

On the use of operations research and management in public education systems

Karen Smilowitz

Department of Industrial Engineering and Management Sciences

Northwestern University, Evanston, Illinois, United States

Email: ksmilowitz@northwestern.edu

Samantha Kepler

Technology and Operations, Stephen M. Ross School of Business

University of Michigan, Ann Arbor, Michigan, United States

Email: srmeyer@umich.edu

Abstract: The tutorial explores a series of questions related to the use of operations research (OR) and operations management (OM) to improve public education systems. These questions are: (i) In the past century, how have changes in education policy prompted new OR/OM areas of research? (ii) What are the current research areas of central importance in education operations? and (iii) How can researchers interested in studying education operations establish partnerships with districts, schools, or education companies? Drawing on key papers from the literature and our experiences, as well as experiences of other researchers and practitioners who have engaged in partnerships, we answer these questions and highlight emerging areas of research in public education. We trace the history of OR/OM research in education chronologically, starting from the 1950s, and discuss the legal, political, economic, technological, and social events that have marked major shifts in public education provision. Though still developing at the time of publication, we discuss COVID-19 and its tremendous impact on education provision. Our broad review explores the links between issues in education over time and OR/OM research interests, methods, and technology.

Keywords: education operations; school bus routing; school location; school redistricting; school choice; edtech; educational inequality

1 Introduction

In this tutorial, we trace the history of operations research and operations management (OR/OM) in the context of public education chronologically, starting from the 1950s to the present. We discuss the legal, political, economic, technological, and social events that have marked major shifts in public education provision. We take a chronological approach to underscore how these events have influenced OR/OM research directions in the area of education operations. The sections of the chapter are divided by major themes in education policy. While the research presented in each section span many decades, the sections themselves are presented chronologically. Section 2 reviews research generated by global changes in beliefs beginning in the 1950s through the 1970s about *school access*, such as school redistricting to racially desegregate public education in the United States and new school location in rural areas of India. Section 3 reviews studies generated by changes in the 1980s when governments faced economic incentives to improve *fiscal efficiency* in schools and districts. Examples include efforts to reduce transportation costs and efforts to more fairly allocate limited resources to districts that serve a high proportion of students from low-income families. Section 4 reviews studies that have examined education operations in the aftermath of *market reforms* of the 2000s that have increased school-versus-school competition for students and resources. Finally, Section 5 reviews research that has examined the growing *edtech* marketplace, which throughout the 2010s has sought to personalize education provision with technologies such as online courses and machine learning algorithms. In this section, we also discuss the challenges faced in education that we are witnessing as a result of the COVID-19 pandemic.

As we highlight the expansion of relevant OR/OM problems in education over time, we also highlight the growing diversity of solution approaches and methodologies. The OR/OM literature in education is extensive and this tutorial is not intended to be exhaustive; this is simply not possible in the scope of a tutorial chapter and many surveys exist on particular subsets of this work; e.g., Johnes (2015), Witte and López-Torres (2017), Park and Kim (2010) and Ellegood et al. (2020). Our focus is on the ways in which OR/OM work in education has evolved through the parallel lenses of educational policy and advances in OR/OM techniques. For example, we discuss the changes in computing power in the 1990s that improved our ability to handle large-scale problems in transportation and logistics. We describe two studies that have used qualitative and field-based methodologies to better understand the non-pecuniary aspects of education provision, such as trust and service quality. Furthermore, we discuss the implications of “big data” in education, such as those generated by online courses that include records of individual learning over time at a large scale.

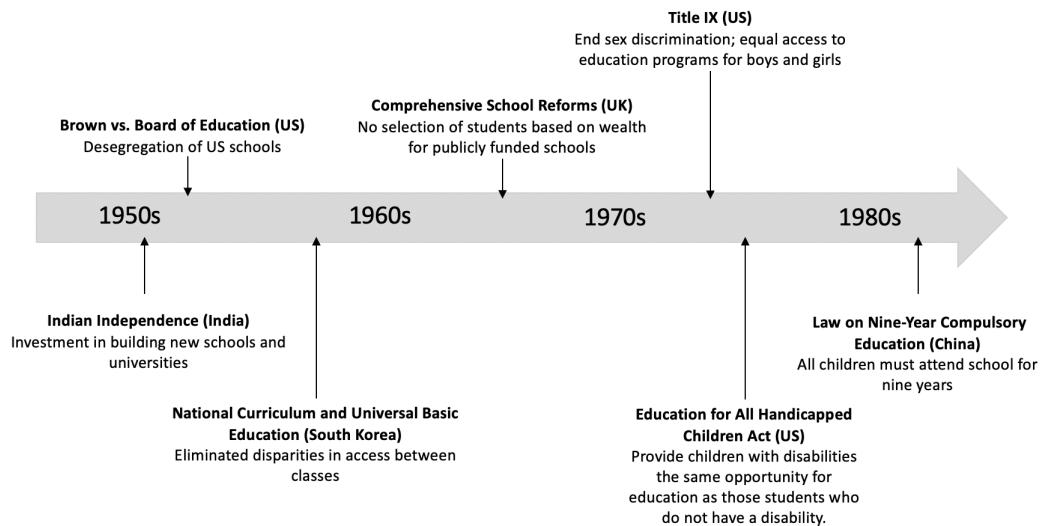
We conclude the tutorial with two forward-looking discussions in Section 6. The first identifies promising open problems in education operations and the second discusses practical concerns about how to partner with school districts or education companies/non-profits to conduct work in this

space.

2 School Access

Some of the earliest examples of OR/OM in education are found in the area of school access. Concerns about school access increased in many countries beginning in the 1950s, as people came to realize the economic and moral reasons to ensure all children had an opportunity to attend a quality school. In some countries, the focus was on dismantling entrenched discrimination by race or gender in educational access, whereas in other countries the focus was more so on reducing historical inequality in access for children of different social classes or different family wealth. The time line in Figure 1 highlights a sample of some of the national efforts to improve school access from the 1950s to 1980s.

Figure 1: Examples of National Efforts from the 1950s to 1980s to Improve School Access



Three OR/MS problem frameworks dominated the ways in which researchers responded to calls for improved school access: (1) redistricting, (2) transportation, and (3) facility location. Within each, researchers expanded existing OR/OM work to consider fairness and social justice issues in education, in addition to considerations of costs and operational constraints.

2.1 Redistricting School Boundaries

Operations research methodologies have been used to identify and evaluate solutions to the reconfiguration of public school attendance area boundaries for over sixty years. In broad terms, the school redistricting problem seeks to find capacity-feasible assignments of students in a school district to a local school (also referred to as “rezoning” or “optimizing attendance area boundaries”).

Given the movement to desegregate schools in the United States that began in the 1950s, most research on school redistricting focuses on the United States. In 1954, the United States Supreme Court ruled in *Brown v. the Board of Education* that segregating schools by race was unconstitutional and that schools should be desegregated “with all deliberate speed”. This ruling was followed by the 1968 *Green v. New Kent* ruling that called for significantly faster progress on desegregation. Around that time, papers began to appear in the OR/OM literature exploring analytical approaches to school desegregation. The titles of such papers reflect the focus on desegregation; e.g., “*School rezoning to achieve racial balance: A linear programming approach*” by Heckman and Taylor (1969), “*An operations research approach to racial desegregation of school systems*” by Clarke and Surkis (1968), and “*A network-flow model for racially balancing schools*” by Belford and Ratliff (1972).

As shown in the above titles, these papers made use of advances in linear programming approaches of the time. Beginning with this early work and continuing today, assignment variables in the redistricting problem are typically defined by street segments or blocks where students live (referred to as tracts in early papers which were much larger in geographic size due to computational limitations) and the schools to which students can be allocated. The above papers modeled the assignment variables as continuous variables, allowing split allocation of students in one location among different schools. Stimson and Thompson (1974) highlights the implications of the early simplifying model assumptions regarding assignment variables and racial balancing constraints, which were necessary given computational limitations at the time, but could lead to undesirable solutions. Over time, it became more common to explicitly model the need to assign students on the same block or street segment to the same school through integer programming approaches, beginning with Liggett (1973). This is one example of the convergence in the growing ability to solve large-scale integer programs and evolving societal views on how students should be assigned to schools. Even with acknowledged computational limitations at the time, the redistricting models of Franklin and Koenigsberg (1973) were used in *Johnson v. Richmond Unified School District* on behalf of the plaintiffs to show that feasible desegregation plans existed for the district to achieve the necessary level of diversity in schools.

The papers in the late 1960s and 1970s that applied emerging operations research techniques to school redistricting identified the unique features of the problem; many focused on the trade-offs between achieving racial balance and minimizing travel distance for students. The dual considerations of achieving racial balance and minimizing travel distance naturally motivated new ways to communicate proposed solutions to decision makers. Work by Ferland and Guénette (1990) is an early example of the move to improve visualization and interactive capabilities. Building on those ideas, a key feature in work over the past three decades has been the integration of optimization tools with Geographic Information Systems (GIS), recognizing the need to more seamlessly inte-

grate geo-data with mathematical models of redistricting and the need to provide decision makers with more powerful visualization tools for interactive decision-making with school districts; see for example, Armstrong et al. (1993) and Caro et al. (2004).

Another challenge is the need to more explicitly consider the geography of the resulting school boundary areas, including the compactness and contiguity of school neighborhoods. Such features are important in creating neighborhood cohesion and minimizing bus transportation. Compactness and contiguity are modeled in different ways in the political districting literature, but have received less attention (from a modeling perspective) in the school redistricting literature. Early work, including Holloway et al. (1975) and Diamond and Wright (1987) among others, included objectives to minimize the sum of the squared distance from student location to school, which results in more compact regions when compared with minimizing the sum of distances. In Diamond and Wright (1987), contiguity is considered through a set of constraints that ensure that a location in the district (in their case, defined as a cell in the region) is assigned to the same school as at least one neighbor. In Caro et al. (2004), contiguity is modeled more directly through the use of internal paths connecting locations to schools. Increased computational power has been a factor in more explicit modeling of compactness and contiguity.

2.2 Improving Access to Education through Transportation

Busing was a key component of desegregation in the United States. *Brown v. the Board of Education* ended *de jure* segregation (segregation of schools by law) in the United States. However, due to prevailing residential segregation, many initial redistricting solutions did not significantly reduce segregation in schools. Efforts then focused on *de facto* segregation (segregation not by law but by social factors). Subsequent court rulings, including *Swann v. Charlotte-Mecklenburg* in 1971, introduced interventions such as busing to advance school desegregation. The School Bus Routing Problem (SBRP) emerged around this time, first introduced in Newton and Thomas (1969); see Desrosiers (1980) and Park and Kim (2010) for reviews. As noted in these reviews and the papers cited wherein, the SBRP can be divided into five sub-problems: *data preparation*, *bus stop selection*, *bus route generation*, *school bell time adjustment*, and *route scheduling*. Most of the more recent work on the SBRP focuses on increasing transportation efficiency, which we discuss in Section 3.1; here, we focus on transportation to increase access. It should be noted that many papers focus on a mix of the two, and we categorize based on the primary focus of the work.

Angel et al. (1972) appears to be the first school bus routing paper to explicitly talk about busing to achieve desegregation in a project supported by a school district in Indiana. The paper focuses on data collection and a simple algorithm for schedule generation. Lee and Moore (1977) uses goal programming to balance the conflicting objectives of minimizing bus cost and distance traveled by students while providing equal access to quality schools for students of all races. This

paper falls into the assignment-type modeling framework discussed in Section 2.1 as bus routes are not explicitly generated in the paper.

Besides desegregation, the SBRP also has been used to study education access of students with disabilities. In the US, for example, *The 1975 Education for All Handicapped Children Act* (renamed the *Individuals with Disabilities Act* in 1990) called for the provision of school transportation for children with special needs. Given the specialized requirements and geographic distribution for such transportation, approaches developed for the general SBRP are not always applicable and school systems looked to researchers for assistance; the first such paper appears to be Russell and Morrel (1986) in collaboration with a school district in Oklahoma in 1985. While related work has been limited, there have been some recent studies (e.g., Braca et al. 1997, Kamali et al. 2013, Caceres et al. 2019). Notably, Kamali et al. (2013) combine bus routing with school assignment for students with special needs, recognizing that decomposing these decisions can lead to long travel distances for students. Transportation for students with special needs continues to be an important area of study given the high costs incurred for budget-constrained school districts.

2.3 Locating New Schools

Outside the US, work in the OR/OM literature focused on locating new schools as a way to expand access and educational equality. In the middle of the 20th century, many countries were facing an increased demand for schooling. This stemmed both from policies that made school newly available to certain groups of children, as well as economic growth that generated higher birth rates and thus more children to educate (Pizzolato et al. 2004). Facility location problems are central to the operations research community both in modeling and solution approaches. While researchers have used the fundamental models of facility location to study the school location problem, as with the redistricting work, “the school location problem in practice is rarely a ‘pure’ cost minimization location problem” (Giesen et al. 2015). Researchers have developed models to help districts and governments decide where to build schools to increase access to education, with a focus on the many complex factors involved in such decisions. Tewari (1992) provides an insightful examination of how location-allocation models can improve decision making related to the location of schools in India. For example, by using a p -median model to maximize the population covered within a fixed distance. Pizzolato and Da Silva (1997) and Pizzolato et al. (2004) study school location in poor areas of Brazil with a p -median approach, under additional constraints that (i) all children walk to school and (ii) preference be given to expand school capacity in addition to building new facilities. Antunes and Peeters (2000) propose a dynamic facility location model to provide guidance with the implementation of a plan in Portugal to expand elementary education.

Although overall demand for schooling increased in this era, demand shifts were also an issue. Many school districts struggled with re-configuring attendance area boundaries to meet changing

needs of their communities, as schools closed due to low utilization. Holloway et al. (1975) introduced integer variables to model the opening and closing of schools over time. School closings were also considered by Diamond and Wright (1987), which examined decisions in a multi-objective framework, reflecting the complex mix of factors impacting school closing and reassignment decisions. While some models were designed for annual reconfiguration of enrollments (e.g., Caro et al. (2004)), many school districts anticipate a commitment of at least ten years for new configurations. However, modeling uncertainty in enrollment changes at a level suitable for redistricting models is challenging. Enrollment predictions for school districts are typically performed at more aggregated levels (i.e., the school or the entire district), whereas redistricting models consider more granular geographic units such as blocks or street segments. Armstrong et al. (1993) addresses some of the challenges of including population projections in attendance area models. Reconciling the differences between district-level enrollment predictions and block-level decision variables is an open area of research.

2.4 Reflections on School Access Research

Stimson and Thompson (1975) reflect on the role of operations research using the *Brown v. Board of Education* decision and the school busing problem to illustrate a larger point about the intersection of education and operations. They refer to two different world views (“Weltanschauung”) which impact how one approaches the school busing problem: either as a desegregation problem in which one redistricts and transports students to schools to achieve racial balance objectives or as an integration problem in which one considers redistricting and bus transportation decisions within a larger set of educational decisions, recognizing that “desegregation is a spatial concept, whereas integration goes beyond space to other dimensions such as relative status and power, consideration and caring.” The paper challenges OR/OM researchers to continue to engage in critical societal issues, but with a broader lens that can change the ways in which researchers approach these problems. Stimson and Thompson (1975) then led to an exchange between researchers about the extent to which there is a truly a distinction between the two world views; see Cooper (1976) and Stimson (1976), which focus specifically on the work of Lutz et al. (1972) to desegregate high schools in Oklahoma City in response to a court ruling which proposed not to change home schools, but only to transport students across schools for extracurriculars.

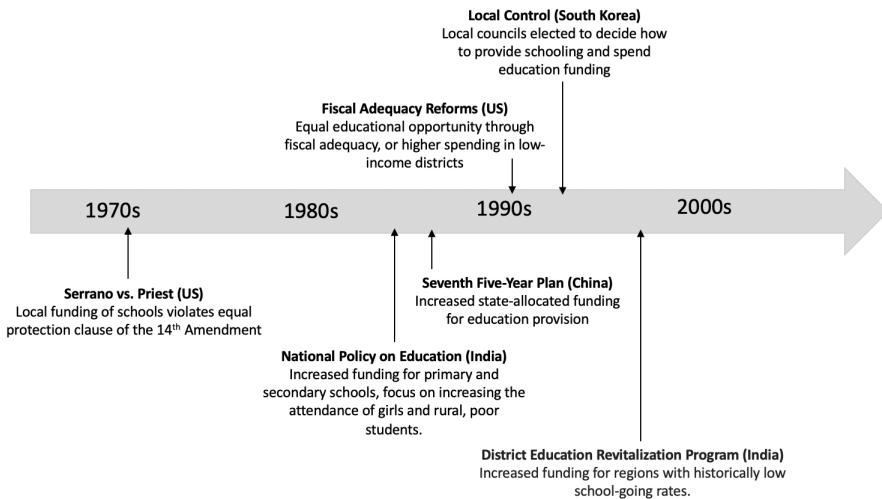
Sixty years on, issues of school access persist around the world. As such, we continue to see new work related to school access and redistricting. Recent work has focused on economic integration of schools as research continues to show how income level impacts access to high quality education, see Bouzarth et al. (2018). Today, trade-offs between providing equitable access to high quality schools and minimizing travel distances to schools are, in many ways, similar to the early conflicting objectives. However, modern versions of the problem are complicated by their intersection with

other dominant realities of education provision: persistent budget inequities, growing competition among schools, and increasingly disruptive technologies. The remainder of this tutorial is dedicated to these other operational challenges in education, and the amplification of their importance due to COVID-19.

3 Fiscal Efficiency

Starting in the 1980s, the global conversation around education increasingly included discussions of fiscal efficiency. Of great importance became how districts were spending their allocated budgets. The timeline in Figure 2 highlights a sample of the global policy changes around district budgets. Many efforts increased allocations to districts that served high proportions of low-income students. The challenge to researchers was how to spend budgets efficiently. Early on, researchers connected lean and process improvement ideas from manufacturing with education. Mood (1969) is among the earliest examples. Johnes (2015) presents a comprehensive review of operations research in education planning, efficiency and performance and logistics, covering several of the topics in this section. Witte and López-Torres (2017) focuses specifically on research in efficiency in education, looking at the economics of education in depth, rather than operations.

Figure 2: Examples of National Efforts from the 1980s to 2000s regarding Fiscal Efficiency



Two operations perspectives became the dominant way researchers responded to calls for improved fiscal efficiency: (1) cost reduction in logistics and (2) school or district efficiency evaluation.

3.1 Cost Reduction in Logistics

A major concern in education spending is to maximize the budget allocation to expenditures that impact the classroom directly, namely teacher and school staff salaries, classroom supplies and

materials. District spending on operational costs, namely transportation, inventory, and food services, are perceived as coming at the expense of the classroom. The belief is cost reductions in operational expenditures move valuable resources into the classroom, and make a direct impact on learning.

While some of the early school bus routing research focused on school desegregation, as discussed in Section 2.2, another stream of work began in the 1970s focused on improving efficiency of bus routes; see Park and Kim (2010), which reviewed 29 related papers from 1969 to 2009. As public schools faced significant budget constraints, bus transportation became a focus of cost-cutting, leading to a large rise in SBRP research in the OR/OM literature. In a recent study of the SBRP looking at work since the publication of Park and Kim (2010), Ellegood et al. (2020) cite sixty publications on the SBRP in the last decade from five continents and over twenty countries. Most papers reviewed in these surveys focus on increased cost efficiency, while some also consider equity in service provision as well. Given the wide range of school systems (e.g., by country, urban or rural), the papers represent a rich collection of SBRP variants.

Efficiency-focused OR/OM work has focused on four of the five SBRP subproblems cited in Section 2.2 with authors typically considering a single subproblem at a time or connected pairings of subproblems (e.g., bus stop selection and route generation or school bell time adjustment and route scheduling). In these papers, authors typically model SBRP components with mixed integer programming (MIP). For example, Gavish and Shlifer (1979) developed a non-linear MIP model for the bus route generation problem for a single-school and Bowerman et al. (1995) developed non-linear MIP models for the combined bus stop selection and route generation problems. Such model formulations are NP-hard and thus can lead to significant challenges for implementation at the school district level. Early work implemented traditional routing heuristics such as tour partitioning (Newton and Thomas (1969)) and the savings algorithm (Dulac et al. (1980)). Later work continued on designing more sophisticated heuristics to tackle the computational complexity, such as a genetic algorithm (Díaz-Parra et al. 2012), tabu search (Pacheco et al. 2013), and a randomized adaptive search procedure (Schittekat et al. 2013). Heuristic solution approaches for the paired subproblems of bus stop selection and route generation mainly follow two strategies: the location-allocation-routing strategy which sequentially selects the bus stops, assigns students to the bus stops, and designs bus routes (e.g., Bodin and Berman 1979) and allocation-routing-location strategy which first groups students into clusters, selects stops and generates routes for each cluster, and then assigns students to the stops (e.g., Park and Kim 2010, Bowerman et al. 1995). Advancements in optimization and new technologies of the time have allowed large school districts to significantly improve the efficiency of their operations. Notable examples can be found in the work with New York City Public Schools in the 1990s which used GIS and optimization to build software for transportation (Braca et al. 1997) and most recently with Boston Public Schools

(Bertsimas et al. 2019, 2020). The Boston Public Schools work is discussed further in Section 4.2. Uncertainty and robustness is another important factor in the SBRP which has been incorporated in recent work. Caceres et al. (2017) consider uncertainty in demand and route duration using chance constraints.

Scheduling of routes and school start times are additional components of the SBRP that have been addressed by the OR/OM community. School systems can reduce the fixed costs of school buses by staggering school start times to allow one bus to serve multiple schools; see Raff (1983), Desrosiers et al. (1986), and Fügenschuh (2009). As recent studies have shown links between school start times and student behavior and academic performance (e.g., Owens et al. 2010, Edwards 2012, Hansen et al. 2005), school systems have begun to reexamine school start times, and recent OR/OM work has focused on the equitable distribution of start times (e.g., Banerjee and Smilowitz 2019, Bertsimas et al. 2019).

Supply logistics represent another critical cost component for schools. Eisenstein and Iyer (1996) partnered with Chicago Public Schools to improve the efficiency of delivery systems, which resulted in a reconfiguration of delivery to consider items separately by physical volume, sales volume, and transportation requirements. Separating these logistics flows achieved a higher service level and reduced costs. Food services also present logistical challenges. For example, school lunch delivery is particularly challenging in rural schools requiring food to stay fresh for a long period of time. Sahin et al. (2007) model the effect of temperature monitoring technology on production and inventory decisions. Kretschmer et al. (2014) study school lunch delivery that operates as part of a humanitarian aid distribution effort in Laos involving the United Nations and UNICEF. Using this country setting as a case study, they recommend humanitarian efforts that transition global efforts to local suppliers, to reduce dependence on external resources and know-how over time. As described in Section 6.2, partnerships focused on improved operational efficiency are of critical importance to school systems.

Other researchers have studied the logistics of allocating budgets themselves. Dimitrov et al. (2014), for example, study the problem of how to fairly allocate a fixed district budget over heterogeneous schools with varying proportions of low-income students to increase school performance. Their allocation approach finds a tractable solution for a school district composed of multiple schools, where the decision is the funding allocation policy that maximizes the total discounted expected reward over a finite planning horizon. This study highlights another angle from which to approach the question of fiscal efficiency in education.

3.2 School and District Evaluation

Evaluation is the second dominant framework used to study fiscal efficiency (Witte and López-Torres 2017). Data envelopment analysis (DEA) is commonly used to relate effectiveness outcomes

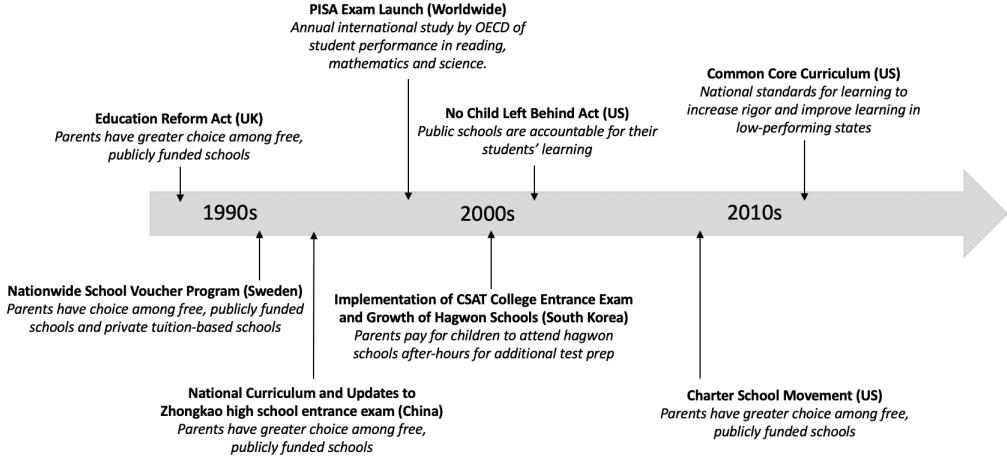
to features of organization design (Lovell et al. 1994, Seiford 1997, Bowlin 1998). DEA is an empirical methodology that generates estimates of a production frontier by comparing among a set of different producers (in this case, schools or districts). In this way, DEA assesses how well schools convert inputs to outputs, which is helpful broadly for identifying schools where improvements can be made. For example, Ruggiero and Vitaliano (1999) employ DEA to measure the efficiency of different districts in New York State, and found an average inefficiency of 14 percent; the most efficient districts were some of the larger districts in the state (Buffalo, Rochester, and Syracuse). Taking a different approach, Grosskopf et al. (1999) employ DEA to examine the impact of a proposed policy: eliminate restrictions on the allocation of school personnel. They use data from Texas to simulate the effect of such a reform, and find reforms would, on average, encourage many school districts to substitute teacher aides for teachers or professional staff.

Since this work, however, studies have shown limitations of DEA when applied to the education sector. Bifulco and Bretschneider (2001) find “in complex data sets typically used in educational research; i.e., data sets characterized by substantial measurement error and endogeneity, simple versions of DEA...do not provide adequate measures of efficiency. It would be difficult to defend implementing performance-based financing or management programs with estimates of school performance whose rank correlation with true performance is no higher than 0.24, and where no more than 31% of schools are placed in the correct performance quintile” (p. 427). For this reason, among others, we observe other empirical methodologies better equipped to deal with endogeneity emerge in response to the more recent shifts in education provision. This includes adapted DEA methods, such as those developed in Johnson and Ruggiero (2014).

4 Market Reforms

Attention to operational efficiency evolved into attention to the education marketplace. Market reforms, or reforms that relied on the market to determine school success by giving parents more choice over where they could send their child to school, were introduced in several countries in the 1990s onward. With market ideals in mind, a global movement emerged that placed greater importance on test scores as a metric of student learning and school effectiveness. It was in the 1990s that the Organization for Economic Cooperation and Development (OECD) began its Programme for International Student Assessment (PISA) global testing and country rankings.

Figure 3: Examples of National Market-Based Education Reforms from the 1990s through 2010s



Three operations areas, still emerging in the education context, have studied education provision in the aftermath of market reforms: (1) mechanism design for school choice, (2) district-level decision-support tools, and (3) service relationships with non-profits.

4.1 Mechanism Design for School Choice

Mechanism design has primarily been used to study school choice policies. Choice policies typically involve allowing parents to choose between different public schools, and/or between public and charter schools, rather than assign them to only one neighborhood school. Choice is often implemented through a lottery, at least when more parents select the school than space would permit the school to admit. Given its relevance to other lottery programs, mechanism design was a natural fit for school choice policies.

The mechanism design approach to school choice was first formulated in economics by Balinski and Sönmez (1999) and Abdulkadiroğlu and Sönmez (2003). Researchers in the field of operations have started to adapt economics models for mechanism design. One example in our field is Ashlagi and Shi (2014). Their study is motivated by school choice lotteries in Boston, Massachusetts. Schools choice lotteries operate when parents and students have choice over which public school to attend, but some schools have far more demand than they have capacity; a lottery is a fair way this is overcome. However, Ashlagi and Shi argue that there is a social cost to school choice: children do not attend the same school as their neighbors, which reduces community cohesion. They propose a correlated lottery which includes a community cohesion metric and makes assignments with students' home communities in mind, and find that they can implement such a lottery that maintains the original assignment probabilities (without community cohesion).

This work, among others, spurred additional operations studies on mechanism design in the education context. Feigenbaum et al. (2020) focus on high school choice allocations in New York

City, specifically the question of how to re-allocate assignments when spots open up at more highly-preferred schools (when, for example, an assigned student elects to attend a private school instead). They introduce a new class of re-assignment mechanisms, the permuted lottery deferred acceptance (PLDA) mechanisms. The PLDA mechanisms compute re-assignment by running a permuted lottery that prioritizes the students who received the poorest assignments (least preferred schools) in the initial assignment. The permuted lottery moves those few students many schools up their preference list rather than many students a few schools up, eliminating costly cascades of reassignment. They find that under reasonable assumptions, reversing the lottery in the re-assignment preserves *ex ante* allocative efficiency and minimizes reassignment. These studies by Ashlagi and Shi (2014) and Feigenbaum et al. (2020), and others that have built upon them (e.g., Ashlagi and Shi 2016), exemplify a broader integration between economics and operations management, but also evidence the specific application and relevance of mechanism design for the improvement of education systems in the aftermath of market reforms.

4.2 District-level decision-support tools

Another growing research area in education operations is the development and application of large-scale simulations for district-level decision support. Market reforms, particularly in combination with tight budget constraint, have put tremendous efficiency pressures on districts. These pressures must be addressed with operational changes, but not at the expense of equity and efficacy. Anticipating how changes will impact these different dimensions in a complex system with thousands of students and hundreds of schools is difficult; researchers have built simulations to support these efforts. Shi (2015) presents a collaboration with Boston Public Schools related to choice menus to offer to families (i.e., subsets of schools from which parents can choose). Shi developed a discrete choice model using historical choice data and a simulation engine that predicted outcome metrics. With this tool, the district was able to realistically evaluate different menus and guide deliberations, particularly with the public. This process ultimately led to consensus and a vote to adopt a home-based plan that offered families a menu of school choices.

Several years after Shi's collaboration with Boston Public Schools, the district was considering a new reform: school start times. As noted in Section 3.1, education research has found that early school start times for high school student is detrimental to their performance, as well as their mental and physical well-being (e.g., Owens et al. 2010, Edwards 2012, Hansen et al. 2005). Consequently, Boston considered shifting school start times; however, the shift would alter bus routes and transportation costs considerably. Bertsimas et al. (2019) studied the school bus routing and school start time problem together, which in the aftermath of market reforms was made more complex by school choice that required a significant subset of students be transported to schools far from home (See also Bertsimas et al. 2020). Their approach involved first integrating the

two objectives, optimizing school start times and transportation costs together. They created a simulation to facilitate the evaluation of current performance of the system relative to these objectives. The researchers, however, went a step further to consider parent preferences; after all, for such a reform to pass parents needed to be satisfied with the changes to start times and transportation.

They explain, “School districts have no hope of satisfying all or even most of their constituents. Moreover, the cost of even trying to satisfy the individual preferences of parents and staff can be prohibitive...each additional point of community satisfaction in Boston can cost dozens of additional buses and tens of millions of taxpayer dollars” (Bertsimas et al. 2019, p. 5947). The simulation tool helped the district visualize the relationship between start times, transportation costs, and parent preferences, which allowed leaders to thoroughly evaluate their options and make decisions based on objective standards.

4.3 Service Relationships with Non-Profits

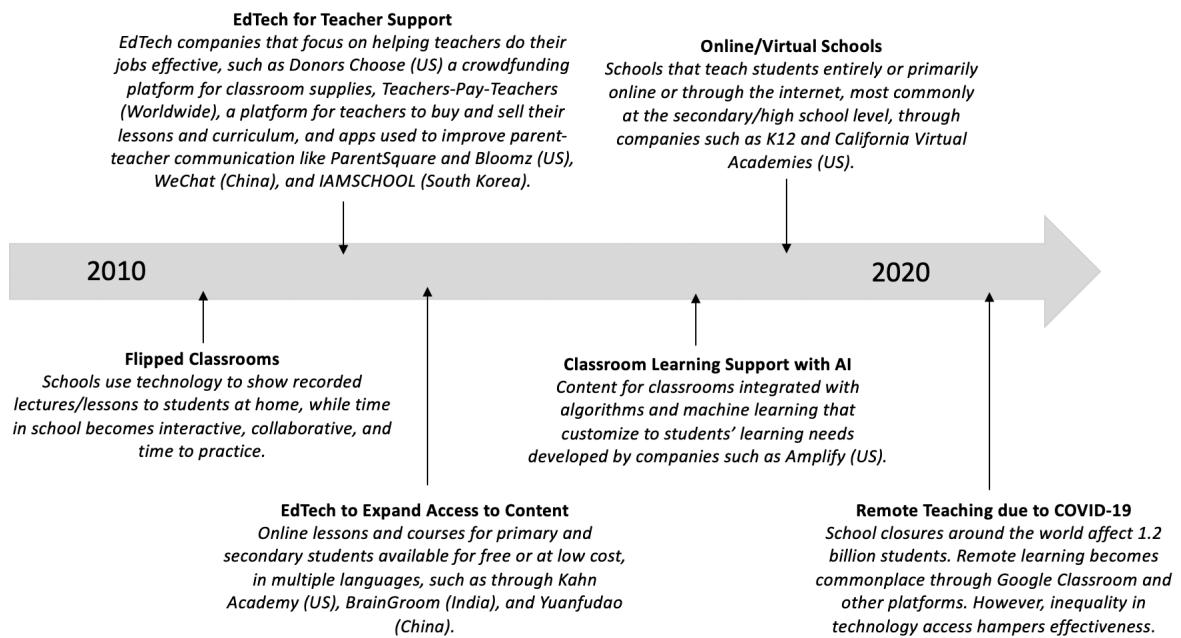
Service contracting is of general interest in operations, as it reflects the networked nature of service provision (Field et al. 2018). Service providers rarely operate in isolation and often have organizational partners or relationships (much like retailers and manufacturers have suppliers) that play an integral role in service provision (Roels et al. 2010). The networked service structure in education is increasingly common in major US cities (Bridwell-Mitchell 2017). Schools can differentiate themselves and provide a higher quality service through their service relationships, such as those with non-profits, business, donors, museums, or libraries. Keppler et al. (2020) study this contracting behavior in the education sector with a multiple case study of public, charter, and private schools in Chicago, Illinois. They explore how schools discern trustworthy partners, given school networks are not yet well-defined and potential service partners vary considerably in size, experience, location, and type (for-profit, non-profit, government agency). Schools identify potential partners as either intrinsically trustworthy or contextually trustworthy (meaning trustworthy depending on contextual factors like shared affiliations or school characteristics), but surprisingly do not avoid contextually trustworthy partners entirely. There is a competitive advantage to contextually trustworthy partners because such partners may act more trustworthy towards one school than toward other schools. Thus, schools pursue, but start with low-stakes relationships, partners they perceive to be contextually trustworthy with hopes of developing a close relationship that other schools cannot easily replicate. This work exemplifies the value of a service contracting perspective for understanding and improving education service provision.

Overall, the three operations perspectives on market reforms - mechanism design, district-level decision-support tools, and service contracting - are still nascent. More research is needed, which we discuss in Section 6.

5 EdTech and COVID-19

The most recent transformation in education has been technological. This was true even before COVID-19. Shown in Figure 5, online platforms have been developed to deliver content, support students and teachers, flip classrooms, and provide schooling virtually before 2020. However, research on EdTech from an operations perspective is limited. Some operations studies have examined technology in the higher education space (e.g., Zhang et al. 2017, Du and Li 2020) and identified opportunities for future research (e.g., Chen et al. 2020). However, the field has only just started to examine the operational impact of EdTech at the primary and secondary levels.

Figure 4: Examples of EdTech Innovations from 2010 through 2020



COVID-19 has expanded the notion of EdTech beyond a supplementary or complementary tool to a core method of education provision - even at the K-12 level. By the end of spring 2020, the positive and negative implications of this have been made clear. In some cases, teachers are innovating and students are learning remotely through technology (Li and Lalani 2020). However, more evident are the significant constraints and negative impacts of remote learning. Inequality of technology access means large number of students cannot access learning opportunities (Polikoff et al. 2020). Inequality in home environments, including children in single parent homes and children of essential workers who work outside the home, limits home learning opportunities in younger children that require more adult guidance (Widiss 2020). The same can be true when parents are working in the home (Harris 2020). In NYC, such inequalities led to four-fold increase in enrollment in remedial courses (remote, of course) in the summer of 2020. Questions of education

access and equality are re-configuring to be more around technology and remote learning, and less around school location and transportation. Moving forward, just as OR/OM researchers have done with other educational operations problems like redistricting, school assignment, transportation, procurement, and service contracting, we will need to consider equity and fairness in EdTech moving forward.

6 Next Steps

The changes in education provision over the past half-century have motivated innovative research in the area of education operations. There remain important open problems, many surfaced after COVID-19 school closures made them apparent. As this tutorial has shown, strong partnerships have been and will continue to be critical to realizing the potential of OR/OM approaches to improve public education systems. In this section, we explore open problems and partnerships.

6.1 Emerging Research Directions in Education Operations

In this section, we highlight four research directions that we believe are (1) pressing for practitioners, (2) well-suited to OR/OM methodological toolkits and expertise built to solve problems both in and outside the education sector, and (3) feasible given the increasing availability of data and partnerships in the education sector. We expect these directions to be important to practitioners in the short and long term as they deal with COVID-19 and its wide-ranging impacts.

Education and Public Health. There is a known relationship between educational outcomes and health outcomes (Lleras-Muney 2005, Cutler and Lleras-Muney 2006, Berkman et al. 2011). However, few studies have considered schools as a location where health care is provided. Schools are places where health issues, such as physical or mental illnesses, can be identified by teachers who are interacting with students every day. Providers come to schools to deliver care such as vaccines, dental care, eye exams, and so on. It is also a place where students receive meals, often breakfast and lunch. The health role schools play was made more apparent after COVID-19 school closures (Keppler 2020). Before COVID-19, Deo et al. (2013) and Savelsbergh and Smilowitz (2016) studied mobile health care vans that provide care to students with asthma in the Chicago Public School system. They study the problem of how to assign capacity under the unique conditions of chronic disease and limited resources. These studies can be used as launching points for future work on health delivery via schools. Such studies could help schools and health care providers address challenges around how to pool resources effectively and efficiently and better serve students and their families. Such problems are of critical importance as we manage COVID-19 after schools re-open. School re-openings will also bring challenges of risk mitigation as schools are also places with students and teachers in close proximity, increasing the likelihood of disease

transmission. For example, achieving social distancing in schools is a question of logistics and space management (Keppler and Smilowitz 2020), as is when and how often to test students for exposure (particularly for schools with limited resources). Broader operational questions such as how to allocate health resources across schools, and how public sector agencies, non-profits, and hospitals can coordinate efforts to maximize health outcomes, are also important. Taken together, such realities motivate empirical and analytical opportunities to study how health care delivery –and public health objectives – can be best accomplished through school access points.

Education and Sustainability. The climate crisis demands that the education sector become more sustainable. Although the education sector is by no means the largest burden on our environment, it is one where sustainable initiatives can expose children at a formative age to care for their environment and reduce waste. Districts and schools already have partner opportunities to implement environmental initiatives, such as those encouraging students to walk or bike to school (WalkBikeToSchool.org) or recycle (PepsiCoRecycling.com). Such efforts intersect with operational concerns such as transportation, supplies and inventory management, attendance, and material waste. However, an open question is how sustainability efforts enable or constrain the achievement of district and school objectives such as cost saving, equity, or performance. Exploring such a question is well-suited to a *behavioral* perspective. Behavioral studies that examine sustainability initiatives in supply chains could be a relevant starting point, such as visibility and transparency in sustainability efforts (Kraft et al. 2018, Buell and Kalkanci 2019, Kalkanci and Plambeck 2020) or the influence of mandatory versus optional disclosures of sustainability initiatives (Kalkanci and Plambeck 2020). Furthermore, studies have shown a strong link between a school's physical environment, specifically pollution and toxins, and student performance (Persico et al. 2016, Heissel et al. 2019, Persico and Venator 2019). Such evidence could promote a new wave of school (re)location problems, overlaying environmental risks with other location factors such as enrollment/population or transportation. We could also adapt resource allocation models, or build machine learning approaches based on data, to help cities make better decisions about where to invest environmental clean up efforts. Schools in Puerto Rico were significantly impacted by Hurricane Maria in 2017 and have struggled to reopen since (Alhindawi and Katz 2019). As climate issues create long-term schools closures like the school closures for COVID-19, it is critical to study the role of EdTech and remote learning. The open problems in this area have a particular sense of urgency.

Education and Human Resources. Given the important role schools play economically and socially, retaining quality principals and teachers is critical. Education researchers have found teacher (and principal) turnover, attrition, and absences are high in many countries (i.e., Ingersoll 2001, Kremer et al. 2005, Ronfeldt et al. 2013, Hanushek et al. 2016). The operational implications of turnover remain understudied in the education setting; in other words, our field is well-equipped

to study *how* turnover influences school processes. Studies have examined turnover and learning in other contexts, such as call centers (Gans and Zhou 2002), as well as the value of flexible workers that can conduct multiple tasks (parallel perhaps to teachers who can teach multiple grades) (Pinker and Shumsky 2000, Hopp et al. 2004). Beyond these examples, the theoretical link between human resource management and operations is well-established (Boudreau et al. 2003). The education context offers unique theoretical additions, such as social influence (people turnover in batches), different organizational types (public, charter, private schools with different pay and work conditions), trade-offs between hiring one experienced teacher at higher pay or two inexperienced teachers at lower pay, and the cost of switching a teacher to new grade they have not taught before. One can imagine models of an entire state or large district system, or models that capture within-school dynamics. In sum, other fields such as psychology, sociology, or management may be well-equipped to theorize why turnover, attrition, and absences happen, but our field is well-equipped to theorize how to manage processes before and after turnover.

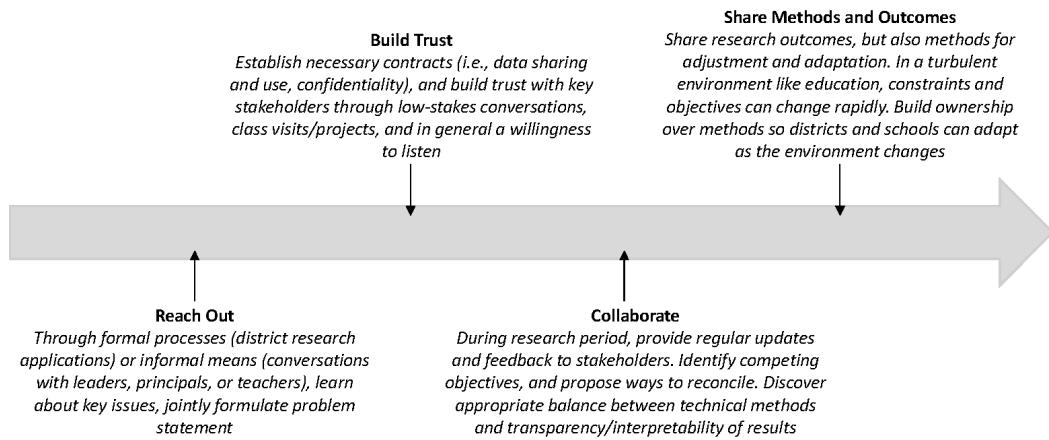
Decision Support Tools and Computational Needs. As evident in the diverse references in this chapter, operational decisions in public education systems are multidimensional. School bus transportation decisions, for example, require thinking of issues like racial equity, travel time equity, environmental sustainability, community and social cohesion, school start and end times, residential development, student mental/physical health, parent work schedules, and test score performance. Given that issues around students, families, and schooling are highly interconnected, representing decisions more completely requires modeling a highly complex decision problem and advanced tools to visualize and analyze proposed solutions. Representing the complex problem is the first challenge; the second is the computational needs and methodologies to approach such problems. As long as methods and computational power constrain our ability to analyze complex decisions in the education domain, our ability to meaningfully contribute will be constrained. Increased computational capabilities allow researchers to simultaneously consider many interdependent operational problems which were previously decomposed due to computational limitations. Equally important, solutions must be communicated clearly and allow users to explore variations as real-world constraints necessitate some degree of human manipulation of solutions. This is an important open area of research to really move the work forward.

6.2 Data and Partnerships

To meaningfully tackle open problems, researchers and practitioners need to form strong, mutually-beneficial partnerships. We discuss next best practices learned from prior work. The discussed best practices are based on our own experiences, as well as the experiences of other researchers in education operations and the experiences of districts that have partnered with researchers. We collected these experiences through an open-ended partnership survey distributed during the de-

velopment of this tutorial. The survey asked questions such as “how did your partnership begin?” and “what advice would you give to researchers and school systems starting out in this area of research?” The survey gathered behind-the-scenes information about research partnerships from five public school districts, Boston, New York City, Chicago, Evanston/Skokie and Denver, involving researchers from seven universities, Georgia Tech, Northwestern, MIT, Stanford University, Columbia University, City University of New York, and the University of Michigan. We translated our findings on partnerships into a guiding timeline of the key steps in such collaborations, shown in Figure 5.

Figure 5: Partnership Timeline



Reach Out. The first consideration to make when initiating a district relationship is whether it will be a *data partnership* or a *problem solving partnership*. A data partnership is one where the researchers have identified a particular problem of interest, and are seeking data from the district to explore it. Examples of this are Feigenbaum et al. (2020) working with NYC and Keppler et al. (2020) working with Chicago. Increasingly, state and district education departments are creating open data portals with reports on schools, test performance, attendance, and other metrics. Open data portals may not provide sufficient information for OR/OM researchers interested in operational problems in education; however, researchers can use open data together with other data from the district, or use the open data to explore and test research ideas. In the case of a data partnership, it is common for districts to have a formal research proposal process for collecting primary data or sharing secondary data that the district already collected. For securing primary data, researchers need to submit to the district’s internal review board to get approval before research can begin. For sharing secondary data, the process is in the form of a data request. Regardless of the type of data, the process can be time-consuming and require tacit know-how of the system. *A best practice is to contact your university’s education or public policy school for a sample proposal that they have submitted with repeated follow-up.* It is important to explain

the focus on operational data rather than student performance data. There can be an advantage in the data request approval process in that studies that explore operational problems (rather than questions related to curriculum and pedagogy) since student-specific information which must be rigorously protected under law is often not needed. Some universities have initiated formal partnerships at the university-district level to streamline the process; see for example <https://ocep.northwestern.edu/data-partnerships/index.html>.

For a problem solving partnership, the initiation process is different. In such partnerships, it is important to build personal relationships with those working at the district and establish strong, collaborative partnerships. Researchers in education operations are rarely successful when reaching out blind. *A best practice is to leverage an existing relationship to facilitate. Examples include researchers in the education or public policy schools at your university or researchers who have worked with related government divisions (e.g., city transportation).* As such partnerships become more common, we have seen partners from education attend OR/OM conferences; consider attending education-related sessions as this may be a source of connections. Within an organization, the question of whom to contact can also be challenging, particularly in districts without an operations division. However, most large districts now have a data and analytics team that would either be interested in learning more about solving operations problems themselves or would know the right division to contact (if the district is smaller and does not have a data and analytics team, consider instead contacts at the state level). *Another best practice is to reach out without a specific problem in mind, instead just a general area of interest (e.g., food and paper waste, student health and transportation).* Initial interactions should be focused more on hearing about the problems the district faces. As a representative from a large urban school district responded in the survey, “On the district side, it’s almost impossible to spend enough time scoping out the project. Getting clear about what, exactly, the problem you need to solve, what data you have that can be used to address the problem, and including your proposed thoughts for how to solve the problem is crucial.” Researchers will be more successful in starting a relationship by evidencing their willingness to listen. *A final best practice is to leverage local affiliations.* We have seen that MIT has worked with Boston Public Schools, Columbia University and City University of New York with NYC Public Schools, and Northwestern University with Chicago and Evanston Public Schools. There is a baseline level of familiarity and trust among local institutions that can be leveraged to motivate an introductory conversation. One way to do this is to attend public meetings as a local citizen to learn about the problems of interest to the district and community, as done by Shi (2015).

Build Trust. Problem solving partnerships require interpersonal trust among collaborators. Building trust early on is critical. *One best practice is start with low-stakes interactions.* In Russell and Morrel (1986), the authors explained how they began their collaboration with a small pilot project: “The transportation management was interested in experimenting with the smaller

routes for severely handicapped students. If the results were successful, the routing analysis was to be applied to the regular special-education routes in the following year” (p. 59). Class projects can be another low-stakes way to initiate a partnership. Denver Public Schools, for one, had worked with capstone design projects at Georgia Tech before larger-scale projects were undertaken (See <https://www.isye.gatech.edu/news/here-comes-bus-and-its-time>). Data competitions present another opportunity, such as the transportation one run by Boston Public Schools that began their relationship with MIT (Bertsimas et al. 2020). For researchers, offering low-stakes opportunities and identifying low-stakes opportunities offered by the district can build some trust between people before they take on a large-scale collaborative effort.

Once a research relationship is in discussion, early discussions should include topics of compliance and legal protections for the university and the district. *A best practice is for researchers to lead the way, and be amenable when asked, to complete compliance steps that establish legal protections for the district.* Possible steps include (1) university and district research review boards, (2) data use agreements (DUAs) between the university and district that specify what data will be shared and how it will be used, (3) non-disclosure agreements (NDAs) that establish confidentiality of some or part of the project, (4) and intellectual property (IP) if commercial technologies, patents, or companies are developed as part of the project. Such protections are at times important for districts to feel comfortable sharing data or collaborating on a project.

Another way to develop a mutual understanding of a collaborative project is to write a grant proposal. Outside funding serves multiple purposes: (i) alleviating the financial burden on the district, (ii) defining project scope before starting the project, (iii) enforcing both parties’ commitment to the project until completion. *One best practice for grant writing is to ask the district to write a letter of support that clearly states their willingness to partner on the proposed project.* The letter can be included with grant proposals written by the researchers, which makes the proposals considerably stronger. It is possible, too, that the district has a grant-writer on staff that could serve as a useful contact in the proposal development process, and may even be willing to help write the grant. The National Science Foundation has a supplemental funding program to support the engagement of teachers in research as well (<https://pubsonline.informs.org/do/10.1287/orms.2019.04.03/full/>).

Collaborate. Once a research project is underway, it is important to discuss deliverables. *A best practice is to jointly define a tangible deliverable early on.* One district representative wrote in our survey, “As a general rule, academic white papers end up gathering dust within school systems. If the focus of a partnership is on writing a white paper, districts will tend to lose interest over time. If the focus is, instead, on solving an interesting technical problem in a sustainable way, the district will stay involved.” This suggests that promising a technical solution, whether it is an interface, analytics tool, algorithm, or measurable strategy, is key to making the partnership

meaningful for the district. Defining the output early on, and key benchmarks in the process along the way, ensures the district and researchers have a shared vision for the project that motivates them to invest continually over time.

For a data partnership in particular, it is tempting to let the partnership evolve separately without the close involvement of district representatives. *A best practice is to schedule joint research meetings regularly throughout the project.* Although researchers may have interacted primarily with the research division to secure the data, insights from the project are relevant to another area of the district, such as transportation, budget, attendance/enrollment, nutrition, safety, or diversity. It is best to connect with those divisions and ask if they would be willing to provide feedback on the insights from the research project, and not just at the end. It is also possible the insights are relevant for school-level decisions; in that case, principals or even teachers can be contacted directly. For problem solving partnerships, where researchers and district representations are working closely together, it is tempting to divide the work in such a way that researchers do more technical computational work and practitioners provide context and feedback. This is reasonable, but can hamper the district's ability, in the end, to continue the work on their own. It makes them wholly dependent on the researchers. *A best practice is to build knowledge with the district, not only provide solutions.* The district should, over the course of the project, develop capabilities to adapt and advance the solution approach or decisions made to a changing environment.

Share Methods and Outcomes. Besides a research publication, projects should generate practical recommendations for states, districts, schools, or education policy-makers. *A best practice is to provide a functional output at the end of a research project, such as a simulation, tracker, interface, etc., that translates research insights into a helpful tool.* If it is adaptable and the district has built a sufficient depth of understanding about the technical workings of the tool, a functional tool can help the district implement change, and perhaps extend the impact of the project more broadly. Simchi-Levi (2014) describes how they collaborated with NYC public schools on school bus routing, including by involving the district in any technology or tool development process (See p. 4). Finally, if a future project is on the horizon, it can be advantageous to keep the relationship going even after the project is completed. *A best practice is to revert back to low-stakes interactions (i.e., classroom visits and presentations) to stay in touch with partners.* We are optimistic that given the operational challenges facing districts due to COVID-19, districts will agree with researchers as to the value of continuing these projects and staying in contact.

Other Types of Research Relationships. District partnerships are not the only option for researchers interested in studying open problems in education operations. The K-12 education sector is full of tech or traditional companies, for-profit and non-profit, serving students, parents, schools, and districts. These companies are sometimes small and flexible, have large amounts of data, and are highly motivated to improve effectiveness and efficiency. One way to

categorize companies in the education space is by whether they provide *within-system solutions* versus *outside-system solutions*. Companies that provide within-system solutions are ones that serve the brick-and-mortar school system. Examples include tech companies that build apps to support teacher-parent communication, help track student attendance and behavior, help students prep at home for standardized tests, or are traditional companies like test or textbook companies. Companies that provide outside-system solutions are ones that aim to disrupt the traditional brick-and-mortar schooling model. Examples include online learning tools for K-12 and virtual/online schools. Researchers can begin by identifying which type of company solution is of greater interest and relevance to their research agenda, and build relationships with those companies. As these companies are at times more comfortable with our methodological toolkit, they are likely to respond to friendly cold-call emails. In this case, relationships can begin just by reaching out.

7 Conclusion

In this tutorial, we discuss how OR/OM research has responded to major events in world history that changed public education provision. One of the events, COVID-19, occurred during the writing of this tutorial. Still in the midst of the COVID-19 crisis, we highlight areas of particular importance for future work. The tutorial provides evidence that education operations has generated value helping education leaders, districts, and schools solve problems that arose as a consequence of reforms. Yet, operational problems of education provision have only grown more complex and more important, and thus a larger team of researchers is needed. This tutorial provides insights for researchers interested in joining.

References

- A. Abdulkadiroğlu and T. Sönmez. School choice: A mechanism design approach. *American Economic Review*, 93(3):729–747, 2003.
- D. Alhindawi and J. Katz. The disappearing schools of Puerto Rico. *The New York Times*, September 2019. URL <https://www.nytimes.com/interactive/2019/09/12/magazine/puerto-rico-schools-hurricane-maria.html>.
- R. Angel, W. Caudle, R. Noonan, and A. Whinston. Computer-assisted school bus scheduling. *Management Science*, 18(6):B-279, 1972.
- A. Antunes and D. Peeters. A dynamic optimization model for school network planning. *Socio-Economic Planning Sciences*, 34(2):101–120, 2000.
- M. P. Armstrong, P. Lolonis, and R. Honey. A spatial decision support system for school redistricting. *URISA Journal*, 5(1):40–52, 1993.

I. Ashlagi and P. Shi. Improving community cohesion in school choice via correlated-lottery implementation. *Operations Research*, 62(6):1247–1264, 2014.

I. Ashlagi and P. Shi. Optimal allocation without money: An engineering approach. *Management Science*, 62(4):1078–1097, 2016.

M. Balinski and T. Sönmez. A tale of two mechanisms: Student placement. *Journal of Economic Theory*, 84(1):73–94, 1999.

D. Banerjee and K. Smilowitz. Incorporating equity into the school bus scheduling problem. *Transportation Research Part E: Logistics and Transportation Review*, 131:228–246, 2019.

P. C. Belford and H. D. Ratliff. A network-flow model for racially balancing schools. *Operations Research*, 20(3):619–628, 1972.

N. D. Berkman, S. L. Sheridan, K. E. Donahue, D. J. Halpern, and K. Crotty. Low health literacy and health outcomes: An updated systematic review. *Annals of Internal Medicine*, 155(2):97–107, 2011.

D. Bertsimas, A. Delarue, and S. Martin. Optimizing schools' start time and bus routes. *Proceedings of the National Academy of Sciences*, 116(13):5943–5948, 2019.

D. Bertsimas, A. Delarue, W. Eger, J. Hanlon, and S. Martin. Bus routing optimization helps Boston Public Schools design better policies. *Interfaces*, 50(1):37–49, 2020.

R. Bifulco and S. Bretschneider. Estimating school efficiency: A comparison of methods using simulated data. *Economics of Education Review*, 20(5):417–429, 2001.

L. D. Bodin and L. Berman. Routing and scheduling of school buses by computer. *Transportation Science*, 13(2):113–129, 1979.

J. Boudreau, W. Hopp, J. O. McClain, and L. J. Thomas. On the interface between operations and human resources management. *Manufacturing & Service Operations Management*, 5(3):179–202, 2003.

E. L. Bouzarth, R. Forrester, K. R. Hutson, and L. Reddoch. Assigning students to schools to minimize both transportation costs and socioeconomic variation between schools. *Socio-Economic Planning Sciences*, 64:1–8, 2018.

R. Bowerman, B. Hall, and P. Calamai. A multi-objective optimization approach to urban school bus routing: Formulation and solution method. *Transportation Research Part A: Policy and Practice*, 29(2):107–123, 1995.

W. F. Bowlin. Measuring performance: An introduction to data envelopment analysis (DEA). *The Journal of Cost Analysis*, 15(2):3–27, 1998.

J. Braca, J. Bramel, B. Posner, and D. Simchi-Levi. A computerized approach to the New York City school bus routing problem. *IIE Transactions*, 29(8):693–702, 1997.

E. N. Bridwell-Mitchell. Them that's got: How tie formation in partnership networks gives high schools differential access to social capital. *American Educational Research Journal*, 54(6):1221–1255, 2017.

R. W. Buell and B. Kalkanci. How transparency into internal and external responsibility initiatives influences consumer choice: Evidence from the field and lab. *Harvard Business School Technology & Operations Mgt. Unit Working Paper*, (19-115), 2019.

H. Caceres, R. Batta, and Q. He. School bus routing with stochastic demand and duration constraints. *Transportation Science*, 51(4):1349–1364, 2017.

H. Caceres, R. Batta, and Q. He. Special need students school bus routing: Consideration for mixed load and heterogeneous fleet. *Socio-Economic Planning Sciences*, 65:10–19, 2019.

F. Caro, T. Shirabe, M. Guignard, and A. Weintraub. School redistricting: Embedding GIS tools with integer programming. *Journal of the Operational Research Society*, 55(8):836–849, 2004.

Y. J. Chen, T. Dai, C. G. Korpeoglu, E. Körpeoglu, O. Sahin, C. S. Tang, and S. Xiao. OM Forum—Innovative online platforms: Research opportunities. *Manufacturing & Service Operations Management*, 22(3):430–445, 2020.

S. Clarke and J. Surkis. An operations research approach to racial desegregation of school systems. *Socio-Economic Planning Sciences*, 1(3):259–272, 1968.

W. Cooper. Communication to the editor—the cluster plan for desegregating public schools, or a little less weltanschauung, please, for some more adequate descriptions. *Management Science*, 22(11):1282–1283, 1976.

D. M. Cutler and A. Lleras-Muney. Education and health: Evaluating theories and evidence. Technical report, National Bureau of Economic Research, 2006.

S. Deo, S. Iravani, T. Jiang, K. Smilowitz, and S. Samuelson. Improving health outcomes through better capacity allocation in a community-based chronic care model. *Operations Research*, 61(6):1277–1294, 2013.

J. Desrosiers. *An Overview of School Busing System*. Montréal: Université de Montréal, Centre de recherche sur les transports, 1980.

J. Desrosiers, J.-A. Ferland, J.-M. Rousseau, G. Lapalme, and L. Chapleau. TRANSCOL: A multi-period school bus routing and scheduling system. *TIMS Studies in the Management Sciences*, 22:47–71, 1986.

J. T. Diamond and J. R. Wright. Multiobjective analysis of public school consolidation. *Journal of Urban Planning and Development*, 113(1):1–18, 1987.

O. Díaz-Parra, J. A. Ruiz-Vanoye, Á. Buenabad-Arias, and F. Cocón. A vertical transfer algorithm for the school bus routing problem. In *2012 Fourth World Congress on Nature and Biologically Inspired Computing (NaBIC)*, pages 66–71. IEEE, 2012.

N. B. Dimitrov, S. Dimitrov, and S. Chukova. Robust decomposable Markov decision processes motivated by allocating school budgets. *European Journal of Operational Research*, 239(1):199–213, 2014.

L. Du and Q. Li. A data-driven approach to high-volume recruitment: Application to student admission. *Manufacturing & Service Operations Management*, 2020. URL <https://doi.org/10.1287/msom.2019.0779>.

G. Dulac, J. A. Ferland, and P. A. Forgues. School bus routes generator in urban surroundings. *Computers & Operations Research*, 7(3):199–213, 1980.

F. Edwards. Early to rise? The effect of daily start times on academic performance. *Economics of Education Review*, 31(6):970–983, 2012.

D. D. Eisenstein and A. V. Iyer. Separating logistics flows in the Chicago public school system. *Operations Research*, 44(2):265–273, 1996.

W. A. Ellegood, S. Solomon, J. North, and J. F. Campbell. School bus routing problem: Contemporary trends and research directions. *Omega*, 2020. URL <https://doi.org/10.1016/j.omega.2019.03.014>.

I. Feigenbaum, Y. Kanoria, I. Lo, and J. Sethuraman. Dynamic Matching in School Choice: Efficient Seat Reassignment after Late Cancellations. *Management Science*, 2020. URL <https://doi.org/10.1287/mnsc.2019.3469>.

J. A. Ferland and G. Guénette. Decision support system for the school districting problem. *Operations Research*, 38(1):15–21, 1990.

J. M. Field, L. Victorino, R. W. Buell, M. J. Dixon, S. M. Goldstein, L. J. Menor, M. E. Pullman, A. V. Roth, E. Secchi, and J. J. Zhang. Service operations: What’s next? *Journal of Service Management*, 29(1):55–97, 2018.

A. D. Franklin and E. Koenigsberg. Computed school assignments in a large district. *Operations Research*, 21(2):413–426, 1973.

A. Fügenschuh. Solving a school bus scheduling problem with integer programming. *European Journal of Operational Research*, 193(3):867–884, 2009.

N. Gans and Y.-P. Zhou. Managing learning and turnover in employee staffing. *Operations Research*, 50(6):991–1006, 2002.

B. Gavish and E. Shlifer. An approach for solving a class of transportation scheduling problems. *European Journal of Operational Research*, 3(2):122–134, 1979.

R. Giesen, P. R. E. Oliveira, and V. Marianov. Rural school location and student allocation. In H. Eiselt and V. Marianov, editors, *Applications of Location Analysis*, pages 273–289. Springer, 2015.

S. Grosskopf, K. J. Hayes, L. L. Taylor, and W. L. Weber. Anticipating the consequences of school reform: a new use of DEA. *Management Science*, 45(4):608–620, 1999.

M. Hansen, I. Janssen, A. Schiff, P. C. Zee, and M. L. Dubocovich. The impact of school daily schedule on adolescent sleep. *Pediatrics*, 115(6):1555–1561, 2005.

E. A. Hanushek, S. G. Rivkin, and J. C. Schiman. Dynamic effects of teacher turnover on the quality of instruction. *Economics of Education Review*, 55:132–148, 2016.

E. A. Harris. ‘It was just too much’: How remote learning is breaking parents. *The New York Times*, April 2020. URL <https://www.nytimes.com/2020/04/27/nyregion/coronavirus-homeschooling-parents.html>.

L. B. Heckman and H. M. Taylor. School rezoning to achieve racial balance: A linear programming approach. *Socio-Economic Planning Sciences*, 3(2):127–133, 1969.

J. Heissel, C. Persico, and D. Simon. Does pollution drive achievement? The effect of traffic pollution on academic performance. Technical report, National Bureau of Economic Research, 2019.

C. A. Holloway, D. A. Wehrung, M. P. Zeitlin, and R. T. Nelson. An interactive procedure for the school boundary problem with declining enrollment. *Operations Research*, 23(2):191–206, 1975.

W. J. Hopp, E. Tekin, and M. P. Van Oyen. Benefits of skill chaining in serial production lines with cross-trained workers. *Management Science*, 50(1):83–98, 2004.

R. M. Ingersoll. Teacher turnover and teacher shortages: An organizational analysis. *American Educational Research Journal*, 38(3):499–534, 2001.

J. Johnes. Operational research in education. *European Journal of Operational Research*, 243(3):683–696, 2015.

A. L. Johnson and J. Ruggiero. Nonparametric measurement of productivity and efficiency in education. *Annals of Operations Research*, 221(1):197–210, 2014.

B. Kalkanci and E. L. Plambeck. Reveal the supplier list? A trade-off in capacity vs. responsibility. *Manufacturing & Service Operations Management*, 2020. URL <https://doi.org/10.1287/msom.2019.0795>.

B. Kamali, S. J. Mason, and E. A. Pohl. An analysis of special needs student busing. *Journal of Public Transportation*, 16(1):2, 2013.

S. Keppler. 7 things public schools do besides teach kids academic basics. *The Conversation*, April 2020. URL <https://theconversation.com/7-things-public-schools-do-besides-teach-kids-academic-basics-134622>.

S. Keppler and K. Smilowitz. School buildings need more space to safely reopen. *The Conversation*, May 2020. URL <https://theconversation.com/school-buildings-need-more-space-to-safely-reopen-138401>.

S. Keppler, K. Smilowitz, and P. Leonardi. Contextual trustworthiness of organizational partners: Evidence from nine school networks. *Manufacturing and Service Operations Management*, 2020. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2944087.

T. Kraft, L. Valdés, and Y. Zheng. Supply chain visibility and social responsibility: Investigating consumers' behaviors and motives. *Manufacturing & Service Operations Management*, 20(4):617–636, 2018.

M. Kremer, N. Chaudhury, F. H. Rogers, K. Muralidharan, and J. Hammer. Teacher absence in India: A snapshot. *Journal of the European Economic Association*, 3(2-3):658–667, 2005.

A. Kretschmer, S. Spinler, and L. N. Van Wassenhove. A school feeding supply chain framework: Critical factors for sustainable program design. *Production and Operations Management*, 23(6):990–1001, 2014.

S. M. Lee and L. J. Moore. Multi-criteria school busing models. *Management Science*, 23(7):703–715, 1977.

C. Li and F. Lalani. The COVID-19 pandemic has changed education forever. This is how. *World Economic Forum*, April 2020. URL <https://www.weforum.org/agenda/2020/04/coronavirus-education-global-covid19-online-digital-learning/>.

R. S. Liggett. The application of an implicit enumeration algorithm to the school desegregation problem. *Management Science*, 20(2):159–168, 1973.

A. Lleras-Muney. The relationship between education and adult mortality in the United States. *The Review of Economic Studies*, 72(1):189–221, 2005.

C. K. Lovell, L. C. Walters, and L. L. Wood. Stratified models of education production using modified DEA and regression analysis. In *Data Envelopment Analysis: Theory, Methodology, and Applications*, pages 329–351. Springer, 1994.

R. Lutz, M. Devine, H. Kumin, and W. Smith. An application of operations research to school desegregation. *Management Science*, 19(4-part-2):P100–P109, 1972.

A. M. Mood. Macro-analysis of the American educational system. *Operations Research*, 17(5):770–784, 1969.

R. M. Newton and W. H. Thomas. Design of school bus routes by computer. *Socio-Economic Planning Sciences*, 3(1):75–85, 1969.

J. A. Owens, K. Belon, and P. Moss. Impact of delaying school start time on adolescent sleep, mood, and behavior. *Archives of Pediatrics & Adolescent Medicine*, 164(7):608–614, 2010.

J. Pacheco, R. Caballero, M. Laguna, and J. Molina. Bi-objective bus routing: An application to school buses in rural areas. *Transportation Science*, 47(3):397–411, 2013.

J. Park and B.-I. Kim. The school bus routing problem: A review. *European Journal of Operational Research*, 202(2):311–319, 2010.

C. Persico, D. Figlio, and J. Roth. Inequality before birth: The developmental consequences of environmental toxicants. Technical report, National Bureau of Economic Research, 2016.

C. L. Persico and J. Venator. The effects of local industrial pollution on students and schools. *Journal of Human Resources*, pages 0518–9511R2, 2019.

E. J. Pinker and R. A. Shumsky. The efficiency-quality trade-off of cross-trained workers. *Manufacturing & Service Operations Management*, 2(1):32–48, 2000.

N. D. Pizzolato and H. F. Da Silva. The location of public schools: Evaluation of practical experiences. *International Transactions in Operational Research*, 4(1):13–22, 1997.

N. D. Pizzolato, F. B. Barcelos, and L. A. Nogueira Lorena. School location methodology in urban areas of developing countries. *International Transactions in Operational Research*, 11(6):667–681, 2004.

M. Polikoff, A. Rosefsky Saavedra, and S. Korn. Not all kids have computers – and they're being left behind with schools closed by the coronavirus. *The Conversation*, May 2020. URL <https://theconversation.com/not-all-kids-have-computers-and-theyre-being-left-behind-with-schools-closed-by-the-coronavirus-137359>.

S. Raff. Routing and scheduling of vehicles and crews: The state of the art. *Computers & Operations Research*, 10(2):63–211, 1983.

G. Roels, U. S. Karmarkar, and S. Carr. Contracting for collaborative services. *Management Science*, 56(5):849–863, 2010.

M. Ronfeldt, S. Loeb, and J. Wyckoff. How teacher turnover harms student achievement. *American Educational Research Journal*, 50(1):4–36, 2013.

J. Ruggiero and D. F. Vitaliano. Assessing the efficiency of public schools using data envelopment analysis and frontier regression. *Contemporary Economic Policy*, 17(3):321–331, 1999.

R. A. Russell and R. B. Morrel. Routing special-education school buses. *Interfaces*, 16(5):56–64, 1986.

E. Sahin, M. Z. Babaï, Y. Dallery, and R. Vaillant. Ensuring supply chain safety through time temperature integrators. *The International Journal of Logistics Management*, 18(1):102–124, 2007.

M. Savelsbergh and K. Smilowitz. Stratified patient appointment scheduling for mobile community-based chronic disease management programs. *IIE Transactions on Healthcare Systems Engineering*, 6(2):65–78, 2016.

P. Schittekat, J. Kinal, K. Sörensen, M. Sevaux, F. Spieksma, and J. Springael. A metaheuristic for the school bus routing problem with bus stop selection. *European Journal of Operational Research*, 229(2):518–528, 2013.

L. M. Seiford. A bibliography for data envelopment analysis (1978-1996). *Annals of Operations Research*, 73:393–438, 1997.

P. Shi. Guiding school-choice reform through novel applications of operations research. *Interfaces*, 45(2):117–132, 2015.

D. Simchi-Levi. OM forum—OM research: From problem-driven to data-driven research. *Manufacturing & Service Operations Management*, 16(1):2–10, 2014.

D. H. Stimson. Communication to the editor—The dangers of research based on one telephone call to Dallas. *Management Science*, 22(11):1283–1285, 1976.

D. H. Stimson and R. P. Thompson. Linear programming, busing and educational administration. *Socio-Economic Planning Sciences*, 8(4):195–206, 1974.

D. H. Stimson and R. P. Thompson. The importance of “weltanschauung” in operations research: The case of the school busing problem. *Management Science*, 21(10):1123–1132, 1975.

V. K. Tewari. Improving access to services and facilities in developing countries. *International Regional Science Review*, 15(1):25–37, 1992.

D. Widiss. Parental leave laws don’t do enough for single moms – but there’s a way to fix that. *The Conversation*, May 2020. URL <https://theconversation.com/parental-leave-laws-dont-do-enough-for-single-moms-but-theres-a-way-to-fix-that-137360>.

K. D. Witte and L. López-Torres. Efficiency in education: A review of literature and a way forward. *Journal of the Operational Research Society*, 68(4):339–363, 2017.

D. J. Zhang, G. Allon, and J. A. Van Mieghem. Does social interaction improve learning outcomes? Evidence from field experiments on massive open online courses. *Manufacturing & Service Operations Management*, 19(3):347–367, 2017.