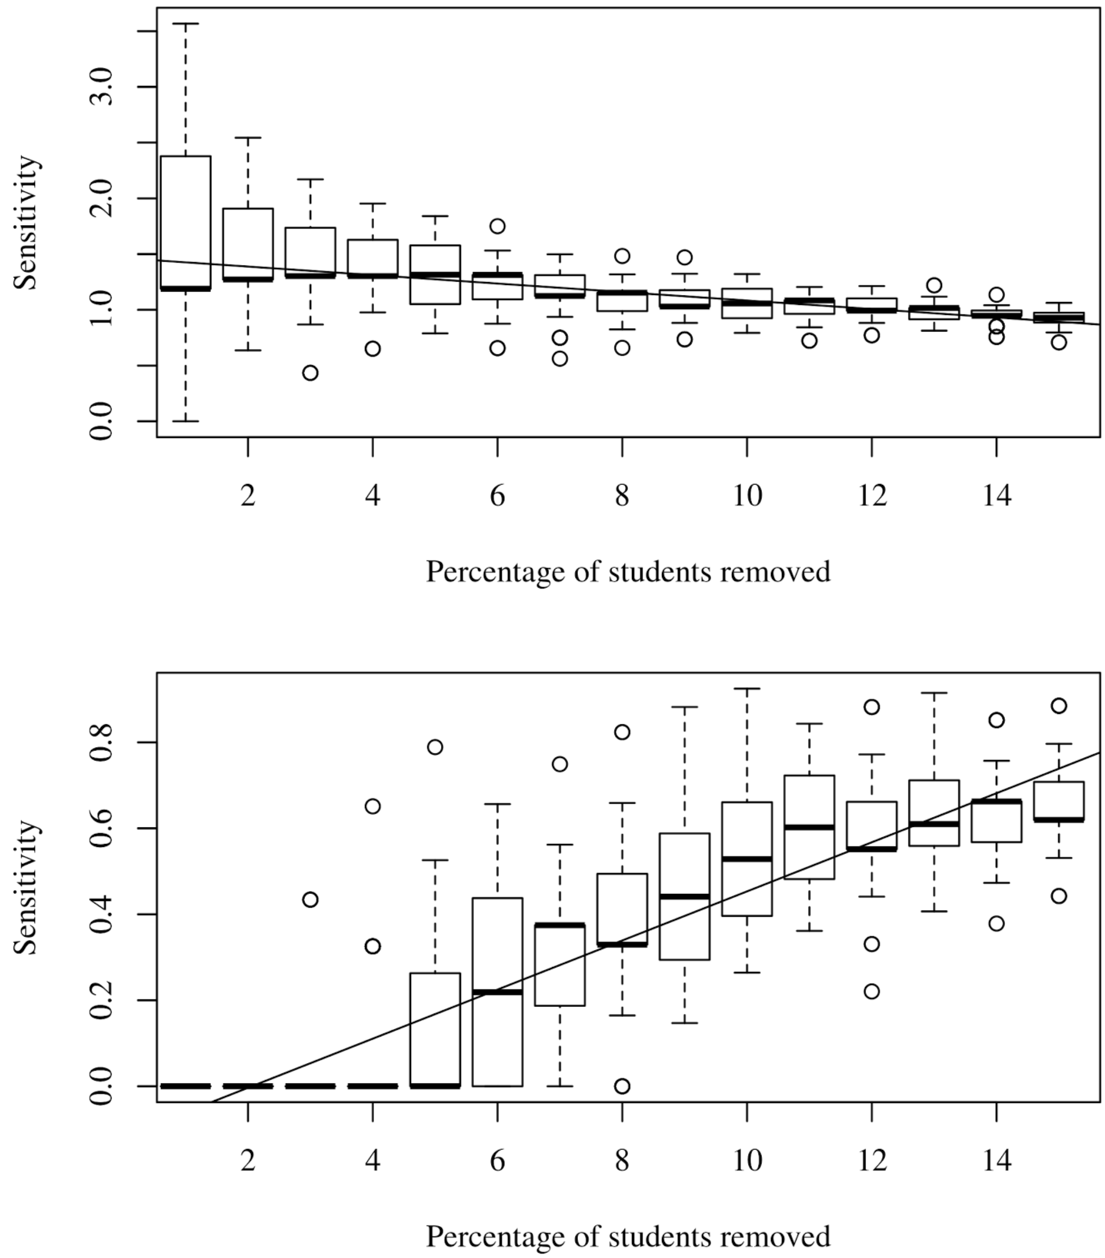

Figure 4.4. Box and whisker plots of the sensitivity of the Fallbrook Union Elementary school district with varying percentages of student removed for participation in a school choice program. The upper plot was created using the district ceiling rule, and the lower plot was created using the school ceiling rule. In both cases students were removed from grades one through three.

Free lunch students only. Insufficient accurate data exists in a publicly-available format at the grade level for the number of students eligible for free lunch programs to include free lunch eligibility as a parameter in the simulation with complete accuracy. Free lunch data only exists at the school level, so simulating a school choice program that only removes students from some grades requires the translation of school-level free lunch eligibility to grade-level free lunch eligibility. In this case I will assume that each grade in a school has the same percentage of students eligible for free lunch as the school as a whole.

As stated above, it is possible to program the simulation to only remove those students eligible for free lunch by setting the weight of those students not eligible for free lunch to zero. By setting weights in this way, once a school has been chosen in the two-step pseudo-random selection process—the likelihood of which will be based only on the number of students eligible for free lunch—those students not eligible for free lunch have a 0% chance of being selected. If a school choice program intends to remove a percentage of students that would be larger than the number of students eligible for free lunch, the simulation will simply remove all of the eligible students. In this case the percentage of students removed to participate in the non-district school choice program would be smaller than that intended by the school choice program. Figure 4.5 shows the difference between a school choice program that removes 5% of all students using both a district ceiling and a school ceiling, and the identical programs that only selects students eligible for free lunch.

Figure 4.5 shows that the probability of outcomes indeed depends on the population selected for the school choice program ($t(1982) = -5.244$, $p < .001$ for the district ceiling rule, and $t(1793) = 28.745$, $p < .001$ for the school ceiling rule). When the non-district school choice program only selects those students eligible for free lunch (the lighter lines in Figure 4.5), the lower-sensitivity outcomes occur less frequently under the school ceiling rule, but more frequently under the district ceiling rule. These results are not generalizable to other districts, especially in the case of the soft district rule; see A Variety of Districts for more discussion about these measurements. Even though the results from a particular district are not generalizable, I have included this detailed description in order to make the comparison of districts in the A Variety of Districts section more meaningful.

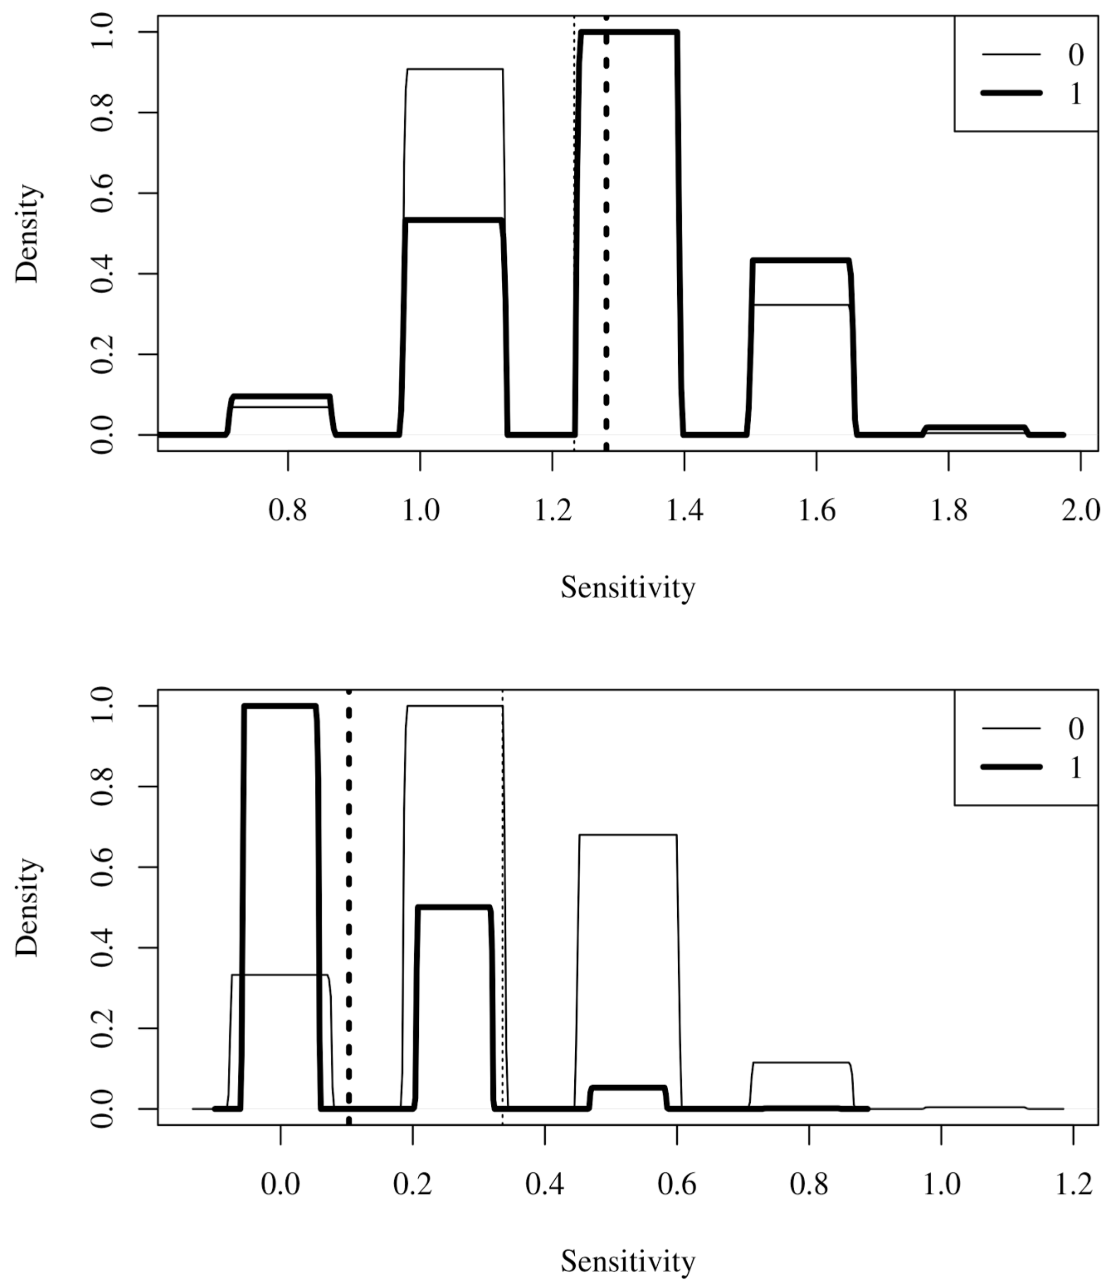

Figure 4.5. Histograms showing the frequencies of faculty/enrollment sensitivities while removing 5% of grades one through three in the Fallbrook Union Elementary school district. The dashed lines show the mean for each group of faculty/enrollment sensitivities. The top plot shows the district ceiling rule and the bottom plot shows the school ceiling rule. The heaviness of the lines show the weight of the students that are ineligible for free lunch.

Table 4.4

Enrollment and percent eligible for free lunch programs in the Fallbrook Union Elementary school district in grades one through three.

School	Total Enrollment[a]	% Eligible for Free Lunch[b]	Free Lunch Enrollment
San Onofre	307	20%	60
Iowa Street	138	5%	7
Frazier	279	59%	165
Fallbrook	315	72%	227
Live Oak	293	59%	171
Pendleton	412	14%	56
La Paloma	307	64%	198

[a]The total enrollment is for grades one through three only.

[b]The percentage of students eligible for free lunch is school wide, and may not be completely accurate for this subset of grades.

Understanding why only selecting subpopulations changes the sensitivity of a district's faculty size to small changes in enrollment requires looking at how the subpopulation is distributed in the district. As shown in Table 4.4, Live Oak Elementary has a smaller total enrollment than Pendleton Elementary, and a greater number of students eligible for free lunch. Because all students eligible for free lunch are equally likely to be chosen for participation in the school choice program, Live Oak Elementary is over three times more likely to lose a student than Pendleton Elementary. This occurs because Pendleton Elementary has relatively few student eligible for free lunch, despite its higher enrollment. Concentrating the students removed for participation tends to increase the number of students removed in a subset of schools in the district, which increases the likelihood that one of those school will lose enough students to decrease its faculty size.

Level 4: A Variety of Districts

All of the conclusions drawn above are based on individual districts in California, and should not be generalized beyond those particular districts. In this section, I will discuss how frequently the conclusions drawn thus far hold in the districts in California, and look for correlations between district characteristics and outcomes from this model.

I will look at five district characteristics in attempting to explain variation in sensitivity. First, I will look at district size, in terms of the number of schools. Second, I will look at the overall percentage of students eligible for free lunch. This characteristic will serve as a proxy for district affluence.

Third, I will look at a more complicated characteristic described by Heinesen (2004): capacity utilization effects (p. 443), or CU. I will define the CU of a grade in a school as the size of the least-full class (assuming all other classes are completely full) compared to the target class size. In many cases the CU is given by,

$$CU = \frac{N_s \bmod n}{n} \tag{4.3}$$

This equation holds except in the case where all the classes are full; when all the classes are full, N_s mod n equals zero, but this situation represents perfect use of capacity, so the CU must be manually reset to one. For example, the second grade at Fallbrook Street Elementary in the Fallbrook Union Elementary district had 106 students in 2011. For the purpose of defining the CU of this district I will assume that this grade has three full classes of 32 (assuming a class size target of 30 and a softness of 2) and one partially full class of ten students. This grade then has a CU of 31% (10 / 32 = 0.31) using a soft ceiling. If we do not allow any softness in the ceiling, then the same grade has three classes of 30 students and one class of 16 students, and a CU of 53%. The average CU of a district is calculated by summing the number of students in the partially full classes and dividing by the total capacity of those classes.

Fourth, I will look at a related measure, the weighted capacity non-utilization (WCNU), which measures the relationship between the percentage of students eligible for free lunch in a particular school to that school's CU. Finally, I will consider the minimum capacity utilization (MCU), which represents the lowest CU of all schools in the district.

For the purposes of this simulation, it does not matter if the second grade at Fallbrook Street Elementary actually has three classes of 32 and one class of ten, or two classes of 26 and two classes of 27. In either case the number of teachers required to staff the grade will reduce by one if ten students leave to participate in a non-district school choice program (assuming a district ceiling rule, or 16 students assuming a hard ceiling rule), but not if nine students leave. In terms of step costs as described previously, the CU measures the location of the grade on the step in terms of distance from the low end.

Capacity utilization differs importantly from mean class size in a district, although both measures show the relationship between actual class sizes and target class sizes. The 2011 class sizes for first grade and fifth grade at San Onofre Elementary in the Fallbrook Union Elementary district are 106 and 74 students, respectively. These grades have average class sizes of 26.5 students and 24.7 students, respectively, assuming a soft ceiling of 32. The first grade classrooms seem better occupied, which might matter in a strictly accounting-based discussion. In terms of allocation effects caused by school choice programs, though, we should view both grades as containing a few full classrooms with 32 students (three in first grade and two in fifth) and one class each with ten students. So these two grades with different average class sizes have identical CUs. This matters because each grade would have to lose the same absolute number of students—ten students in both cases—before a teaching position could be eliminated, even though ten students represents a different percentage of their enrollment.

Removing all students with equal probability. The following discussion will concern a non-district school choice program that removes 5% of students from grades one through three. We would not expect that either average free lunch eligibility or free lunch eligibility distribution would matter when free lunch eligibility plays no part in deciding which students to remove. Interestingly, though, the percentage of students eligible for free lunch does explain a significant part of the variation for both the district ceiling ($F(1,536) = 4.838$, $p = 0.028$) and the school ceiling ($F(1,536) = 15.13$, $p < .001$). In both cases the percentage of students eligible for free lunch is positively related to faculty/enrollment sensitivity, indicating that districts with lower average socioeconomic status can respond more sensitively than districts with higher average socioeconomic status. In neither case is the portion of the variation explained large, though ($R^2 = .007$ and $R^2 = 0.026$, respectively).

This result may seem surprising at first, but further examination offers both a statistical and contextual explanation for the existence of this relationship. Statistically, this relationship arises because of a weak, negative underlying relationship between percentage of students eligible for free lunch in a district and capacity utilization ($F(1,536) = 4.526$, $p = 0.034$); districts with a higher percentage of students eligible for free lunch tend to use their capacity less efficiently. Because of the relationship between

CU and faculty/enrollment sensitivity, the prediction that districts with lower socioeconomic status use their capacity less efficiently also leads to the prediction that those districts will have faculties that are more sensitive to losing a small number of students to a non-district school choice program. In fact, controlling for CU in the linear model before including percentage of students eligible for free lunch renders the relationship between faculty/enrollment sensitivity and the latter insignificant ($F(2, 536) = 2.516$, $p = 0.113$) when using the district ceiling rule. Percentage of students eligible for free lunch remains significant ($F(2, 536) = 10.514$, $p = .001$) even after controlling for CU using the school ceiling rule. Duncombe (2006) offers an explanation for why this situation might occur: more affluent districts also tend to have more involved voters, which produces a sort of "monitoring" (p. 19) effect that can pressure school boards and administrators to operate more efficiently. The observation here—that more affluent districts have higher CU and thus lower sensitivity—confirms this reasoning.

District size (as measured by the number of schools in the district) does not explain a significant portion of the variation in sensitivity ($F(1, 536) = 0.003$, $p = 0.958$ and $F(1, 536) = 3.611$, $p = 0.058$ using a district and school ceiling rule, respectively). District size does, however, explain where the variation occurs. Figure 4.6 shows that districts with fewer schools have a much greater variation in sensitivity; large districts tend to have a sensitivity much closer to the overall mean than small districts. The lines of best fit in Figure 4.6 depict linear models of the absolute magnitudes of the observations (i.e., the positive size of the distance from the mean) as predicted by the number of districts. The negative slope of this line indicates that variation in sensitivity decreases as the number of schools in the district increases for both the district ceiling rule ($F(1,536) = 17.119$, $p < .001$) and the school ceiling rule ($F(1,536) = 4.680$, $p = .031$). The figures stated above use all observations. If the largest district in California, Los Angeles Unified, is removed, the model explains much more of the variation: from $R^2 = 0.029$ to $R^2 = 0.100$ for the district ceiling rule and $R^2 = 0.007$ to $R^2 = 0.035$ for the school ceiling rule. Los Angeles might reasonably count as an outlier, as the number of schools it contains, 916, lies more than 21 standard deviations above the mean for California as a whole.

Capacity utilization contributes most to our understanding of the variation in faculty/enrollment sensitivity across districts. We would expect that districts with a high CU would have lower faculty sensitivities to the removal of a small number of students: high CU indicates that a school is far from the low end of the step, and needs to lose a large fraction of the step size before faculty reductions are possible. If many schools in a district have this same quality, the removed students would not likely be concentrated enough to allow enrollment to cross to the next step down. Figure 4.7 shows that the simulation confirms this logic. CU explains a significant portion of the variation in sensitivity in simulations that use both the district ceiling rule ($F(1,536) = 65.98$, $R^2 = 0.108$, $p < .001$) and the school ceiling rule ($F(1,536) = 204.8$, $R^2 = 0.275$, $p < .001$). Additionally, the coefficients in each case are negative, indicating that districts with a higher CU tend to respond to the removal of students less sensitively.

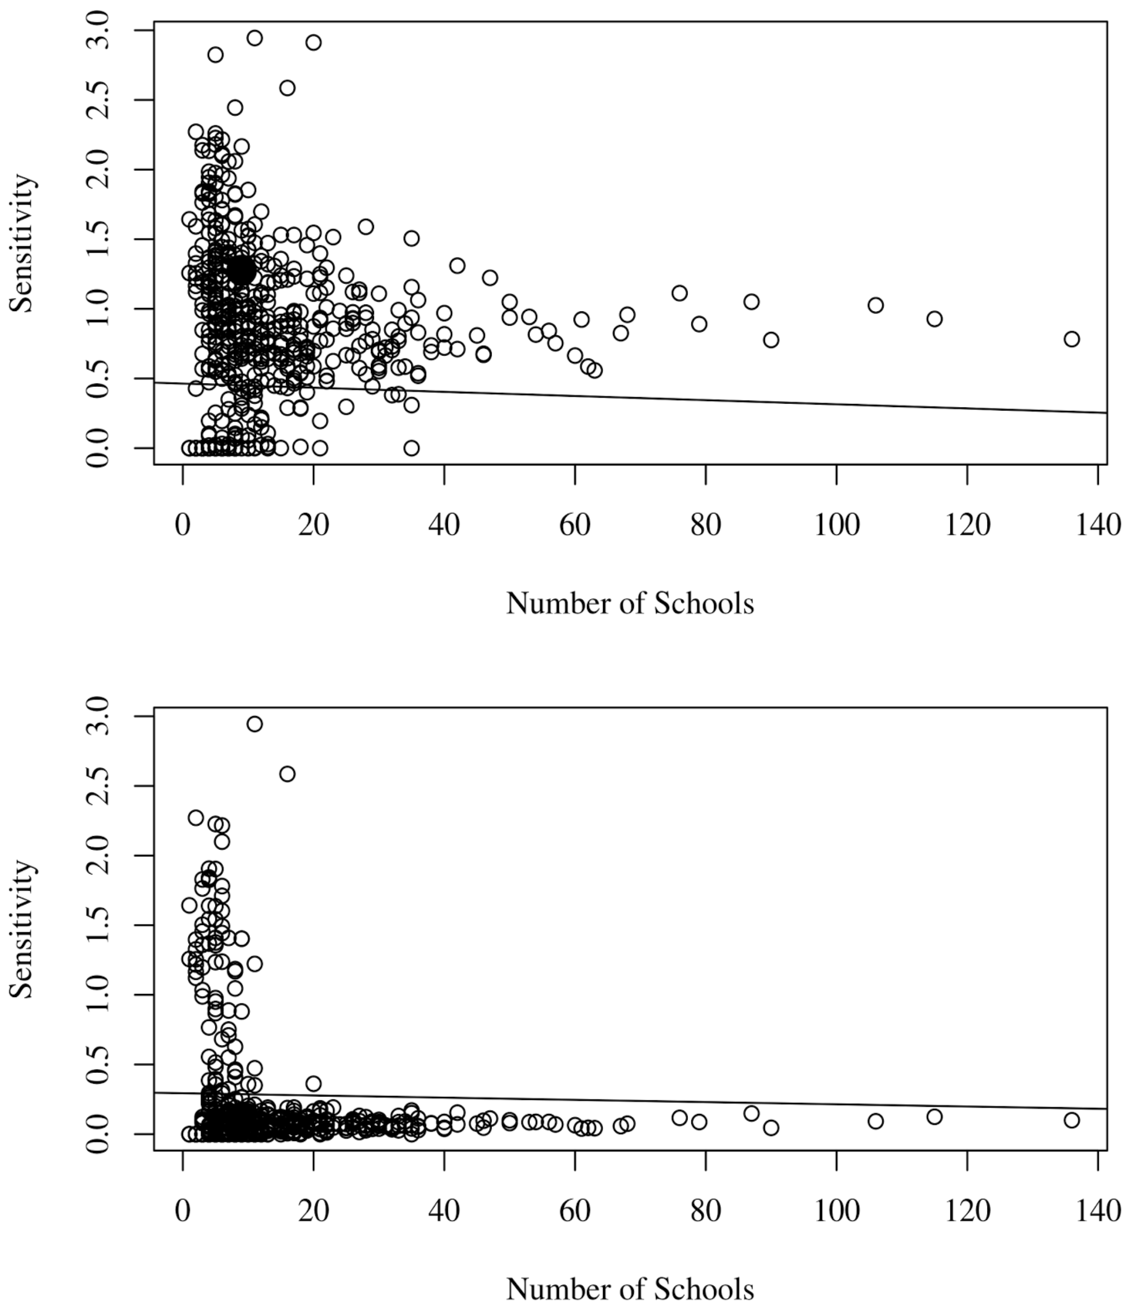

Figure 4.6. The relationship between number of schools in a district and sensitivity for both a district ceiling rule (top) and a school ceiling rule (bottom) when 5% of students leave from grades one through three. The linear model shows positive distance from the mean modeled by number of districts. To increase readability, the scales of these plots are reduced to omit the two largest districts in California, though they were present for the calculations.

Free lunch students only. Like the simulation that removes all students with equal probability, the faculty/enrollment sensitivity of a simulation that only removes students eligible for free lunch depends heavily on capacity utilization in the district. It also depends on how the students eligible for free lunch are distributed, particularly, whether the students eligible for free lunch are concentrated in schools with low CU. To explore this relationship, I will create a measure of capacity non-utilization (i.e., the distance from the high end of the step instead of the low end) that is weighted by the proportion of students eligible for free lunch in the district that can be found in that grade, or WCNU.

$$\overline{CU_F} = (1 - CU)\frac{F_s}{\sum_s F} \tag{4.4}$$

When choosing only students eligible for free lunch, a school is most likely to eliminate a teaching position when the CU is low (i.e., it is close to the low end of the step) and when the school has a large fraction of the students eligible for free lunch, as this increases the expected number of students who will leave that school. I use the distance from the high end of the step, *1 - CU*, so that both terms that are multiplied together will have larger values when the probability of eliminating a teaching position is high. Figure 4.8 shows plots of the mean faculty/enrollment sensitivity modeled by

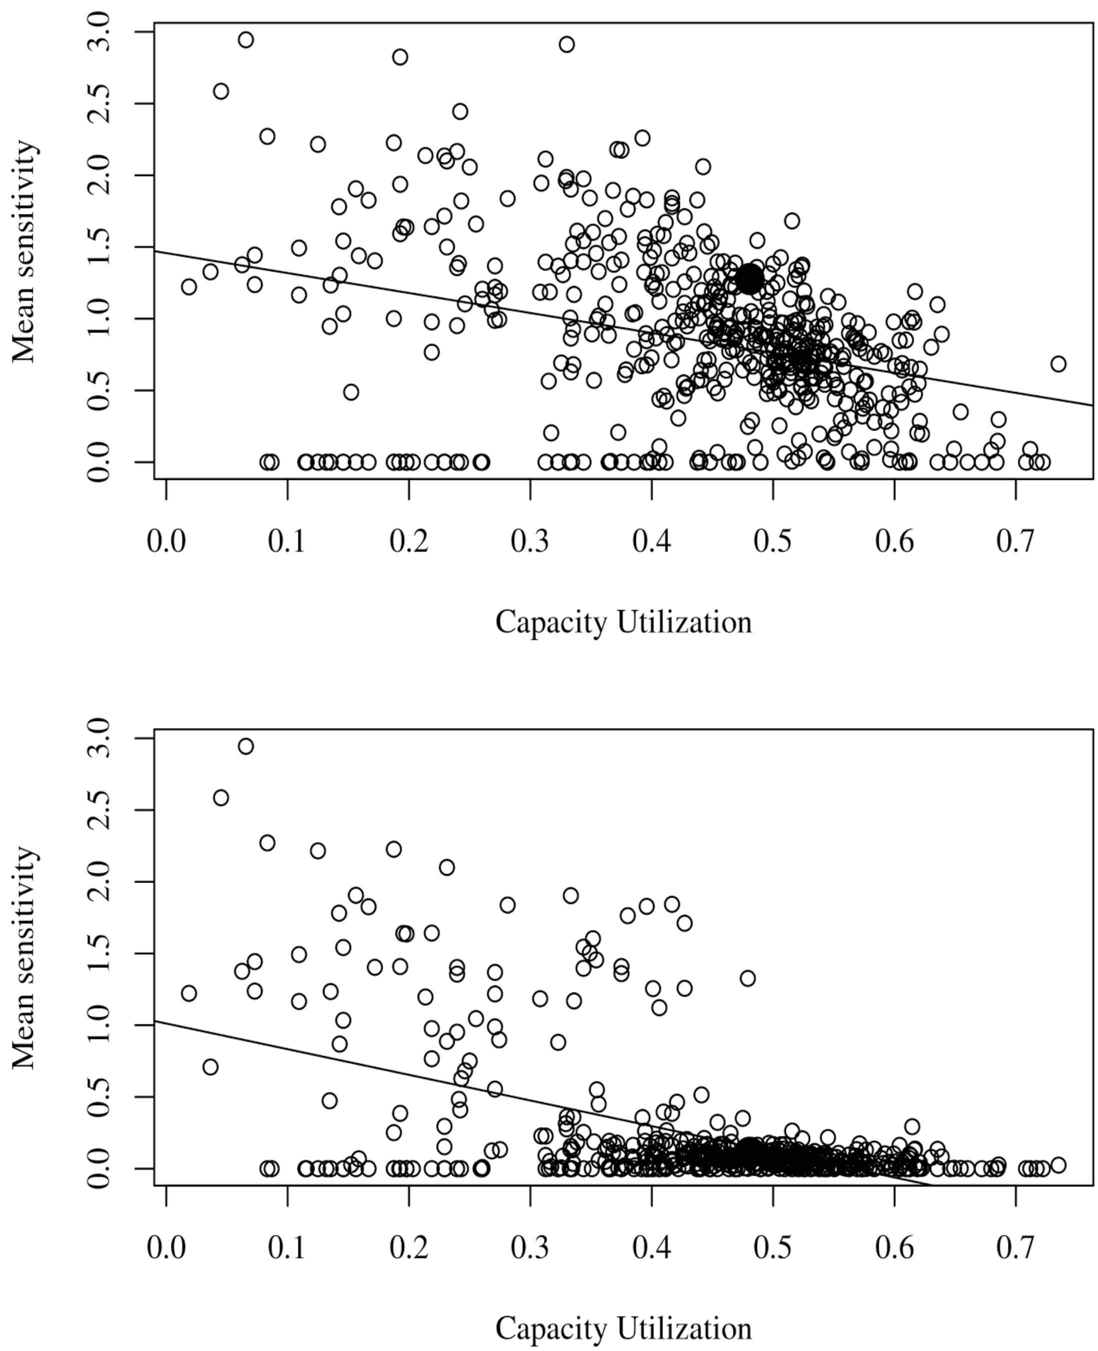

Figure 4.7. Faculty/enrollment sensitivity to removing 5% of students in grades one through three modeled by CU in school districts in California using a district ceiling rule (top) and a school ceiling rule (bottom). The large dot represents the Fallbrook Union Elementary school district.

WCNU from 500 runs of the simulation using a district ceiling, removing 5% of students from grades one through three, and choosing only students eligible for free lunch. Table 4.5 describes the relationships between WCNU and faculty/enrollment sensitivity.

Table 4.5
A summary of linear models that describe the relationships between faculty/enrollment sensitivity and weighted capacity non-utilization.

Rule	Included Data	F	R^2	*p*
District ceiling	All data	55.62	0.108	< .001
District ceiling	Significant data	184.6	0.441	< .001
School ceiling	All data	25.11	0.051	< .001
School ceiling	Significant data	54.37	0.203	< .001

One of the two trend lines on Figure 4.8 depicts linear models for all of the data; the other trend line shows a model that only includes those observations for which the mean sensitivity calculated when selecting only students eligible for free lunch differs from the mean sensitivity found when selecting all students with equal probability; significance is defined at an $\alpha = 0.05$ level using the Bonferroni-Holm correction for multiple comparisons.

The model that uses all data would be most helpful for policy makers interested in predicting the economic impact of a proposed school choice program, because the model does not rely on our ability to predict which districts will have a significant difference in mean sensitivity between school choice programs that target all students or only those eligible for free lunch. Even though this finding may have little direct application in policy discussions, it does contribute significantly to our understanding of how removing a small number of students affects the number of teaching positions in a district: *When*

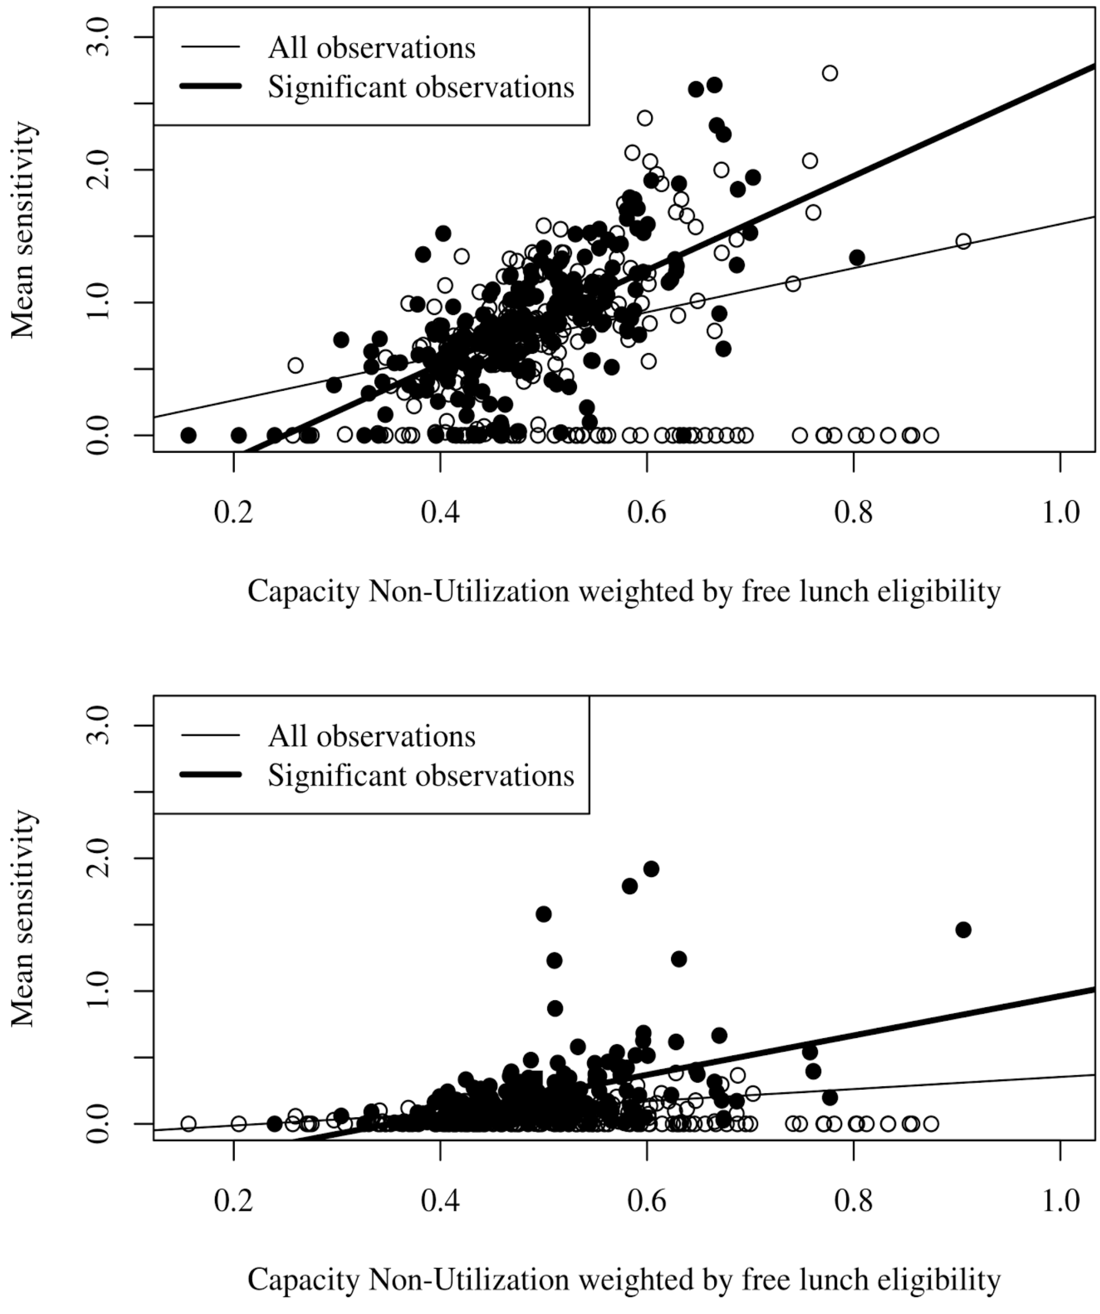

Figure 4.8. Faculty/enrollment sensitivity to removing 5% of students, targeted at students eligible for free lunch in grades one through three modeled by WCNU in school districts in California using a district ceiling rule (top) and a school ceiling rule (bottom).

targeting students eligible for free lunch matters, the degree of concentration of students eligible for free lunch in low-CU schools affects the faculty/enrollment sensitivity a great deal.

Different "small numbers" of students. Figure 4.9 shows that the percentage of students removed from the Fallbrook Union Elementary district significantly impacts the sensitivity of the faculty size to changes in enrollment. In the case of the Fallbrook Union Elementary district, the smallest percentages (i.e., those percentages below 3%) resulted in sensitivities near zero using the school ceiling rule and above one using the district ceiling rule. As the percentage of students removed increased, the sensitivity increased sharply using the school ceiling rule and decreased somewhat, though significantly, using the district ceiling rule. As we have seen in this section, though, districts can behave very differently in the simulation, and that finding holds for the percentage of students removed as well.

To examine the consistency of the influence of percentage of students removed in the simulation on sensitivity I created a simulation to run 50 trials of removing a percentage of students and determining faculty size using both the district and school ceiling rules. The percentages of students removed were integer percents between 1% and 15%, inclusive. This simulation then ran over all districts in California with more than one elementary school. For each district, I recorded the slope and intercept of the line of best fit. Figure 4.9 shows a pseudo-random selection of 30 of these districts. Plotting all 537 lines makes the figure very hard to read, but the selection of lines shown in Figure 4.9 is representative of the group of lines as a whole; the slopes of the lines in the selection do not meaningfully differ from the slopes of the whole group ($t(31.69) = 0.564$, $p = 0.577$, and $t(34.08) = -0.302$, $p = 0.765$ for district and school ceiling rules, respectively). Figure 4.9 shows that the impact of percentage of students removed on faculty/enrollment sensitivity is more consistent for the school ceiling rule than for the district ceiling rule. The collection of slopes for the school ceiling rule is significantly different from being randomly distributed around zero ($t(536) = 10.83$, $p < .001$); the 95% confidence interval [2.579, 3.722] shows that the trend in the slopes is to be significantly larger than zero. This indicates that as the percentage of students removed as part of a school choice program decreases, so does the sensitivity of faculty size.

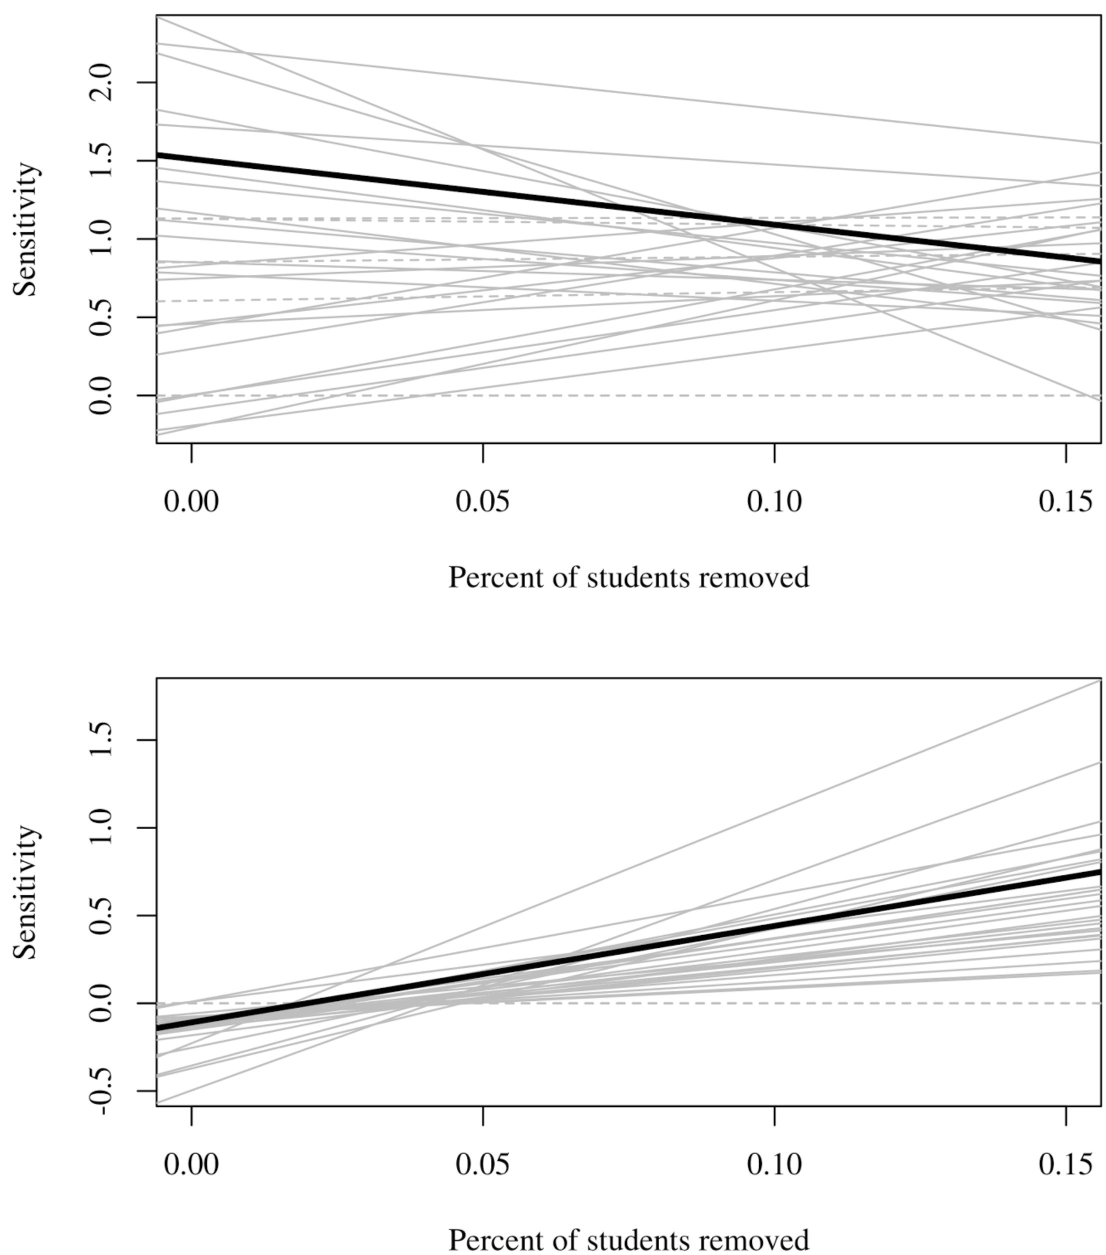

Figure 4.9. Lines of best fit based on simulations that remove between 1% and 15% of students from 30 districts in California. The top plot was created using the district ceiling rule, and the bottom plot was created using the school ceiling rule. For reference, the heavier line is based on the Fallbrook Union Elementary district. The dashed lines represent lines of best fit that are not significant at an $\alpha = 0.05$ level.

The slopes found using the district ceiling rule, on the other hand, tell a different story. The 95% confidence interval of slopes [-0.840, 0.437] does contain zero, indicating that we cannot state that the impact of removing different percentages of students is different than zero ($t(534) = -0.621$, $p = 0.536$). This does not occur, however, because percentage of students removed tends not to influence the sensitivity in different districts, but because the magnitude (i.e., steepness) and direction (i.e., positive or negative) of the influences occur with fairly uniform frequency. Even more importantly, a strong, negative relationship ($R^2 = 0.895$, $F(1,533) = 4542$, $p < 0.001$) exists between the slope and the intercept of the lines of best fit generated using the district ceiling rule. This indicates that the variation in slope is not due to a variation in the faculty/enrollment sensitivities at high percentages of students removed, but to a variation in sensitivities at very low percentages of students removed. That is, districts with a negative relationship between percentage of students removed and faculty/enrollment sensitivity tend to have that relationship because they have a high faculty/enrollment sensitivity at very low percentages of students removed, and not because they have a particularly high sensitivity at large percentages of students removed. In short, without knowing anything about a district, we can state that if the way of deciding on faculty size is best modeled by the school ceiling rule, then sensitivity will decrease as percentage of students removed decreases. If faculty size is best modeled by the district ceiling rule, then sensitivity will become more dramatic and less predictable as the percentage of students removed decreases.

We can use hierarchical linear modeling (HLM) to relate district parameters to the magnitude and direction of the relationship between percentage of students removed and faculty size sensitivity. Two district characteristics account for a significant part of the variation in slopes of the lines of best fit created by simulating different percentages of students removed in California school districts. Not surprisingly, CU is one and a related measure, minimum capacity utilization (MCU) is the other. Where CU represents the average capacity utilization of all schools in a district, MCU represents the lowest capacity utilization of all schools in a district. For example, if the third grade in a ten-school district has nine schools where the smallest classes each contain 29 students (with a maximum class size of 30) and one school where the smallest class contains 2 students, the CU for that grade would be 0.877, but the MCU would be 0.067. As it turns out, when removing a very small number of students, the MCU can be more important than the CU. Table 4.6 shows an ANOVA of the slopes of lines of best fit created by simulating the removal of varying percentages of students. Note that of the 537 districts with multiple elementary schools, the linear models discussed here only use the 458 districts that do not have to add extra teachers after initially placing teachers using the district or ceiling rule—recall that those two rules are identical before students are removed by the simulation. In 79 of the districts, the average class size in at least one grade would be greater than the target district class size, 30 in this case. I excluded these districts because they behave similar to a hard ceiling at 30 students per class instead of a soft ceiling at 32 students; in a real and meaningful sense, these districts are different than the other 458 and should not be included when attempting to discern the behavior of a school district during a school choice program.

Table 4.6

ANOVA of the slopes of lines of best fit created by removing between 1% and 15% of students from each district in California using both the school ceiling rule and the *district ceiling rule.*

Rule	Parameter	Coefficient	*p*	R^2
District Ceiling				0.152
	Intercept	5.887	< .001	
	CU	-13.30	< .001	
	MCU	20.49	< .001	
School Ceiling				0.655
	Intercept	14.78	< .001	
	CU	-18.93	< .001	
	MCU	-4.846	< .001	

[a]The word "intercept" here refers to the expected value of the slope when CU = 0 and MCU = 0, and not, for example, the location of where the lines of best fit on Figure 4.9 cross the y-axis.

Note. These linear models were created using only districts that do not have to consider the district average class size when placing teachers.

The direction (i.e., positive or negative) of the coefficients suggest interesting behavior for the district and school ceiling rules. For the district ceiling rule, both the expected slope[2] when CU = 0 and MCU = 0 and the coefficient for MCU are positive, while the coefficient for CU is negative. When the least efficiently utilized school in a district is particularly poorly utilized (i.e., a small MCU), the slope will tend to be smaller. This makes sense because of the strong, negative relationship between slope and intercept for the lines of best fit found using the soft district rule: a smaller MCU means a smaller slope, which means a larger intercept, and thus a greater sensitivity to very small enrollment changes. In other words, districts that contain a poorly utilized school can be overly sensitive to very small enrollment changes.

To understand the HLM coefficients for the school ceiling rule, we have to look at the other end of the range of percentages. When using the very smallest percentages of

2 The "expected slope" is called the "intercept" in Table 4.6, but I will avoid that term here because I use "intercept" in a different capacity in this section, as well.

students removed, the sensitivities calculated using the school ceiling rule are typically zero with a few outliers at the next smallest value that represents decreasing the faculty size by one teacher. Because a sensitivity of zero is the absolute minimum sensitivity in this simulation—which is to say that a district will never hire teachers after losing students to a non-district school choice program—variation in slope occurs because of differences in sensitivity at the high-percentage end of the range. The expected slope when CU = 0 and MCU = 0 for the school ceiling rule is positive, and the coefficients for CU and MCU are both negative. Because of the absolute minimum sensitivity of zero at the low percentages, we can interpret the negative coefficients for the slopes as meaning that high-CU and high-MCU districts have smaller sensitivities for larger percentages of students removed, which confirms the logic described elsewhere in this study.

Monetary Accounting Model

Faculty/enrollment sensitivity, although interesting, serves as an intermediate measure between the behavior of the schools during a non-district school choice program and what advocates and decision makers actually disagree about: money[3]. The simulation described above can model district finances by viewing students as the source of district funding and teachers as comprising one of four categories in a district's budget. As a way of exploring the role of allocation effects on district finances, I will create a financial model of a district by assuming that students are the only source of funding and teachers represent one of the district's significant expenditures. This monetary accounting model differs from a cost accounting model (Levin & McEwan, 2001) in that the monetary accounting model described here aims to show the budget of the school district, and not a

3 All of the monetary amounts in this section are expressed in 2011 dollars.

larger gamut of monetary and non-monetary costs. The cost and funding categories that I will work with are shown in Table 4.7.

Table 4.7
Funding and expenditure in the financial model of the district.

Category	Funding/Expenditure	Behavior
Funding		
Per pupil expenditure	$9,139	Proportional to enrollment
Expenditure		
Plant and administration	p% of budget	Fixed
Instructional faculty[a]	$63,903 per teacher	Proportional to faculty size
Other faculty	Varies	Inversely related to faculty size
Non-instructional services	Varies	Inversely related to other spending

[a]This is the mean faculty salary in California from 2011. In some trials of the simulation described below, this average salary was used to generate salaries for individual teachers. The particulars of creating faculty salaries for use in this simulation is described below.

The financial model employed in this inquiry assumes that all of the money that provides for a district's spending is determined by the number of students that attend school in that district. Each student enrolled brings the district the same amount of money, which I refer to as the per pupil expenditure (PPE). The funding that comprises the PPE might come from various sources, like the city, state, or federal government, local levies, district power equalization or block grants, etc. The nature of funding sources may, in fact, play an important role—as shown by Nechyba (1999)—but that behavior falls outside the scope of this study. I will consider the funding for education in California as a constant dollar amount allocated per pupil. In 2011, the PPE in California was $9,139; this differs from the per pupil revenue, which was $11,048 (Holeywell, 2013). Spending differs from revenue in that revenue includes factors such as "capital outlays and other costs" (Holeywell, 2013, p. 1). Spending only includes those budgetary components that go directly to the educational service of the students. The decision to use spending versus revenue will not impact the absolute financial calculation of the impact of a school choice program, because the $1,909 per pupil counted as revenue but not included as spending would count as part of the fixed cost in this monetary accounting model, but it will matter when the financial calculation is expressed as a portion of the PPE. Using spending to determine the size of the budget will show the higher allocation cost (or benefit) of a school choice program as a portion of PPE: if the simulation finds that each participant creates an allocation cost of $200, that allocation cost will be a larger percent of the spending (2.2% of $9,139) than the revenue (1.8% of $11,048). I will use spending as the measure of the size of a district's budget because using spending makes the percentage of the district budget allocated for faculty salaries closer to the stated average of approximately 60% (Holeywell, 2013)—the portion of the district budget dedicated to instructional costs calculated here is approximately 32% when using spending, and approximately 26% when using revenue. A possible explanation for the difference in average reported faculty expenditure and the faculty expenditure calculated here is discussed below.

Most voices on both sides of the school choice debate agree that the operation of a school includes some fixed costs, or at least costs with a sufficiently large step size to behave as though they are fixed for the purpose of this simulation. Costs that fit into this category include building maintenance, some utilities, and some personnel costs. Small changes in enrollment would not allow a district to close and sell entire schools, so this category of costs will not change under any circumstances that will arise here. The size of a district's fixed costs will vary, so, instead of trying to determine an appropriate percentage of the budget dedicated to fixed costs that will apply to all districts, I will select the portion of a district's budget that is fixed from the range [10%, 80%] of non-faculty spending; see below for a discussion of how the size of fixed costs affects the financial outcomes of this simulation. In practice this works out to between 4% and 32% of the district's total budget.

The most important cost category in this study is instructional faculty salaries. I will determine the cost of a teaching faculty in a district in two ways. First, using the average salary—$63,903 in California in 2011—and multiplying by the number of teachers in the district and a benefit ratio. The benefit ratio (always assumed to be 1.25 in this simulation) accounts for the non-salary aspects of teacher compensation, such as retirement contribution and health benefits. The second method involves using the starting and maximum salary in California in 2011—$41,372 and $81,573, respectively — to create a hypothetical list of salaries for individual teachers. To do this I will center a normal distribution between the ends of the salary scale, at $61,472, and the set the standard deviation at one fourth of the range, or $10,050, and then pseudo-randomly generate a list of salaries from this distribution. If any salaries generated in this way fall below the minimum or above the maximum, the algorithm will reset them to the nearest boundary. This method allows for the determination of the change in total faculty cost if the teachers cut from the district are chosen by seniority (i.e., removing the newest teachers with the lowest salaries) or by retirement (by letting the most experienced teachers with the highest salaries retire). Note that the center of the range of salaries lies below the average salary in California. To account for this, I will multiply the cost of faculty by a constant determined before students are counterfactually removed from the district and the size of the faculty changes. After the students are removed, I will multiply the calculated faculty cost by this same constant. See Table 4.8 for an equation that describes these methods. The fact that the center of the salary range and the mean faculty salary differ indicates that the distribution of salaries is not normal. Assuming a normal distribution, while likely inaccurate, does still allow this simulation to select teachers from the high and low end of the salary range, and so still accomplishes the goal of creating the distribution.

Table 4.8

Equations that show how faculty cost is determined from aggregate salary information and the number of teachers.

Rule	Equation
Mean salary rule	$C_{faculty} = \overline{S} \times T \times r_{benefit}$
Salary distribution rule	$C_{faculty} = \sum_{T}(U(\frac{S_{max} + S_{min}}{2}, \frac{S_{max} - S_{min}}{4})) \times r_{benefit}$

Note. $C_{faculty}$ is the cost of the district's faculty, $\hat{S}$ is the mean faculty salary, T is the number of teachers in a district, $r_{benefit}$ is the benefit ratio, $U(m,sd)$ is an element chosen from the uniform distribution centered at m with a standard deviation of sd, S_{min} and S_{max} are the starting and maximum salaries, respectively.

The other faculty category of costs is added to make the faculty spending calculated in this model equal to 60%, the stated average percentage of a district's budget that goes to faculty spending (Holeywell,). The reasons that the calculated faculty spending in this model do not sum to 60% of the calculated budget are numerous, but special education likely accounts for a substantial portion. Special education classes are often smaller than other classes in a school, and some students receiving special education services attend mainstream classes but have a paraprofessional to help ensure they receive an appropriate educational experience. Neither of these factors are accounted for in this model, and both would make actual faculty spending higher than predicted by this model. Accounting for the role of special education in the selection of students and in creating the financial model of the districts will be an important area of future research.

I calculate the cost of the other faculty before removing students for the non-district school choice program, and then leave that cost when determining the financial model afterward. Because this faculty spending is not suggested by the enrollment in the district, I assume that changes in enrollment will not affect this category of the budget.

The final line in the model of a district's budget described here is non-instructional services. I assume that all of a district's money not spent on fixed operating costs or faculty goes to offer students a rich education experience; this represents the full amount of money available to provide services to students. These services take myriad forms: art supplies and musical instruments, equipment for athletics, busing and admission for field trips, just to name a few. Dividing this amount of money by the number of students in the district yields the average amount of money spent on each student, or Services Per Student (SPS). The interplay among enrollment, faculty size, and the percentage of the budget dedicated to fixed costs will likely change the SPS in a district as a result of a non-district school choice program.

The ultimate financial measure that I will calculate is the allocation cost per participant, ACP. The ACP is the reduction in SPS summed over the students who remain in the district during the non-district school choice program, divided among the students who participate in the non-district school choice program, then expressed as a portion of the loss in spending of the students who participate, or

$$ACP = \frac{(SPS_i - SPS_f) \times S_f}{(S_i - S_f) \times PPE} \qquad (4.5)$$

where SPS_i and SPS_f are the calculated SPS values before and after removing students from the district, respectively; and S_i and S_f are the numbers of students in the district before and after removing students from the district, respectively. For example, an ACP of 0.4 indicates that 40% of the PPE for each student that participates in the program should be counted as a cost to the students who remain in the district.

I chose the ACP as the ultimate financial measure in this study for two reasons. First, this number will help policy makers arrive at a monetary value to include in a cost-effectiveness study, and additionally, inflation will not affect this number. Assuming that the relationship between average faculty salary and per pupil expenditure remains roughly similar over time, the ACP predicted for a given school choice program in a given district should not change. Second, some school choice advocates recommend creating vouchers for less than the per pupil expenditure as a way of allowing public schools to continue to cover their costs, thus making voucher programs more palatable. Calculating the ACP allows policy makers to better understand how much of a voucher could be offered while not harming a public schools ability to cover its fixed (and potentially step) costs. Policy recommendations of this sort are discussed in more detail in chapter five.

The influence of the percentage of students removed for participation. The results described in the A Variety of Districts section above show that the percentage of students removed from a district for participation in a school choice program strongly affects the faculty/enrollment sensitivity. Using a school ceiling rule, the faculty/enrollment sensitivity remains at zero up to approximately 4% of students removed and then increases. Using the district ceiling rule, the faculty/enrollment sensitivity varies enormously for removal percentages near zero, and then converges to the mean as the removal percentage increases. Because we expect a relationship between faculty/enrollment sensitivity and allocation cost, running the accounting model with various percentages of students removed should show that the cost depends on the number of students removed. But because the number of students removed will also affect the magnitude of the change in the amount of money available for non-instructional services compared to fixed costs, predicting the relationship between percentage of students removed and allocation cost using logic alone will be challenging, if not impossible.

To explore this relationship, I ran the simulation three times with 100 trials each on every district in California with students in grades one through three; the percentage of students removed in each of the three runs were 2%, 5%, and 8% of students. Figure 4.10 shows kernel density plots for each removal percentage using both the school ceiling rule and the district ceiling rule. The percentage of the budget dedicated to fixed costs in both cases is 25% and the faculty spending was calculated in all trials using the mean teacher salary. Table 4.9 shows a summary of these trials.

Because of the large number of observations—100 trials for each of the 536 districts with more than one elementary school—the difference in ACP values based on percentage of students removed is significant ($F(1,160498) = 121.9$, $p < .001$, and $F(1,160498) = 813.4$, $p < .001$, for the district and school ceiling respectively), but in neither case is the impact of percentage of students removed very important ($R^2 < .001$, and $R^2 = .005$, respectively). We can conclude that other factors have more influence on ACP than percentage of students removed. A relationship exists, statistically speaking, but policy makers can safely ignore that relationship when considering parameters of a school choice program.

Table 4.9
Calculated ACP for all districts in California using the average teacher salary and assuming that 25% of a district's budget goes to fixed costs.

Teacher assignment rule	% students removed	Mean ACP	% less than zero
District Ceiling	2%	0.555	8.79%
	5%	0.571	1.76%
	8%	0.573	0.51%
School Ceiling	2%	0.767	4.88%
	5%	0.775	0.56%
	8%	0.730	0.06%

The influence of the size of fixed costs. Figure 4.11 shows box and whisker plots for the ACP calculated during 100 trials of all districts in California with elementary schools. As stated above, the fixed costs used in these calculation range from 10% of non-instructional spending (4% of the total budget) to 80% of non-instructional spending (32% of the total budget). All trials in this section use the mean teacher salary when calculating faculty spending. This analysis shows that decision makers should predict a higher ACP when the portion of the district's budget that is fixed is large. This behavior

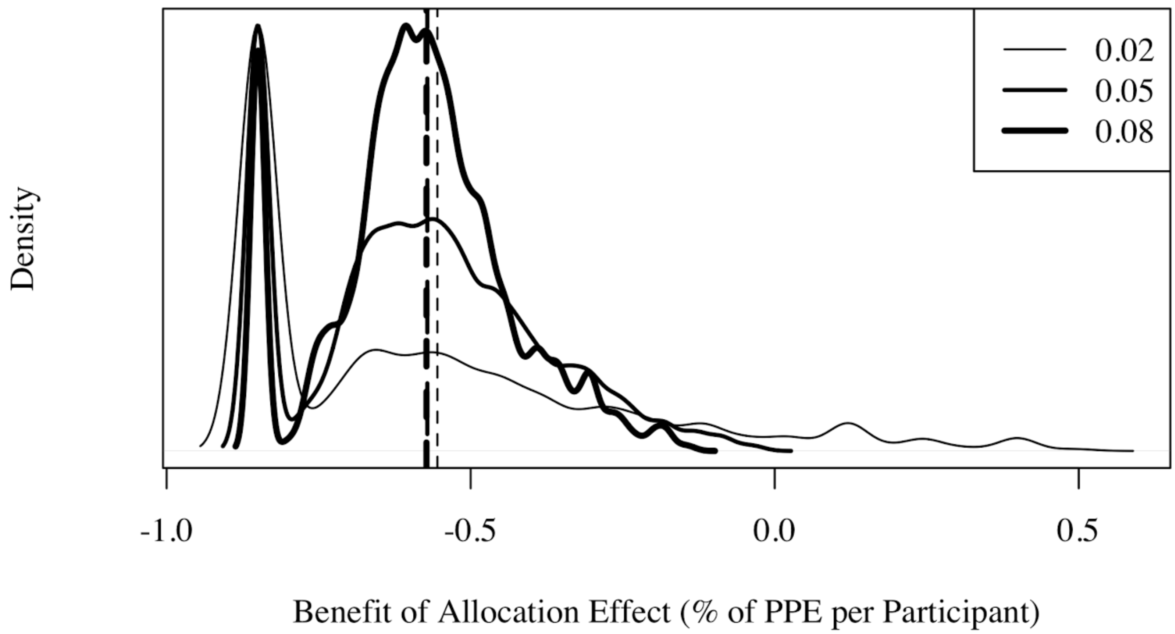

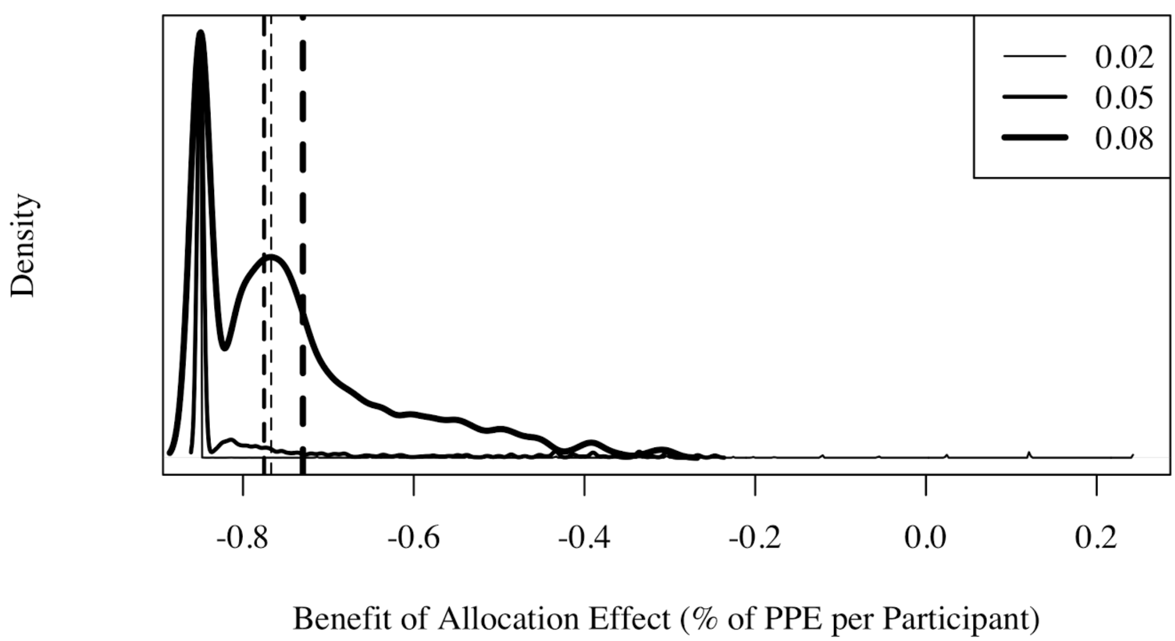

Figure 4.10. Kernel density plots showing the frequency of allocation effect per participant values in all districts in California for three different percentages of students removed for participation in school choice programs. The top plot uses the district ceiling rule, the bottom plot uses the school ceiling rule, and both plots assume that 25% of a district's budget goes to fixed costs.

occurs because the size of the fixed portion of the budget is inversely related to the level of resources available to provide non-instructional services to students. When removing the same absolute amount from two unequally size pools, the relative amount removed from the smaller pool will be larger.

Table 4.10

Calculated ACP for grades one through three in districts in California using the average teacher salary and assuming that different amounts of a district's budget goes to fixed costs.

Teacher assignment rule	Fixed % of budget	Mean ACP[a]	% less than zero
District Ceiling	4%	0.361	5.19%
	11%	0.431	3.31%
	18%	0.501	2.28%
	25%	0.571	1.76%
	32%	0.641	1.19%
School Ceiling	4%	0.565	1.56%
	11%	0.635	1.00%
	18%	0.705	0.57%
	25%	0.775	0.56%
	32%	0.845	0.19%

[a]Because of the large number of observations—100 trials for each of the 537 districts with more than one elementary school—the difference in ACP values based on percentage of the budget counted as fixed is significant ($F(1,267498) = 97550$, $p < .001$, and $F(1,267498) = 57380$, $p < .001$, for the district and school ceiling respectively). Unlike percentage of students removed, the percentage of the budget assumed to be fixed is important ($R^2 = 0.267$, and $R^2 = 0.177$, respectively).

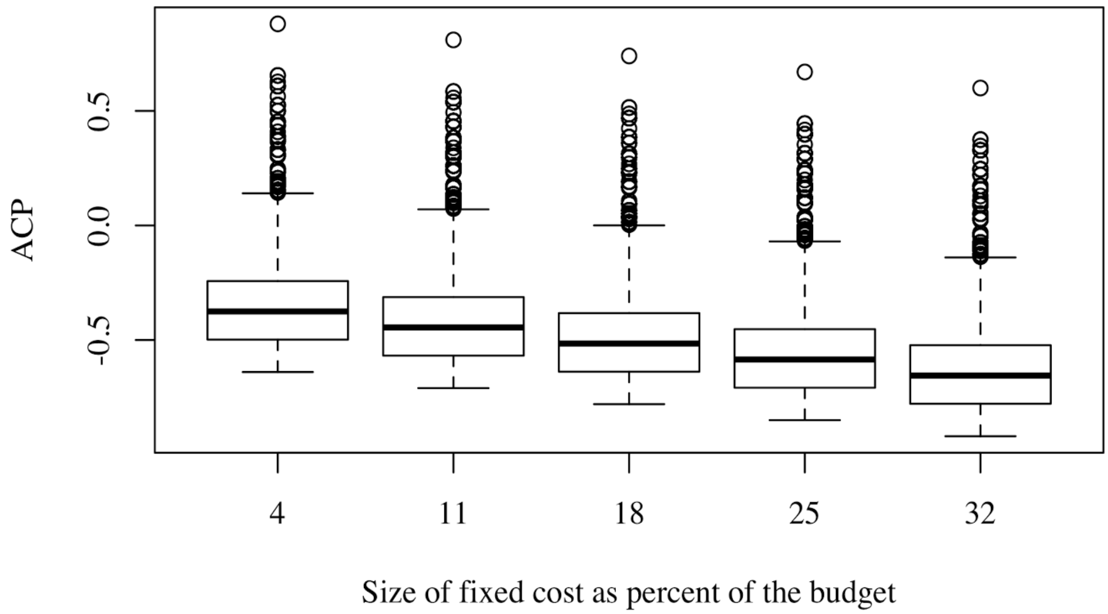

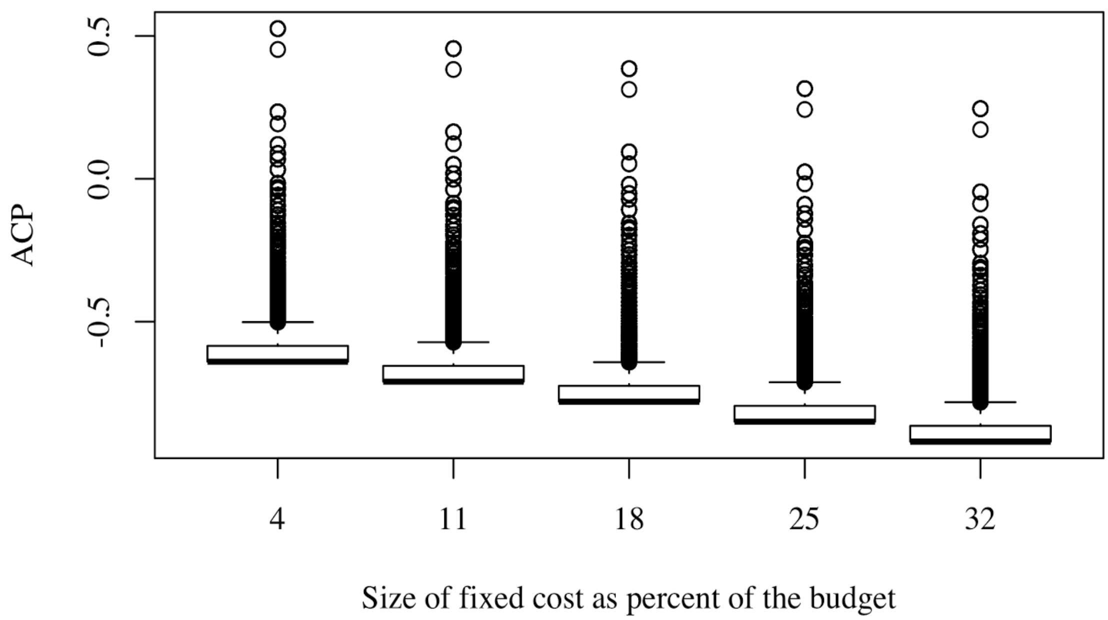

Figure 4.11. Box and whisker plots showing the frequency of ACP values in all districts in California for five assumptions about the portion of a district's budget that is fixed. The top plot uses the district ceiling rule, the bottom plot uses the school ceiling rule, and both plots remove 5% of students for participation in the school choice program.

The influence of different methods of selecting teaching positions for removal.

Just as exactly which students participate in a school choice program can affect the size of allocation effects created by non-district school choice programs, which teachers leave the district when faculty positions are eliminated will affect district finances as well. This occurs because teachers at the high end of the pay scale earn more than twice as much as teachers at the low end of the pay scale. A district that shrinks its faculty size by not filling a position after the retirement of an experienced, highly educated teacher would, we assume, reduce its faculty spending by $101,966 (the maximum salary of $81,573 times the benefit ratio of 1.25). On the other hand, a district that lets a new teacher at the bottom of the pay scale go, say, because of seniority, would only reduce faculty spending by $51,751. The district that reduced its faculty size by retirement would then have $50,251 more to spend on enriching the educational experience of students in that district compared to the district that eliminated a faculty position based on seniority. This additional $50,251 would reduce that allocation cost (or increase the allocation benefit) of the school choice program.

I will compare three different methods for deciding on the specific teachers to remove when the faculty size required in a district decreases. First, I will assume that every teacher earns the 2011 average salary of $63,903. Second, I will assume that the elimination of every teaching position can occur by not filling a position after the retirement of an experienced, highly-educated teacher at the top of the pay scale. To eliminate *n* teaching positions in a district in this way, I will create a distribution of faculty salaries as described above and eliminate the *n* greatest salaries. Finally, I will assume that all faculty reductions occur by laying off new, low-paid teachers. Like reduction by retirement, I will create a distribution as described above and eliminate the *n* lowest salaries. Because the seniority- and retirement-based methods both involve the generation of lists of salaries from a random distribution, I will conduct these calculations for the Fallbrook Union Elementary school district many times using a Monte Carlo simulation; all trials in this simulation assume that the fixed cost is 25% of the total budget and remove 5% of students for participation in the school choice program.

Figure 4.12 shows kernel density plots for 50,000 trials in the Fallbrook Union Elementary school district that remove 5% of students and assume that 25% of a district's budget is fixed. Each distribution in Figure 4.12 shows the distribution of costs using one of the methods for deciding which teachers would no longer work in the district described above.

The bottom plot in Figure 4.12 shows the distributions found using the school ceiling rule. Note that each distribution has an identically-located local maximum at the low end of the spectrum, near an ACP of 0.85. It may at first be surprising that this local maximum should occur at the same ACP value for all three distributions. The reason for this identically-located maximum is that the trials that result in an ACP of 0.85 are the trials where the faculty size cannot decrease—i.e., a faculty/enrollment sensitivity of zero. If no faculty positions can be eliminated, then it does not matter how the salary of a teacher is determined.

Figure 4.12 shows that the ACP experienced in a district indeed depends on the salary of the teachers that leave the payroll when the faculty size shrinks. If faculty reductions occur when more expensive teachers earning near the top of the pay scale retire and those positions are left vacant, the ACP is smaller. Conversely, if new, low-earning teachers are laid off to satisfy faculty reduction, the ACP will be larger. Furthermore, the influence of the rule used to calculate faculty cost savings becomes more pronounced as the number of teaching positions removed increases. The mean number of teaching positions eliminated in these trials using the district ceiling rule is 4.87, and 0.41 using the school ceiling rule and district ceiling rule, respectively. Which rule was used to calculate faculty cost savings explains more of the variation in the former ($R^2 = 0.644$, $p < .001$) than the latter ($R^2 = 0.031$, $p < .001$). Table 4.11 further summarizes these results.

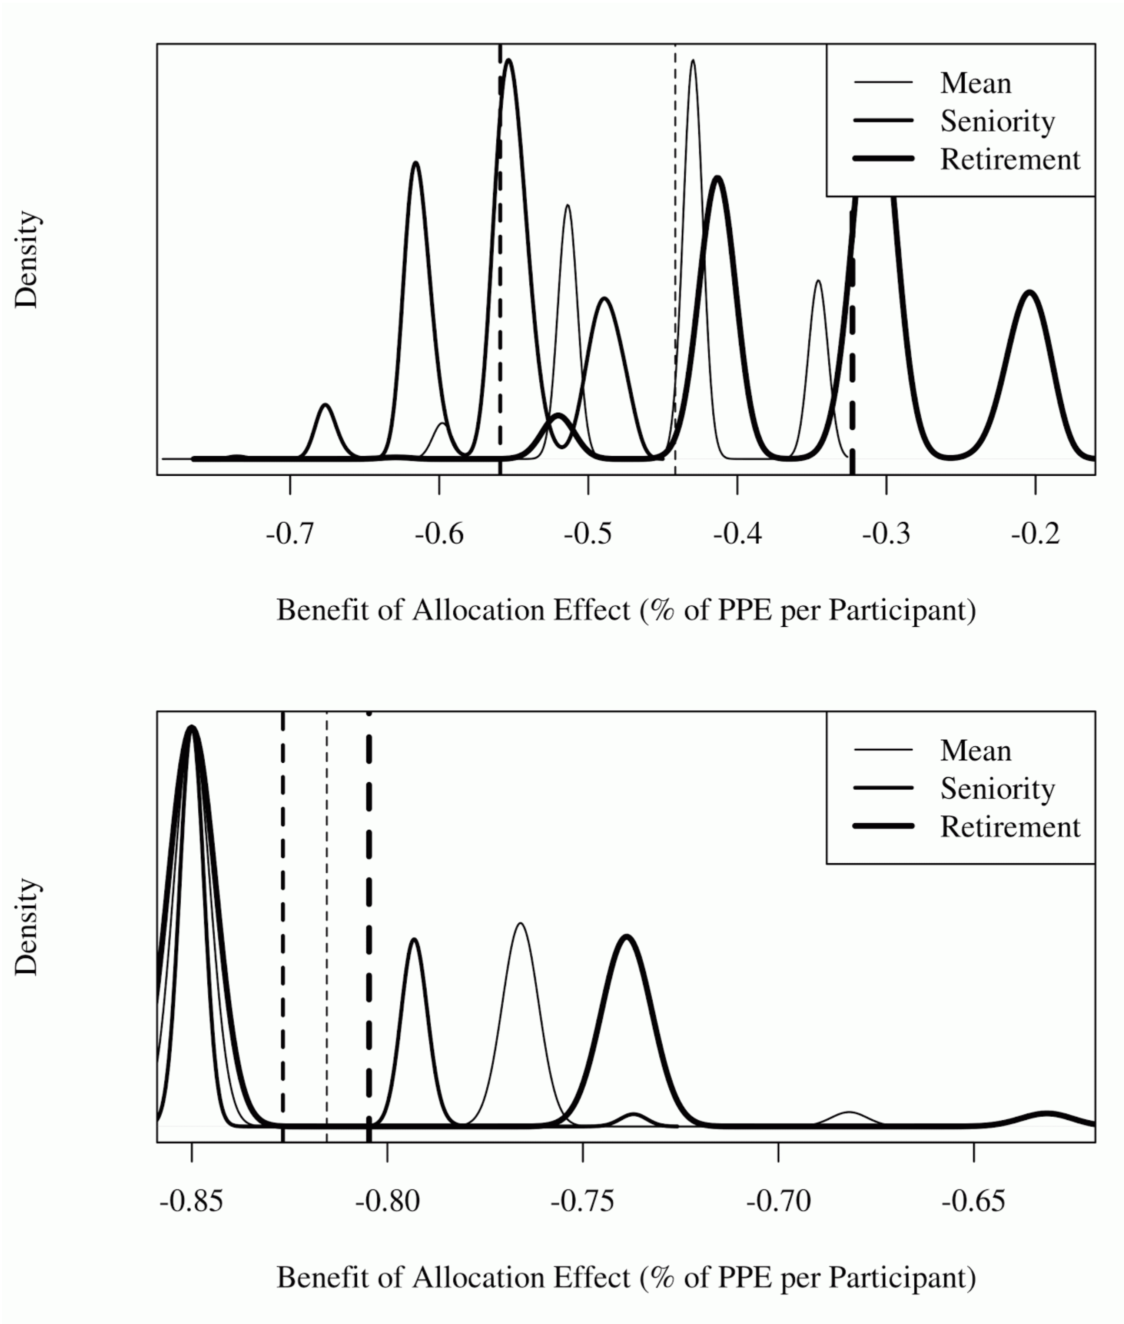

Figure 4.12. Kernel density plots showing the distribution of allocation effect per participant using different methods of determining faculty salary. Each plot was made by simulating the Fallbrook Union Elementary school district assuming that 25% of the district's budget is fixed and 5% of students are removed from grades one through three. The top plot used the district ceiling rule and the bottom plot uses the school ceiling rule.

Table 4.11
Calculated ACP using different methods of determining faculty salary for the Fallbrook Union Elementary school district removing 5% of students in grades one through three and assuming that 25% of a district's budget goes to fixed costs.

Teacher assignment rule	Faculty salary rule	Mean ACP[a]	% less than zero
District Ceiling	By seniority	0.559	0
	By mean salary	0.442	0
	By retirement	0.323	< .001
School Ceiling	By seniority	0.827	0
	By mean salary	0.815	0
	By retirement	0.805	0

[a]The difference in ACP values based on the rule used to calculate faculty cost savings is significant ($F(1,149997) = 135900$, $p < .001$, and $F(1,14997) = 2407$, $p < .001$, for the district and school ceiling respectively).

Conclusion

The simulation of a district as a financial system described here allows us to explore and understand aggregate behavior of school districts based on componential rules. Some of the findings presented here conform to our expectations, while others may surprise us. As a conclusion to this chapter, I will summarize the important behaviors of districts as financial systems that we can take away from this study.

The degree of intradistrict coordination matters.

Better coordinated districts can respond to small enrollment changes more sensitively than uncoordinated districts. I have considered a district well-coordinated when the schools in that district can communicate their class sizes with sufficient timeliness and accuracy to allow some classes to exceed the average class size target based on the knowledge that other classes are small enough to keep the district average below the target. The magnitude of this difference depends heavily on other parameters in the simulation, but we can conclude that better-coordinated districts have higher faculty/enrollment sensitivities and lower ACP values than other districts.

The degree of capacity utilization matters greatly.

Heinesen (2004) found that education systems operating near capacity are less sensitive to small enrollment declines, and the present study confirms this finding. This finding also supports the view that faculty costs behave like step costs in a district: more efficiently utilized schools are less sensitive to enrollment changes. When the school choice program targets a subpopulation of students (e.g., students eligible for free lunch), the combination of a large percentage of students belonging to the subpopulation at a school and low CU increases faculty/enrollment sensitivity. Neither of these characteristics by themselves strongly influence faculty/enrollment sensitivity, though.

The size of the "small" enrollment decline matters.

For well-coordinated districts, very small enrollment declines (near 1% of total enrollment) can lead to near zero, or extremely large, faculty/enrollment sensitivities, and anything in between. Much of the variation can be attributed to differences in CU in the district. For districts better modeled by the school ceiling rule, very small enrollment declines almost universally leads to faculty enrollment sensitivities equal to zero.

Allocation costs often exist even when faculty size responds sensitively to enrollment.

The presence of fixed costs in school districts makes school choice programs lead to an allocation cost even when the faculty/enrollment sensitivity is greater than one. This outcome damages arguments by school choice proponents that districts can respond to small enrollment changes with balanced budgets in a way that does not harm students who remain in the district. Non-district school choice programs almost always lead to a decrease in resources available for those students who do not participate.

Larger fixed costs in a district increase the allocation cost of school choice programs.

When the slice of a district's budget available to provide non-instructional services to students is smaller, a decrease in available funding (that is not offset by a decrease in faculty spending) disproportionately shrinks that slice. Districts with a greater portion of their total budget dedicated to fixed costs will see a larger allocation cost to the students who remain during a non-district school choice program.

The percentage of students removed slightly influences ACP.

The simulation shows that removing more students creates a very small impact on the ACP of the school choice program. Removing more students, however, substantially decreases the likelihood that a district will experience an allocation benefit. Trials that led to an allocation benefit were never commonly observed: in no trial was the percentage of observations resulting in an allocation benefit above 10%. These two findings together suggest that a smaller percentage of students removed for participation in a non-district school choice program leads to more erratic ACP values.

How teachers are removed during faculty reduction matters a great deal.

If we assume that new, low-earning teachers leave when the faculty size shrinks, the calculated ACP will be much larger than if we assume that experienced, high-earning teachers retire and their vacancies remain unfilled. This difference increases as the faculty/enrollment sensitivity increases, because the number of teachers acts as a multiplier for the difference in salaries. After the rule used to determine the faculty size (i.e., the school ceiling rule or the district ceiling rule), the rule used to determine which teachers leave has the greatest impact on ACP.

CHAPTER 5: Discussion

Introduction and Structure

This final chapter will discuss the results of this study. I will begin by using the results to answer the research questions posed in chapter one. I will then place the findings in a greater context of the current state of research and evaluation in the economics of education. A section describing the limitations of this study and a section discussing the generalizability of this study will follow. I will then suggest improvements and directions for future research. Finally, I will offer some concluding thoughts.

Summary and Interpretation of Results

This study primarily intends to further introduce the use of a technical method–using a computer simulation to explore an economic system—to the economics of education; describing the behavior of a district during a school choice program is a secondary goal of this study. In this section I will present answers to the research questions proposed in chapter one. A summary of the findings pertaining to the economics of school choice can be found at the end of chapter four, and will not be repeated here.

Research Question 1: Can the resources (computing power, level of detail of data and information, etc.) available to practicing evaluators generate useful predictions of the cost of school choice programs?

Like so many questions, the best answer to this question is "it depends." But the short answer to this question, if we must decide one way or the other, is "yes." When Hertz (1957) proposed this kind of simulation for financial risk analysis, one of his positions was that numerical predictions should assist human judgment, not replace it. The probability distributions he created showing the frequency of different outcomes of a financial investment were never meant to directly produce a decision; they were meant to help a human being make a decision about a course of action by providing a more thorough understanding of the likelihood of various outcomes. Like Hertz's simulations, the simulation created here produces a distribution of outcome frequencies that policy makers can consider when faced with decisions about implementing a school choice program. We do not expect that policy makers will coldy apply the rules of cost-effectiveness analysis and do whatever the numbers tell them. Rather they will weigh costs and effectivenesses along with harder-to-measure aspects of policies before deciding on a plan. The outcome of simulations like this can provide them with rich information not previously available.

On the other hand, if we require that a useful prediction is one that offers a number representing the ACP to within some small percentage—say 10%—then this method cannot likely provide a useful prediction. This simulation can generate a description of a range of ACP values, but the range produced is not particularly narrow, so a district should not use this simulation to determine, say, how much funding to raise with a bonding bill in advance of enacting a school choice program; the output is not sufficiently specific.

The barrier to generating a specific prediction is not computing power, though; the barrier is our ability to define the rules that govern how the components interact. One common theme in the findings presented here is that the degree of coordination among schools in a district with respect to school wide class sizes versus district wide class sizes matters a great deal in determining faculty/enrollment sensitivity. The two rules used in the simulation, the school ceiling rule and the district ceiling rule, aim to bookend the spectrum of intradistrict coordination. The spectrum of intradistrict coordination does not contain a linear range of numeric locations, however: stating that a district falls 65% of the way from the school ceiling rule to the district ceiling rule makes little sense on its own, and would be essentially impossible to write in computer code. Even if describing a district's level of coordination were possible in the code, ascertaining the degree of coordination would pose a remarkable evaluation challenge. Although references to coordination problems within a school district exist in the literature (for example, Driscoll, Halcoussis, & Svorny, 2003; Borland, Howsen, & Trawick, 2003), I could not find an assessment battery that quantitatively measures the degree of intradistrict coordination. So until school districts begin hiring and firing based on the real-time application of an imutable algorithm, which is not a policy that I espouse in any way, this simulation will not provide decision makers with specific numerical results. In short, a range of outcomes described by the school ceiling rule and the district ceiling rule may bound the precision of the output of this simulation.

Research Question 2: How sensitive are faculty costs to enrollment changes caused by school choice programs?

This number does not exist. Districts and school choice programs vary too much for a single number to capture reality. When a small percentage of students are drawn from well-utilized schools in a poorly-coordinated district, we can safely expect that the faculty/enrollment sensitivity will be near zero and the cost arising per participant due to value lost by the students who remain will be near the per pupil expenditure. If the school choice program targets a subpopulation of students that are concentrated in poorly utilized schools in a well-coordinated district, on the other hand, we would expect a faculty/enrollment sensitivity above one and perhaps even a benefit to the students who remain in the public schools in the district.

The absence of an answer to this question is among the most valuable potential contributions of this research. This finding shows that the cost of a school choice program will vary widely by district. And if the cost of a school choice program changes from district to district, then so will the attractiveness of school choice programs as a policy for addressing low student achievement. Thus school choice is too blunt an instrument to warrant serious consideration as an element of state- or nationwide education policy.

Some aggregate behavior observed in this simulation can contribute to our understanding of how allocation costs arise due to school choice programs. We can conclude that the presence of fixed costs makes the ACP vary within a smaller range than the faculty/enrollment sensitivity. Even when the faculty/enrollment sensitivity is near one—meaning that the percentage of students leaving the district and the percentage of teachers leaving the district are approximately equal—the ACP can still be between 35% and 65% of the per pupil expenditure in the district. In order for a school choice program to result in a benefit to students who remain in the district, the faculty/enrollment sensitivity needs to fall substantially above one, which is not common in any but the most favorable circumstances. We can also conclude that better utilized districts tend to have lower faculty/enrollment sensitivities, and that leads to higher ACP values.

Research Question 3: How much will a proposed school choice program cost due to inefficiencies created by allocation effects?

The answer to this question will sound very similar to the previous one: it depends too much on the particular district and school choice program to come up with a single number. The range of reasonable ACP values is even quite large—somewhere between 0.3 and 0.85—but this still offers two insights.

First, we cannot reasonably expect ACP values of zero, and values near the higher end of the ACP range appear to be more common. The school choice proponents who claim that policies where money follows students result in cost neutrality are wrong about that point. Allocation effects caused by school choice programs will lead to a reduction in resources available to students who remain in the district schools, even though the magnitude of this reduction depends heavily on the district and the school choice program.

Second, this points to an interesting paradox for school choice proponents. Purely theoretical discussions of fixed costs in schools have lead to the creation of school choice policies that let an amount of money smaller than the per pupil expenditure follow students (Gotlob, 2004). Suppose that based on a simulation in a school district, we suspect that the ACP for a school choice program will be 0.4; we could reduce the cost to the students not participating by only letting 60% of the per pupil expenditure go to the school of the participant's choice. The other 40% would remain with the district school to keep the level of resources available per student equal to the level present before the school choice program. Thus neither the students nor the schools feel a cost due to the school choice program. One goal of school choice programs, however, is to incent schools to perform better as a way of retaining students and the dollars that follow them. If the amount of money that follows students is chosen such that no one feels any pain because of the school choice program, then this pressure will not exist: schools will have no financial incentive to try to retain students. We can reason that an ACP greater than zero is a necessary theoretical component of school choice programs, but the distasteful position of taking resources away from students leads some school choice advocates to search for a pain-free policy. The courses of action suggested by these two perspectives mutually exclude each other.

Sending a student to a school with some fraction of the per pupil expenditure leads us to wonder what the school receiving that student can do with it, especially in the case of private school vouchers. If the per pupil expenditure approximates the amount of money necessary to educate a student, and students using vouchers only bring part of that amount of money with them, then we must assume that the rest of the resources required for that student's education must come from somewhere else. This injection of resources would count as a further cost of the voucher program, in addition to the benefit lost by students who do not participate. Furthermore, we expect that as the monetary value of an education voucher decreases, so do the number of families for whom the voucher will tip the scale toward attending a private school. The rest of the voucher program participants would have attended a private school in any event, and the voucher accomplishes little more than putting money back their pockets; Gotlob (2004) refers to the resources received by families planning to attend a private school with or without the voucher as the "deadweight" (p. 1) cost of the voucher program. These three factors—the lack of an incentive for public schools to perform better in school choice programs, the need for further resources at the schools that receive school choice participants, and deadweight costs—suggest that offering vouchers worth part of the per pupil expenditure will not have either the positive effects or lack the negative consequences that school choice proponents suppose.

Significance and Context of Findings

Our society has recently entered the age of Big Data. The combination of electronic data gathering and low data storage costs mean that education systems will store a large and increasing amount of data about their operation. This study has offered further proof of concept that economic tools for other fields (in this case, finance) can be used on the data that education systems now possess to elevate our understanding of schools and districts as complex systems.

The specific results of this study—that we can expect an ACP between 0.3 and 0.85—are of little help beyond piling evidence against the claim that school choice programs are cost neutral. The contribution to our understanding of school district finances in the context of school choice programs—that intradistrict coordination and capacity utilization greatly affect the cost of faculty/enrollment sensitivity—will hopefully elevate policy discussions and make our speculation about school choice programs more accurate. The great lesson pertains to using computer simulation to learn about the education system. Few researchers have adopted this approach. I hope that my contribution, when standing with the work of Nechyba, Romano, and others, will help push the popularity of computer simulation in the field of the economics of education into common use for both research and applied purposes.

Specific implications for policy makers

Education policy makers, while not generally likely to write their own simulations, will encounter them with increasing frequency. This means that policy makers will have to develop skills to understand simulations and their output. I have a few pieces of advice for policy makers as they begin to encounter simulations.

Carefully distinguish specific findings from descriptions of behavior. Suppose we wish to conduct a literature review to help develop some new policy. The aggregate behavior described by a simulation in the literature can inform our discussions about this new policy even if the data used in that simulation do not look much like the data in the policy we are considering. Just because a simulation was created to model a school district in a different location and time does not mean that we should disregard all of the results from the simulation; we should only disregard the specific predictions, or at least carefully evaluate how relevant they are.

My conversations with simulation skeptics lead me to believe that some people inappropriately question the veracity of simulation because they fall back on habits of doubt developed when examining regression. When determining the applicability of findings from a study that attempts to uncover underlying relationships by controlling for a large number of variables, like regression, we need to carefully consider if the population examined in the regression study closely matches the population affected by the policy in question. We should not apply the findings from a regression study to our new policy discussion if information about our policy setting differs substantially from the data used in that study; but this is not true for simulation. The randomization in RCTs effectively ensures that uncontrolled variables will occur evenly in the treatment and control groups, and so can be omitted from the regression model. If our population differs from the one in the RCT, though, we have no way of knowing if the prevalence or absence of uncontrolled variables in the studied population would change the relationship in our population. Simulation does not attempt to control for variables in an attempt to uncover an underlying relationship between variables of interest. Rather simulation helps us understand the aggregate behavior of componential rules. We need to worry more about the assumed componential rules in a simulation than about the data that create its initial state, which leads to my next recommendation.

Notice what is included in the model, and what is lacking. This simulation should be viewed as the first step in an iterative process, not as a final product. Future improvements on this simulation will refine assumptions and mechanisms that I have included, and include components that I have ignored or overlooked. One can safely conclude that no simulation will perfectly represent the system it intends to model, so every simulation lacks something. This simulation, for example, lacks a thorough treatment of special education, which is a very important component of a school's spending. Whoever writes the next simulation can improve upon what I have done by including a rich description of special education.

Creating a simulation that includes special education would be ideal, but acknowledging that the simulation does not include special education still aids our understanding of the behavior of districts during school choice programs. If we believe Biglaiser and Ma (2003) that schools try to attract students that cost less to educate, then we can assume that private and charter schools would attract few students requiring special education services. As students not requiring special education services leave a school, the mean cost of educating the students that remain in the school will increase, and this will increase ACP. Although we do not know the amount that the ACP would increase if we model special education in this way, we would assume that the current numeric range of ACP values underestimates the actual ACP values we would expect for a school choice program.

Pay attention to the definition of "conservative." Consider two individuals who might make a prediction about snowfall amounts: a ski resort owner and the head of snow plow contracting for an apartment complex. These individuals have an opposing financial stake in the amount of snow that the area receives, and thus will define conservative estimates differently. The ski resort owner can expect more customers, and thus more revenue, if more snow falls. If too little snow falls, the ski resort will have a short season and will have to run snow machines at great expense. When making a financial plan for the year, the ski resort owner will likely make a conservative estimate of snowfall near the low end of expected snowfall range. The head of snow plow contracting, on the other hand, will have to pay more for snow plowing if more snow falls. A conservative estimate for the snow plow contractor will lie near the high end of the expected snowfall range. Each individual places a conservative estimate far from their optimal circumstance.

Similarly, advocates on opposite sides of a policy like school choice programs would define optimal differently, and thus would make different conservative estimates. So understanding what 'conservative' means will help policy makers understand whether conservative estimates over- or underestimate important numerical metrics. In this analysis I attempted to include both ends of conservative (e.g., the district and the school ceiling rules, faculty reduction by seniority and by retirement, etc.) so that the range bounded by my estimates would describe the set of reasonably likely outcomes.

Specific implications for evaluators

Simulation of economic systems offers evaluators a powerful tool in cost-effectiveness analysis and feasibility analysis. As simulations gain popularity, I expect that evaluators will more often work with users who wish to incorporate simulations into their mix of methods. I have two specific recommendations for evaluators about using and working with simulations.

Develop the capacity to create simulations. Simulations like the one described here do not require a professional programmer; evaluators with even casual programming skills can learn how to write simple, but effective, simulations. Those evaluators without programming ability or the desire to gain it may want to make connections with professional programmers that they can hire if an evaluation calls for a simulation.

Maintain a utilization focus. Whether a simulation is written by the evaluator or someone outside the evaluation team, all evaluators will need to know how to interact with simulations. Evaluators will have to decide if simulation is an appropriate tool in a given context and will have to interpret the output of simulation and negotiate those interpretations with evaluation users. From a methodological standpoint, however, simulation is no different than any other technique: we need to make all decisions about the use of simulation from a utilization-focused perspective (Patton, 2008). We should avoid using simulation if the results of the simulation could not feasibly help the user answer evaluation questions, and we should not make the simulation any more intricate than the user will find helpful. We can determine the potential utility of a simulation in the same way as other techniques: by mocking simulation results and asking the user to describe how they will react. For example, mocking up the range of outcomes generated by a simulation can be the beginning of an important discussion between the evaluator and the user. Examining mocked probability distributions of outcomes can help the user understand his or her appetite for risk and more easily picture best- and worst-case scenarios.

Limitations

This simulation operates perfectly when running on the processor in a computer. The limitations of this study arise from the input data that initially creates the simulated districts and from the assumptions that describe the rules in the simulation.

Issues with the input data

The data from the California Department of Education contain some errors. These errors only became noticeable when data from one source conflicted with data from another source; for example, when the number of students eligible for free lunch exceeded the total number of students enrolled. This situation indicates either an error in the free lunch eligibility calculation or an error in the total enrollment calculation, and will only be noticeable when the two numbers are similar. If the number of students eligible for free lunch is half of the total enrollment this discrepancy could still exist, but would not raise a flag in the simulation. For this reason, we have no way of assessing the absolute accuracy of the input data.

The lack of a method for accurately assessing the error rate does not render the data from California useless. Trusting readers can assume that the California Department of Education would not allow errors in their data that are both very common and very large. And even though these errors exist, they should not shake our confidence in the findings from this study. Unlike a regression analysis, this simulation does not seek to uncover an underlying relationship in a set of variables. Small deviations in individual data will thus not greatly affect the relationship. This simulation strives to examine the aggregate behavior that occurs due to componential rules, so unless we have some reason to suspect that the erroneous input data would lead to different aggregate behavior, we should not count small errors in the input data as threats to the validity of these results.

Issues with assumptions

Using class size laws to predict the cost of teachers in a school only accounts for approximately half of the actual cost of the teaching faculty. This suggests that the logic used to determine staffing at a school is more complicated than a simple formula based on numbers of students, and we have no reason to suspect that this logic is either uniform to any degree or deterministic enough to fit nicely into a function in a computer simulation. While the simulation used in this research does offer some insight into how the predictable faculty size (the number of teachers that we can predict using class size rules) responds to small enrollment changes, and this insight is useful, the predictable faculty size seems to be only part of the story.

In particular, this study does not take some kinds of instruction into account. The model of the education system used in this simulation assumes that all students (1) attend class staffed by a single teacher, (2) receive instruction only from that teacher, and (3) receive an equal amount of education resources. None of these assumptions match reality well. Some students attend class with a paraprofessional for some or all of a school day. Some teachers do not run classrooms of their own, but rather assist the instruction of other teachers as content specialists. Some students receive more educational resources than others, including but not limited to students who receive special education services. The absence of these intricacies of the American education system both help explain why the calculated salaries in the simulation equal only about half of the actual spending that goes toward instruction. It also highlights the limits of the accuracy of this simulation. I assumed that all of the instructional spending not calculated directly from the enrollment was independent of the number of students. This assumption obviously breaks down for large changes in enrollment; for example, the number of paraprofessionals employed by a school would likely decline if the number of students decreases by half. The assumption that all instructional spending not calculated by enrollment is fixed would functionally increase the percentage of the school's spending that is fixed and overestimate the ACP.

Generalizability

One possible complaint about this study is that the input data are too specific: The simulation is built on three grades in one state from one academic year. This specificity does limit the utility of the numerical findings, but precise numerical findings were never a goal of this work. If instead we look at the input data as a way of starting with a realistic distribution of parameters—a set of numbers of schools in a district, numbers of students per grade, distribution of socioeconomic statuses, etc. that existed somewhere at some time—then these data do satisfy their role. In other years or other locations the mean classes might be a bit fuller, the mean schools a bit bigger, or the mean districts a bit smaller. These situations merely mean that fuller classes, larger schools, and smaller districts are more common than in the data I used. But the data upon which this simulation is built contain full classes, large schools, and small districts, and these were compared to emptier classes, small schools, and large districts. The value of using such a large range of real data as inputs is that they allow for the examination of a correspondingly large range of simulated behavior. For example, we can explore the relationship between simulated faculty/enrollment sensitivity and, say, how full a class is (measured by the variable CU). That relationship is among the most important findings generated by this study, and we would not imagine that a different mix of low-CU and high-CU schools would lead to a different relationship between CU and faculty/enrollment sensitivity. Because the most important response variable—the faculty/enrollment sensitivity—is calculated entirely within the simulation, we do not need to worry about exogenous variables, and thus the specificity of the input data does not detract from the most important findings in this study.

Improvements and Directions for Future Research

This study represents a first step into using simulation to stress-test a school district. The preceding discussion has outlined some of the limitations and shortcomings of this study, and in this section I will describe some small tweaks that could improve the current study and directions for future research. I will begin with modest changes that could improve this simulation and expand to larger research suggestions.

Using different rules to remove students

This study only examined two ways of removing students: selecting students with equal probability and selecting exclusively from a heterogeneously distributed subpopulation. Several other options exist for selecting students that could model interesting real or hypothesized behavior of schools. (1) As students are pseudorandomly selected for participation in the counterfactual school choice program, the other students in that grade could have a slightly increased probability of selection. This would model the case where the teacher quality causes families to more strongly consider changing schools. (2) Biglaiser and Ma (2003) suggest that private schools and charter schools have an economic incentive to attract students that are inexpensive to educate. The simulation presented here examined the behavior of school districts when students are chosen exclusively from a subpopulation—as with an income threshold defined in the policy—but not when students are preferentially chosen from a subpopulation. This could be accomplished in the simulation by giving students in the subpopulation a weight several times greater or smaller than everyone else, but not giving any student a probability of being selected for participation of zero.

Including the variable cost of educating students

The incentive to attract students with low education costs suggests another small improvement. Instead of representing "high cost" students in the simulation merely as having lower probability of being selected for participation in the school choice program, students could be assigned an extra cost of their education. From that extra cost, a probability of being chosen for selection could be determined, but this extra cost could also be represented in the accounting aspects of the model. If the simulation selected a student with high extra education costs for participation—a relatively unlikely event—that student's absence would free up some resources for the students who remain at the school. If the simulation selects only students with low extra education costs for participation, the resources available after paying for fixed costs and instructional costs would be even lower for students who remain.

Writing componential rules into the simulation to account for unequal costs of educating students would not pose a challenge, but two other aspects of including this feature in the model would. First, selecting numbers for “extra education costs” would require careful attention. Some students in special education programs receive education resources many times greater than the per pupil expenditure. The range of costs would be huge and running the simulation with an inaccurate frequency distribution could make the outputs erratic and meaningless. Second, describing students receiving special education services only as more expensive cogs in a financial system seems to run contrary to the spirit of the Individuals with Disabilities Education Act.

Examining allocation costs over time

The simulation created for this research only examines the impact of a school choice program during the first year of its implementation. No regression toward an efficient equilibrium over the course of several academic years is considered. The basic structure of this simulation could be recursively applied, i.e., the output of one year’s simulation could be the input of the next year’s simulation. Some amount of planning on the part of the district could then be incorporated into the model that would change the school assignment of some incoming kindergarten or first grade students so that the schools in the district would again have efficiently utilized capital.

Calibrating the rules used to assign teachers

Perhaps the most important and least accurate aspect of this research is the level of intradistrict coordination bookended by the school ceiling and district ceiling rules.

Interviews of and interactive simulations with principals and district officials would give a much more complete picture of how teachers are assigned to schools based on enrollment and in what cases a teaching position would be eliminated. This kind of inquiry could also shed some light on how the half of instructional spending not predicted by this simulation is spent and how that might behave in response to small enrollment changes. Unlike this simulation, such research would be very specific and would not likely lead to highly generalizable findings.

Updating existing research that uses estimates of the cost of school choice programs

The initial motivation for this study came from Yeh's (2007, 2010) work using cost-effectiveness analysis to compare proposed solutions to low student achievement. In these works Yeh supposes that the faculty size sensitivity to small enrollment changes will be zero and the ACP will be one. This research shows that neither of those suppositions are entirely reasonable. But they are likely closer to accurate than not. The range of ACP values suggested here could be used to update Yeh's reasoning, which would have two impacts. First, the overall cost of school vouchers in his analysis would decrease, but not by much. Second, our confidence in these estimates would increase and school choice advocates would have less opportunity to attack his assumptions.

The improvements to our understanding of allocation costs could raise the level of detail in large cost analyses of school choice programs conducted by Levin (1997). Levin's work examines many sources of cost and lost value associated with implementing a voucher program (e.g., the cost of informing the community, the cost of adjudicating disputes, etc.), and the results of this simulation could refine the section concerning allocation costs.

Including other step costs

I will propose another category of allocation cost that could be included in a simulation like this. This other cost can reasonably be included because it behaves as a step cost to some degree in that it will remain fixed for a range of enrollments, and change relatively dramatically as enrollment crosses a threshold: busing. Adding busing to the simulation would present its own set of challenges. Deciding how many buses a district requires depends on more than just the number of students in the district, it also depends on where those students live. For example, targeting certain income groups for participation in a voucher program may inadvertently target students from certain neighborhoods, which could make the relationship between the number of students in the district and the number of buses required much more complicated. In this simulation busing was included as a fixed cost, so including busing as a step cost could only decrease the projected ACP of a voucher program.

Conclusion

I began this work by describing how an improved estimate of the cost of school choice programs can help ameliorate the problem of low student achievement in some small way. In the ensuing pages I described a technique using a stress-testing model developed in the world of finance to examine allocation costs in a district caused by a school choice program. After executing this simulation, I presented both numeric results about what we can expect in terms of allocation costs, and a qualitative description of how the components of a district interact to create those allocation costs.

We should not lose sight of the fact that this dissertation ultimately aims to help the economics of education keep up with economic methods in other fields. Areas like meteorology, epidemiology, finance, and marketing have realized the benefits of using computer simulation to help understand and master complex systems. I will conjecture that these areas have adopted modern analytical technology like computer simulation because these fields have strong incentives for accurate prediction. Although accurate predictions in education are valuable, the education system does not offer the same kind of incentives for accuracy—no one is directly going to get rich by predicting the cost of an educational policy more accurately than their competitors. I hope that this work will help move computer simulation from the realm of science fiction and into the toolbox for practicing evaluators and education policy makers.

References

Appel, A. W. (1985). An efficient program for many-body simulation. *SIAM Journal on Scientific and Statistical Computing*, *6*(1), 85-103.

Bak, P., Tang, C., & Wiesenfeld, K. (1988). Self-organized criticality. *Physical review A*, *38*(1), 364-374.

Barnes, J. E., & Hernquist, L. (1992). Dynamics of interacting galaxies. *Annual review of astronomy and astrophysics*, *30*, 705-742.

Barrow, L., & Rouse, C. (2008). School vouchers: Recent findings and unanswered questions. *Economic Perspectives, 32(3)*.

Beveridge, T. (2009). No Child Left Behind and fine arts classes. *Arts Education Policy Review*, *111*(1), 4-7.

Biddle, B. J., & Berliner, D. C. (2002) Unequal school funding in the United States. *Educational Leadership, 59(8)*, 48-59.

Biglaiser, G. & Ma, C. A. (2003) Price and quality competition under adverse selection: Market organization and efficiency. *The RAND Journal of Economics, 34(2),* 266-286.

Briggs, A. H., Wonderling, D. E., & Mooney, C. Z. (1997). Pulling cost-effectiveness analysis up by its bootstraps: A non-parametric approach to confidence interval estimation. *Health economics*, *6*(4), 327-340.

Caesar, S., & Watanabe, T. (2011, July 31). Education takes a beating nation wide. *Los Angeles Times.* Retrieved from http://articles.latimes.com/2011/jul/31.

California Education Code (2006) California Education Code Section 41376. Retrieved on 11 April 2014

Chakrabarti, R., & Roy, J. (2011). Do Charter Schools Crowd out Private School Enrollment? Evidence from Michigan. Staff Report No. 472. *Federal Reserve Bank of New York*.

Cibulka, J. G. (1983) Response to enrollment loss and financial decline in urban school systems. *Peabody Journal of Education, 60(2)*, 64-78.

Cordes, J. J., Ebel, R. D., & Gravelle, J. (Eds.). (2005). *The encyclopedia of taxation & tax policy*. The Urban Insitute.

Costrell, R., Hanushek, E., & Loeb, S. (2008). What do cost functions tell us about the cost of an adequate education?. *Peabody Journal of Education, 83(2),* 198-223.

Coulson, A. J. (1998) Mythconceptions about school choice. Retrieved from SchoolChoices.org/roo/myths.htm.

Crowe, G. (2011) *Financing education in Minnesota: 2011-12.* Minnesota House of Representatives Fiscal Analysis Department. Retrieved from http://www.house.leg.state.mn.us/fiscal/files/11fined.pdf.

Darling-Hammond, L., Berry, B., & Thoreson, A. (2001). Does teacher certification matter? Evaluating the evidence. *Educational Evaluation and Policy Analysis, 23(1),* 57-77.

Dougherty, J., Harrelson, J., Maloney, L., Murphy, D., Smith, R., Snow, M., & Zannoni, D. (2009). School Choice in Suburbia: Test Scores, Race, and Housing Markets. *American Journal of Education, 115(4)*, 523-548.

Duncombe, W. (2006). Responding to the charge of alchemy: Strategies for evaluating the reliability and validity of costing-out research. *Journal of Education Finance 32(2),* 137-169.

Eckhardt, R. (1987). Stan ulam, john von neumann, and the monte carlo method. *Los Alamos Science*, *15*, 131-136.

Epple, D., Figlio, D., & Romano, R. (2004). Competition between private and public schools: testing stratification and pricing predictions. *Journal of Public Economics*, *88*(7), 1215-1245.

Epple, D., & Romano, R. E. (1998). Competition between private and public schools, vouchers, and peer-group effects. *American Economic Review*, 33-62.

Epple, D. & Romano, R. (2008) Educational vouchers and cream skimming. *International Economics Review, 49(4),* 1395-1435.

Ferris, J. S., & West, E. G. (2004). Economies of scale, school violence and the optimal size of schools. *Applied Economics, 36(15),* 1677-1684.

Fishman, J. J. (1983) Family choice: An idea whose time has come and gone? *The Urban Lawyer, 15(1),* 113-133.

Friedman, M. (1955) The role of government in education. In R. A. Solo (Ed.), *Economics and the Public Interest.* New Brunswick, NJ: Rutgers University Press.

Fox, J. (2011). *The myth of the rational market: a history of risk, reward, and delusion on Wall Street*. HarperCollins.

Gilbert, N., & Terna, P. (2000). How to build and use agent-based models in social science. *Mind & Society*, *1*(1), 57-72.

Goldhaber, D. D. (1999). School choice: An examination of the empirical evidence on achievement, parental decision making, and equity. *Educational Researcher, 28(9)*, 16-25.

Gottlob, B. J. (2004). *The fiscal impacts of school choice in New Hampshire*. Josiah Bartlett Center for Public Policy in partnership with the Milton & Rose D. Friedman Foundation.

Grebogi, C., McDonald, S. W., Ott, E., & Yorke, J. A. (1983). Final state sensitivity: An obstruction to predictability. *Physics Letters A, 99(9),* 415-418.

Gregerson, H. & Sailer, L. (1993) Chaos theory and its implications for social science research. *Human Relations 46(7)*, 777-802.

Hanushek, E. A., & Lindseth, A. A. (2009). *Schoolhouses, courthouses, and statehouses: Solving the funding-achievement puzzle in America's public schools.* Princeton University Press.

Heinesen, E. (2004). Determinants of local public school expenditure: A dynamic panel data model. *Regional Science and Urban Economics, 34(4)*, 429-453.

Hertz, D. B. (1979) Risk analysis in capital investment. *Harvard Business Review, 57(5)*, 169-181.

Hethcote, H. W. (2000). The mathematics of infectious diseases. *SIAM review*,*42*(4), 599-653.

Holeywell, R. (2013) *Education spending per student by state*. Retrieved from http://www.governing.com/gov-data/education-data/state-education-spending-per-pupil-data.html

Horngren, C. T., Foster, G., Datar, S. M., & Rajan, M. (2009) *Cost Accounting, A Managerial Emphasis*. Upper saddle river: Pearson.

Howell, W. G., Wolf, P. J., Peterson, P. E., & Campbell, D. E. (2000). Test-score effects of school vouchers in Dayton, Ohio, New York City, and Washington, D. C.: Evidence from randomized field trials. Prepared for the annual meetings of the American Political Science Association, Washington, D. C., September 2000

Hoxby, C. (2003) School choice and school competition: Evidence from the United States. *Swedish Economic Policy Review, 10,* 9-65.

Jarman, D. W., & Boyland, L. G. (2011). The Impacts of Budget Reductions on Indiana's Public Schools: The Impact of Budget Changes on Student Achievement, Personnel, and Class Size for Public School Corporations in the State of Indiana. *Current Issues in Education, 14(2)*, 1-9.

Jencks, C. (1970) Giving parents money for schooling: Education vouchers. *The Phi Delta Kappan, 52(1),* 49-52.

Kenny, L. W. (1982) Economies of scale in schooling. *Economics of Education Review, 2(1),* 1-24.

Ladson-Billings, G. (2006). From the achievement gap to the education debt: Understanding achievement in US schools. *Educational Researcher,* 35(7), 3-12.

Leibenstein, H. (1966). Allocative efficiency vs." X-efficiency". *The American Economic Review, 56(3)*, 392-415.

Leombruni, R., & Richiardi, M. (2005). Why are economists sceptical about agent-based simulations?. *Physica A: Statistical Mechanics and its Applications*, *355*(1), 103-109.

Levin, H. M., & Driver, C. E. (1997). Costs of an educational voucher system. *Education Economics, 5(3)*, 265.

Levin, H. M., & McEwan, P. J. (2001). *Cost-effectiveness analysis: Methods and applications* (2nd ed.). Sage Publications, Incorporated.

Lloyd, S. C. (2008). States notch slow, steady progress toward consistent graduation goals. *Education Week, 27(33)*.

Manson, S. M. (2001). Simplifying complexity: A review of complexity theory. *Geoforum, 32(3)*, 405-414.

Marschke, R., Laursen, S., Nielsen, J. M., & Dunn-Rankin, P. (2007). Demographic inertia revisited: An immodest proposal to achieve equitable gender representation among faculty in higher education. *The Journal of Higher Education*, *78*(1), 1-26.

Michigan Department of Education. (2010) *Public School Academy Report to the Legislature*. Retrieved from http://www.michigan.gov/documents/mde/2010_Legislative_Report_-_FINAL_356780_7.pdf. on 7 December 2012.

Minnesota Department of Education. (2011) *Minnesota education statistics summary*. Retrieved from http://education.state.mn.us/MDEAnalytics/Summary.jsp. on 7 December 2012.

National Center for Education Statistics. (2011) Table 101: Number and enrollment of public elementary and secondary schools, by school type, level, and charter and magnet status. *Digest of Education Statistics.* Retrieved from nces.ed.gov/programs/digest/d11/tables/dt11_101.asp

Nechyba, T. J. (1999). School finance induced migration and stratification patterns: the impact of private school vouchers. *Journal of Public Economic Theory*, *1*(1), 5-50.

Nechyba, T. J. (2000). Mobility, targeting, and private-school vouchers. *American Economic Review*, 130-146.

Ostrom, T. M. (1988). Computer simulation: The third symbol system. *Journal of experimental social psychology*, *24*(5), 381-392.

Paine, T. (1791) The rights of man. Retrieved from ushistory.org/paine/rights/singlehtml.htm

Patton, M. Q. (2008). *Utilization-focused evaluation*. Sage.

Riew, J. (1966) Economies of scale in high school operation. *The Review of Economics and Statistics, 48(3),* 280-287.

Rickles, D., Hawe, P., & Shiell, A. (2007). A simple guide to chaos and complexity. *Journal of Epidemiology and Community Health, 61(11)*, 933-937.

Robinson, R. (1993). Cost-effectiveness analysis. *British Medical Journal, 307(6907),* 793-795.

Rowan, B., Camburn, E., & Barnes, C. (2004). Benefiting from comprehensive school reform: A review of research on CSR implementation. *Putting the pieces together: Lessons from comprehensive school reform research,* 1-52.

Scafidi, B. (2012) The fiscal effects of school choice programs on public school districts. *The Friedman Foundation for Educational Choice NATIONAL RESEARCH, March,* 1-36.

Senge, P. M., & Sterman, J. D. (1992). Systems thinking and organizational learning: Acting locally and thinking globally in the organization of the future. *European journal of operational research*, *59*(1), 137-150.

Silver, N. (2012) *The Signal and The Noise: Why So Many Predictions Fail—But Others Don't.* New York: Penguin.

Sorge, M. (2004). Stress-testing financial systems: an overview of current methodologies. *Bank for International Settlement Working Papers, 165*.

Tengs, T. O., Osgood, N. D., & Chen, L. L. (2001). The cost-effectiveness of intensive national school-based anti-tobacco education: results from the tobacco policy model. *Preventive medicine*, *33*(6), 558-570.

Tiebout, C. M. (1956). A pure theory of local expenditures. *The journal of political economy*, *64*(5), 416-424.

Weidlich, W. (2002). Sociodynamics-a systematic approach to mathematical modelling in the social sciences. *NONLINEAR PHENOMENA IN COMPLEX SYSTEMS-MINSK-*, *5*(4), 479-487.

West, A., Ingram, D., & Hind, A. (2006) “Skimming the cream”?: Admission to charter schools in the United States and to autonomous schools in England. *Educational Policy, 20(4)*, 615-639.

Westfall, R. S. (1994). *The Life of Isaac Newton*. Cambridge University Press.

Wilson, E. J., Wilson, R., & Krizek, K. J. (2007) The implications of school choice on travel behavior and environmental emissions. *Transportation Research Part D: Transport and Environment, 12(7)*, 506 - 518.

Yeh, S. S. (2007) The cost-effectiveness of five policies for improving student achievement. *American Journal of Evaluation, 28(4),* 416-436

Yeh, S. S. (2010) The cost-effectiveness of 22 approaches for raising student achievement. *Journal of Education Finance, 36(1),* 38-75.