An Agent-Based Model of Intra-District Public School Choice

Spiro Maroulis* Eytan Bakshy† Louis Gomez‡ Uri Wilensky§

September 24, 2010

Working paper. Comments welcome.

Email: s-maroulis@kellogg.northwestern.edu

Abstract

This paper addresses the difficulties in evaluating the impact of choice-based reforms in education by developing an agent-based model that captures the dynamic processes related to moving from a neighborhood-based to choice-based system. Using data from Chicago Public Schools to estimate and set initial conditions of the model, we illustrate how differences in participation rates and capacity constraints can result in uneven treatment effects across school choice programs, and suggest an approach of indexing results across contexts. Simulations also reveal how a high emphasis on school achievement by individuals choosing schools can paradoxically constrain district-level achievement by limiting the number of new schools that survive in a district.

1 Introduction

Much controversy surrounds choice-based reforms in education. On one hand, proponents of choice-based reforms claim that giving parents the ability to choose the school their children attend provides both access to better schooling for disadvantaged populations as well as the incentives necessary for school reform (Chubb & Moe, 1990). On the other hand, opponents of school choice claim that choice-based programs will not bring about the hoped for improvements in schools, but instead only drain resources from troubled schools that can least afford to lose them (Fiske & Ladd, 2000).

Unfortunately, empirical research on the causal effect of choice programs has yet to resolve the controversy. One line of research investigates whether the students who attend "choice" schools (e.g, private schools, public magnet schools, charter schools) are better off than those who do not.

*Ford Motor Company Center for Global Citizenship, Kellogg School of Management, Northwestern University

†School of Information, University of Michigan

‡School of Education, University of Pittsburgh & Carnegie Foundation for the Advancement of Teaching

§Center for Connected Learning and Computer-Based Modeling, School of Social Policy and Education & McCormick School of Engineering, Northwestern University

Observational studies using a variety of approaches to account for selection bias have resulted in disputed findings, with some finding better educational outcomes for students who attend magnet (Gamoran, 1996) or Catholic schools (Bryk et al., 1993; Coleman et al., 1982; Evans & Schwab, 1995; Morgan, 2001); others finding little to no effect on academic achievement (Alexander & Pallas, 1983; Goldhaber, 1996; Neal, 1997); and yet others raising serious concerns about the approach used to deal with selection bias (Altonji et al., 2002).

Perhaps even more surprisingly, field studies taking advantage of the lotteries put in place to deal with oversubscription to choice programs have not yielded conclusive evidence on the "treatment effect" for choosers. Randomized field trials of pilot voucher programs in Milwaukee (Greene et al., 1997; Rouse, 1998; Witte et al., 1995), New York City, Dayton, and Washington, DC. (Howell et al., 2002), have resulted in effect sizes ranging from small to modest, with debates around methodological issues pertaining to selective attrition in Milwaukee (Witte, 1997), and subgroup definition in New York City (Krueger & Zhu, 2004). Cullen et al. (2006) find that lottery winners in Chicago's public choice program experience improvements in nontraditional outcomes such as reduction in disciplinary incidents, but not in more traditional measures such as test scores and drop-out rates.

Of course, even if one could establish a clear causal relationship between attending choice schools and academic performance, the findings would have to come with the caveat that they are not an evaluation of the systemic effect of choice programs. Improvement could arise from a number of mechanisms that go beyond students sorting themselves into better schools[1], particularly through the competition induced by household choice (e.g., Hoxby, 2003). Indeed, embedded in the logic of school choice reform is a dynamic causal story, which can be stated roughly as follows:

> More choice leads to the movement of students. Schools losing students feel pressure to change in order to attract and keep students, which in turn creates pressure for all schools to change. Choice also creates the opportunity for new schools to enter the market, providing students with a wider range of options, and further increasing the competitive pressure on existing schools. Bad or undesirable schools improve or close; good or desirable ones survive; and new, stable levels of enrollments, school types, and student achievement are reached. If the subsequent resorting of students does not lead to clusters of low achieving students, or if the effects of that clustering are not large, the mean level of student achievement will be higher than before.

[1] A school can be "better" either in an absolute sense, or in the sense that it is a more appropriate match for a given student.

Work attempting to pinpoint the "treatment effect" of choice for choosers provides a valuable input to the causal story – an indication of the extent to which schools differ in their ability to increase student achievement – but it does not address the entire story. Moreover, even on this narrower question of whether current choice schools do a better job of educating students, only by understanding the dynamics that lead to the observed estimates can one reasonably synthesize results from treatment effects studies. For example, one cannot make a valid comparison of the effect size of choice on participants estimated from a large scale choice program like the one in Chicago to the effect size estimated on participants in a small, pilot program like the one in Milwaukee, without knowing the nature and extent of schools' competitive response in each; i.e., it is quite possible that the experience of the control group changed over time in one, narrowing the gap between choosers and non-choosers, but not in the other.

A second line of school choice research attempts to identify a causal effect on system-level performance by exploiting variation in competition across metropolitan areas. These studies typically investigate the association between indicators of competition levels such as market concentration ratios or private school enrollments, and academic outcomes such as test scores and graduation rates. Belfield & Levin (2002) review this literature and find results ranging from no statistical significance to modestly sized associations. The potentially endogenous nature of the competition measures used in this work, however, make pinpointing the causal effect difficult, and have led to substantial controversy over the techniques used to ameliorate such concerns (e.g., Hoxby, 2005; Rothstein, 2005). But again, even with greater confidence in the magnitude of the association between system-level outcomes and competition measures, we would still need to know more about the mechanisms responsible for the association in order to evaluate the entire causal story. For instance, is the aggregate relationship a consequence of threatened schools improving, new and better entrants displacing poor performing schools, students sorting more beneficially across schools, or some combination? Since different choice-based policy prescriptions address each of these mechanisms to a different degree, the answer to this question is at least as important as the estimation of a causal effect.

In this paper, we take a step towards directly investigating the dynamic processes related to school choice reform through the use of agent-based modeling (ABM) – a computational approach often used in complex systems research to study the system-level patterns that emerge over time from the interactions of heterogeneous agents (Epstein & Axtell, 1996; Miller & Page, 2007; Wilensky, 2001). Though more common in the traditional sciences, agent-based modeling cuts across disciplinary boundaries and is becoming increasingly used in social science (e.g., Miller & Page, 2007; Axelrod, 1984; Epstein & Axtell, 1996; Tesfatsion & Judd, 2006). Typically, agent-based models are comprised of three components: agents, environment, and rules (see Introduction, Epstein & Axtell, 1996; Wilensky & Rand, 2009). In the context of a social system, *agents* represent people with heterogeneous attributes. The attributes can be fixed characteristics like race and gender, or attributes that can change over time such as wealth and economic preferences. The environment often takes the form of a lattice of geographic sites, or *patches*; a set of *links* between agents representing social relations; or both. For example, a patch might represent a geographic census block and have an attribute that captures the economic resources of that block. A link might represent a friendship that exists between two people and have an attribute that denotes the strength of that friendship. *Rules* govern the behavior of the agents and the environment. An example of an agent rule might be something like "always purchase the bundle of goods that maximizes your utility." For a patch, a rule might govern some underlying rate of growth or decay of the resources at that site. Rules can also govern the interaction between agents and their environment (e.g., "move to the patch with the best schools"), as well as the formation of new links (e.g., "form a friendship with a friend of a friend").

By varying the rules of the agents and the environment, such a framework can be used to study a large range of questions related to school choice reform. In this paper, we focus on gaining a better understanding of (i) the factors that give rise to different treatment effects between chooser and non-choosers, and (ii) the sensitivity of district-level outcomes to differences in student preferences. More specifically, we develop a model of a school district's transition from a catchment area system to a common form of school choice known as "open enrollment", where households can choose among existing public schools in the district, but do not receive vouchers to go outside of the public system. The agents in our model are schools and students. Students vary in ability and background. Schools vary in quality and building capacity. Students rank schools by using a preference function based on the mean achievement and geographic proximity of a school. Their academic achievement results from a combination of individual traits and the "value-added" parameter of the school they attended.

We use data from Chicago Public Schools to initialize the model in a plausible manner.[2] Analysis of the model reveals how differences in participation rates and capacity constraints can result in uneven treatment effects across school choice programs. Simulations also reveal how a high emphasis on school achievement at the household-level can constrain district-level achievement by limiting the number of new schools that can survive in a district – a paradoxical mismatch between micro- and macro-levels of behavior.

While not in the same modeling tradition, our study is related to computational general equilibrium (CGE) models of school finance (Epple & Romano, 1998; Nechyba, 2000; Fernandez & Rogerson, 2003; Ferreyra, 2007) This literature emphasizes that individual reactions to school choice policies can aggregate in ways that impact the larger economic environment (through the housing market, for example), and uses computation to better understand the possible consequences of that aggregation (for review, see Nechyba, 2003). Our study can be viewed as a complement to this work in two broad ways. First, we focus on public school choice – a common variant of school choice reform not directly addressed by most of this work.[3] Second, even beyond the focus on public choice, we tackle a different set of questions than is usually asked in CGE models. CGE models typically calibrate school and student behavior to values consistent to aggregate data, then ask questions about how different school financing systems would impact the level and distribution of school quality. In our model, we initialize students and schools with micro-level data and then ask how educational outcomes under a choice program might differ if the students or schools were to behave differently.

The paper proceeds in four main sections. The first section presents a highly stylized agent-based model of the transition from a catchment area to open enrollment system to illustrate and further motivate the approach. The second section develops a more detailed model of the transition to school choice, using data from Chicago Public School's open enrollment program to initialize parameters, and uses the model to conduct computational experiments similar to the work using randomized lotteries to isolate achievement differences between choosers and non-choosers that can be attributed to the existence of a choice program. The third section extends the previous model by

[2]Although we use data from a particular district, our goal is not to test the outcomes of a specific program or school choice approach.

[3]A notable exception is Epple & Romano (2003), who develop a model where households can choose public schools without changing neighborhoods, and investigate the role of "friction" in that choice (such as transportation costs) in determining the distribution of quality across schools.

allowing new entrants to the district (which one can think of as in a public choice context as charter schools), and then examines the emergence of two district-level outcomes – market concentration and mean achievement. The final section concludes and outlines areas of future work.

2 Model I: An Illustrative Example

In this section, we present a highly simplified agent-based model of a school district which will be elaborated as the paper progresses. Model I (Figure 1) presents a hypothetical school district with two types of agents – schools and students. This hypothetical district is comprised of nine neighborhoods each containing one school, and students residing in a particular neighborhood can only attend their neighborhood school. The goal of Model I is to examine the dynamics of changing from this typical catchment system, to a system where households have choice among all public schools in the district. To provide a simple starting point, Model I leaves out complicating factors such as capacity constraints and heterogeneity in household decision-making and simply asks: What if we could change the world in a way where, from this point forward, all new students were able to attend the school with the highest mean achievement?

2.1 Model Description

The district and its agents. The simulation takes place on a grid of 18 x 18 sites (patches) that represents the geography of the school district.[4] This hypothetical district is divided into nine 6 x 6 neighborhoods, as shown in Figure 2. One school is set in each of the nine neighborhoods of the district, and each school is comprised of four grades, with an equal number of students in each grade. The district is conceived as a public school district, so all schools are tuition-free to all students in model.

In Model I, the district is compromised of 150 total students whose location is determined by a random draw from a uniform distribution of all locations on the grid. In the initial state of the model, all students were required to attend their respective neighborhood school. If given the opportunity to chose, students have perfect information about the mean achievement of each school

[4] All models in this paper were created using the NetLogo Programmable Modeling Environment (Wilensky, 1999). NetLogo comes with an interface to the Mathematica scientific computing environment (Wolfram Research Inc., 2007) which was used to collect and analyze batches of simulation runs under different parameter settings (Bakshy & Wilensky, 2007).

Figure 1: NetLogo Model of School Choice (Model I)

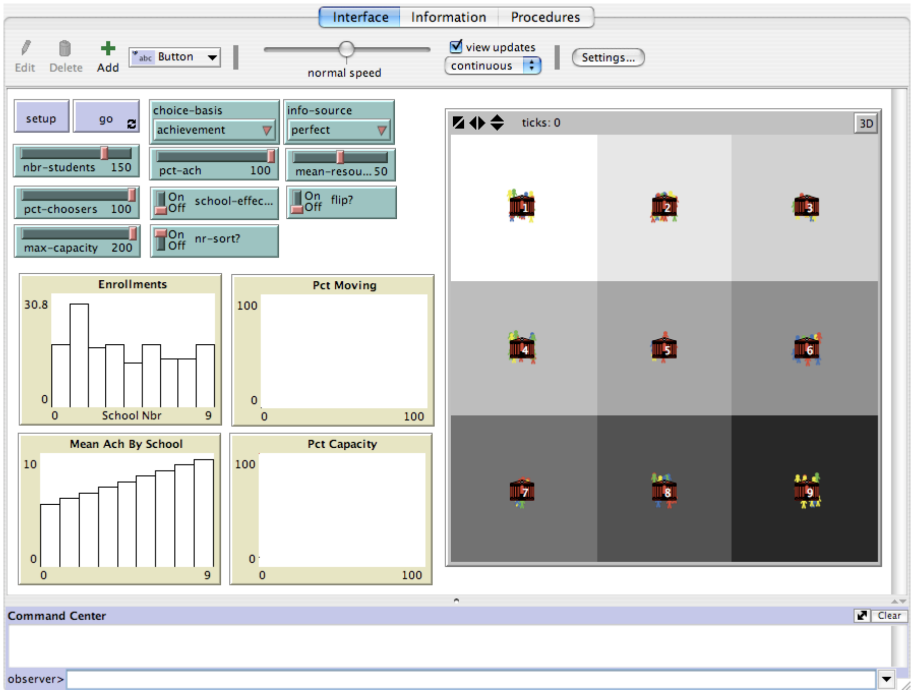

and no mobility constraints.

Each neighborhood differs in its mean level of socioeconomic status. Model I reflects this difference in neighborhoods by initializing students that reside in each neighborhood with differing levels of achievement. The students residing in Neighborhood 9 have the highest mean achievement and the students in Neighborhood 1 have the lowest, with a linear decrease in mean achievement between Neighborhoods 9 and 1. Achievement across students is distributed normally within each neighborhood, allowing for some overlap in distributions between neighborhoods.

A single time period. The simulation takes place in discrete steps, which can be thought of as a time period of one year. Each time period of the model proceeds as follows:

1. New, incoming students to the system must choose a school in the district to attend; one interpretation is that they are 8th graders choosing a high school to attend.
2. The incoming students always choose to attend the school with the highest mean achievement from the last period; they use no other selection criteria are are no longer required to attend their neighborhood school.
3. All schools must accept all students who apply; there are no capacity constraints.
4. Schools update their mean achievement and total enrollment given their new students.

Figure 2: District Geography of Model I

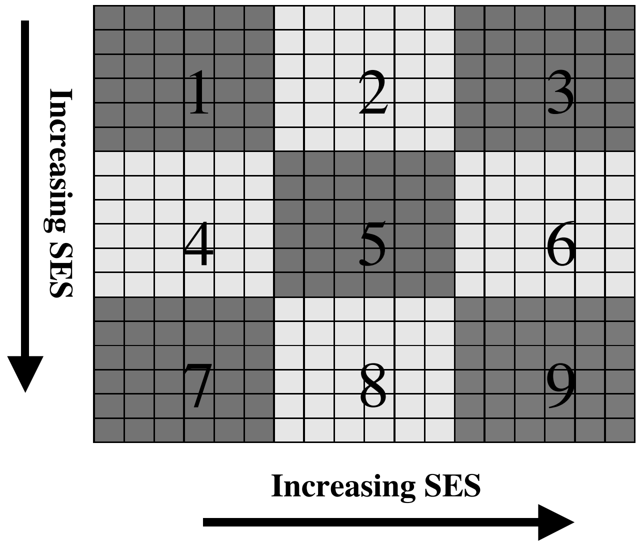

5. Any school whose enrollment drops below three students, permanently closes.
6. Students who have been at a particular school for four years "graduate" or leave the system. They are replaced with an equal number of incoming students who must enroll in one of the nine schools, and whose residence is also distributed randomly across the nine neighborhoods. Students already attending a particular school and not graduated must stay enrolled in the same school.

2.2 Who Survives?

One might think that the results of such a simulation are trivial: if everyone chooses based on achievement, and there is perfect information and mobility, then the school with the highest mean achievement (School 9), will end up with all the students in a just a few time periods. However, Figure 3 shows the frequency of survival by school across 100 simulations. Since there are no capacity constraints, as expected, only one school survives in each run of the simulation. However, aside from Schools 1-5 that never survive, despite being the school with the *highest* initial mean achievement, School 9 is the school that survives the *fewest* number of times.

The explanation for School 9's failure to win on every run is simple, but only after investigating the dynamics of how the results emerge. In a world of perfect information and perfect mobility where everyone chooses, all the new incoming students – including the low achieving ones – rush to the highest achieving school. This lowers the mean achievement of School 9, often enough to give one of the other schools a chance to take the lead. In the next time step, all new incoming students

Figure 3: Frequency of Survival By School (Model I, 100 total runs)

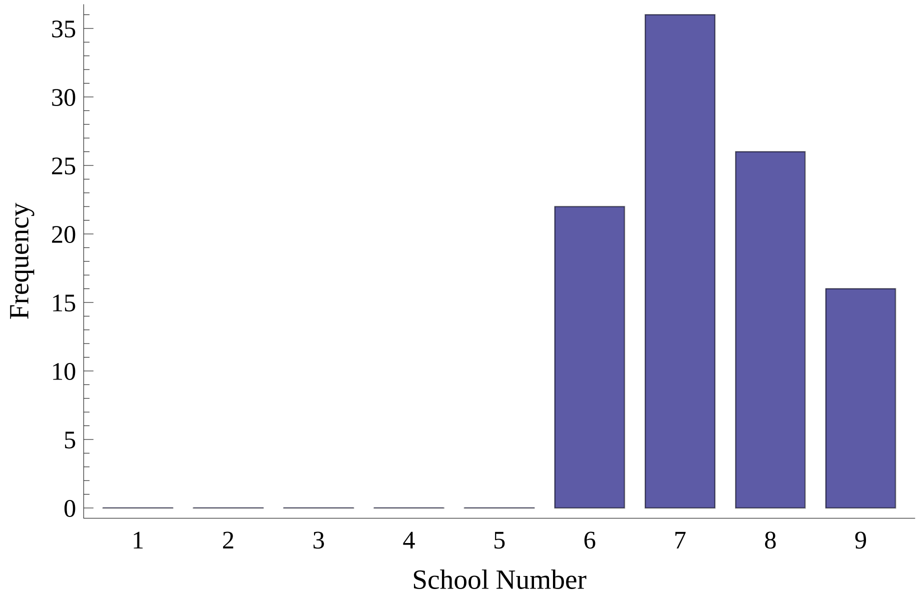

rush to the new leader, producing the same dynamic. Consequently, none of the "top schools are guaranteed a victory.

2.3 Adding School Effects

Currently schools in Model I do nothing to improve the achievement of their students – they simply report the mean of the achievement of the students; there are no "school effects" in the model. An alternative assumption is that schools help produce achievement gains for the students who attend their schools, and differ in their ability to do so. Model I incorporates this notion by giving a school an additional attribute: value-added.

The value-added of a school in Model I is defined as the amount (in standard deviation units) that a student's achievement increases each year they attend a particular school. Like achievement, value-added varies across the neighborhoods; School 9 in Neighborhood 9 is given the highest value-added and School 1 in Neighborhood 1 the lowest, with a linear decrease in between. So now, School 9 is not only the school initialized with the highest achieving students, but it is also that one that indisputably has the highest value-added of any school in the district.

Figure 4 illustrates what happens to school enrollments on a typical run of the simulation with otherwise the exact same rules as before. The lowest achieving schools, Schools 1-6, die out quickly,

Figure 4: Student Enrollments by School (School Effects “On”)

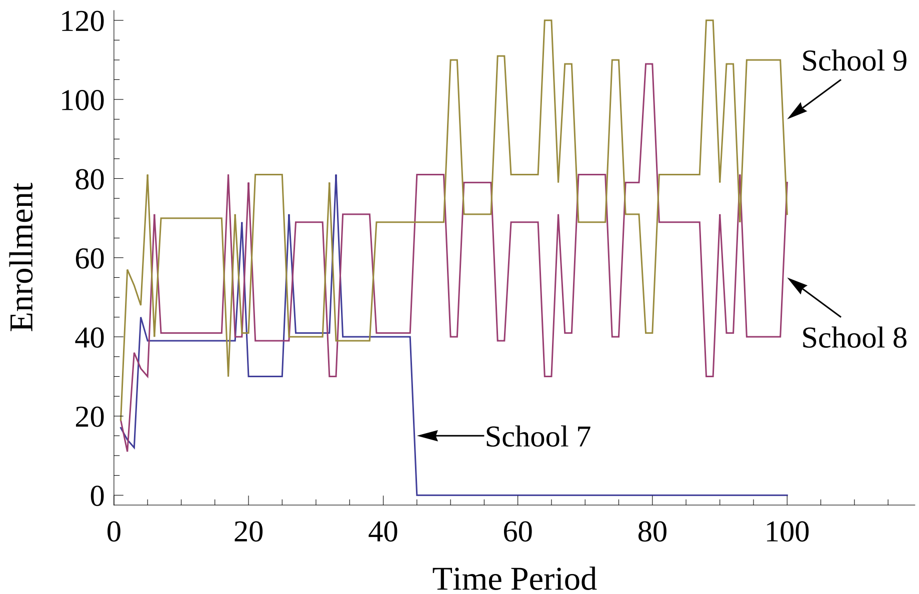

and School 7 remains for approximately 40 time periods. The two highest achieving and highest value-added schools – Schools 8 and 9 – survive and their enrollments fluctuate from year to year.[5] School 9 has a higher initial mean achievement, a higher value-added, and can grow instantly to accommodate all the students in the district. Yet still in one year School 8 has more students, the next few years School 9, then School 8 again, and so on. The reason for this oscillation is the exact same dynamic identified above when there were no school effects – the school with the highest mean achievement draws students that lower its mean achievement, making another school the “best” in the following year. The addition of school effects simply slows down the process enough for one additional school to survive, and biases it towards the higher value-added schools.[6]

In summary, although Model I clearly lacks many mechanisms that operate in a real world school choice program, it nicely illustrates a common occurrence in the study of complex systems that motivates an agent-based modeling approach: even when micro-level behavior is simple and perfectly understood, the emergent macro-level results are still often highly surprising.

[5] Occasionally, School 7 survives and School 8 does not.

[6] In fact, other types of “friction” in the system, like running the simulation with imperfect information about school achievement, would do the same thing.

3 Model II: Initializing with Chicago Public School Data

In Model II, we made two changes important for trying to understand the outcomes that emerge in public choice programs. First, we introduced a constraint on individual school growth. Second, we ensured that the initialization of the model parameters, such as student achievement and school value-added, corresponded with a set of reasonable, real world values – in this case, data on students and schools from Chicago Public Schools open enrollment program during the years of 2000 - 2003. We then used the model to conduct computational experiments comparing the differences between student "choosers" and "non-choosers," in a manner similar to current literature evaluating pilot choice programs.

3.1 Model Description

Model II builds on the same basic idea as Model I: There are two set of agents – students and schools – who operate on a landscape that represents the geography of the district. The simulation begins in a state where all students attended their assigned neighborhood school, and the first time period of the simulation represents the first year when students can choose. Although the model could represent the school choice decision at any level, we treat the schools as high schools, and the students as 8th graders choosing among the schools. An overview of the model is presented first, followed by the precise specification of the agent rules in subsequent sections.

Each time period of the simulation proceeds as follows:

1. The model is populated with 5000 incoming students, and a fraction of them are randomly designated as "choosers"; the fraction is determined by the tunable parameter, $pctChoosers$
2. The "choosers" rank schools in accordance to their preferences, and in random order attempt to attend their top choice school; the remainder of the students attend their assigned neighborhood school.
3. If there are no available spaces at the student's top choice, the student attempts to attend the next school on their ranked-list, and continues to try schools until finding one with room. Regardless of availability, a student's assigned neighborhood school must accept them.
4. Students update their achievement level; the updated achievement depends both on the student's individual-level attributes and the value-added of the school they attend.
5. Schools update their aggregate enrollment and achievement values; they also estimate the number of spaces available for new students next year.

6. Schools that do not meet a minimum threshold of enrollment are permanently closed.
7. Students completing their fourth year in a school, graduate from the system; a student stays at the same high school all four years.

3.2 Student Attributes and Initialization

To populate the model with new incoming students in each time period, Model II samples from a set of 17,131 incoming high school students identified from data made available by Chicago Public Schools (CPS).[7] The goal was to create initial conditions that resembled a district where students attended assigned, neighborhood schools so students who attended magnet and selective schools where not included.

A given student i in the simulation has the following attributes, corresponding to the actual CPS student data:

- $\{xcor_i, ycor_i\}$: the geographic location based on the student's home census block
- $achiev_{1i}$: incoming achievement taken from the student's 8th grade score on the math portion of the Iowa Test of Basic Skills; the scores were standardized across all students in the district.
- $achiev_{2i}$: current achievement as determined in the simulation
- $poverty_i$: the concentration of poverty in the student's home census block
- $assigned_i$: neighborhood school assigned to the student by CPS
- $attended_i$: actual school attended by the student as determined by the simulation
- $white_i$: 1 = white, 0 = other race
- $male_i$: 1 = male, 0 = female

3.3 School Attributes and Initialization

The simulated district consists of 43 neighborhood schools, all of whom had students that were a part of the 17,131 student sample. A given school j in the simulation has the following attributes, corresponding to the actual CPS data:

[7] To enter the set of 17,13, a student had to a) attend 9th grade at a CPS high school between the years of 2000 and 2003, b) attend a CPS school in the 8th grade, c) remain in the same home census block between 8th and 9th grade, and d) have valid and non-missing 8th and 11th grade standardized test scores. As it will be made clear below, the fact that students get re-used in this process is not a problem since the key attributes – current achievement and attended school – are determined by the simulation.

- $\{xcor_j, ycor_j\}$: the geographic location based on the school's address
- $enroll_j$: the total number of students enrolled in the school
- ma_j: the mean achievement of all the students that attend the school
- va_j: the value-added of the school; a fixed value estimated from the CPS data
- dc_j: the design capacity of the school's building[8]
- $maxcap_j$: the maximum enrollment capacity of the school; it equals the design capacity multiplied an exogeneous parameter, *expansion*
- $expassgn_j$: the expected number of students assigned to the school expected to actually attend the school in the following time period; it equals the mean of $enroll_j$ from the last four time periods
- $availsp_j$: the number of available spaces
- $closed_j$: 1 = permanently closed, not accepting new students; 0 = open and accepting new students
- new_j: 1 = school created after time 0, 0 = initial school

An approximate value-added of each school was obtained from the CPS data by using the school-level residuals of a hierarchical linear model estimated for the sample of students used in the simulation, as described in more detail in Appendix A. Schools were otherwise initialized by running the model for five time periods, under the condition that students were forced to attend their assigned school. A representative initialization of the district is presented in Figure 5.[9]

3.4 Student Rules

Achievement growth. Current achievement is updated only once for each student, at the completion of the first year at the school, and remains constant after the initial update. The current achievement for student i attending school j is determined using the following equation estimated from CPS data as described in Appendix A:

$$achiev_{2ij} =$$

$$-0.0956 + 0.6794 * achiev_{1i} + 0.1567 * white_i + 0.1151 * male_i - 0.0629 * poverty_i + va_j + r_{ij},$$

[8]Several schools were missing design capacities in the data. In those cases design capacity of the school was estimated as five times the size of the average incoming freshman class to that school.

[9]In 2000 CPS had already had an open enrollment program for quite some time. Consequently, our value-added estimates would reflect improvements made by schools up to that point in response to the program.

Figure 5: Representative Initialization for Models II and III. A circle represents a school. The circle's size is proportional to the school's enrollment, and its color indicates the mean achievement of the school (green = high, red = low). Schools are placed at the geographic location of their address. The small dots represent students, who are placed within their home census block and attend their assigned school.

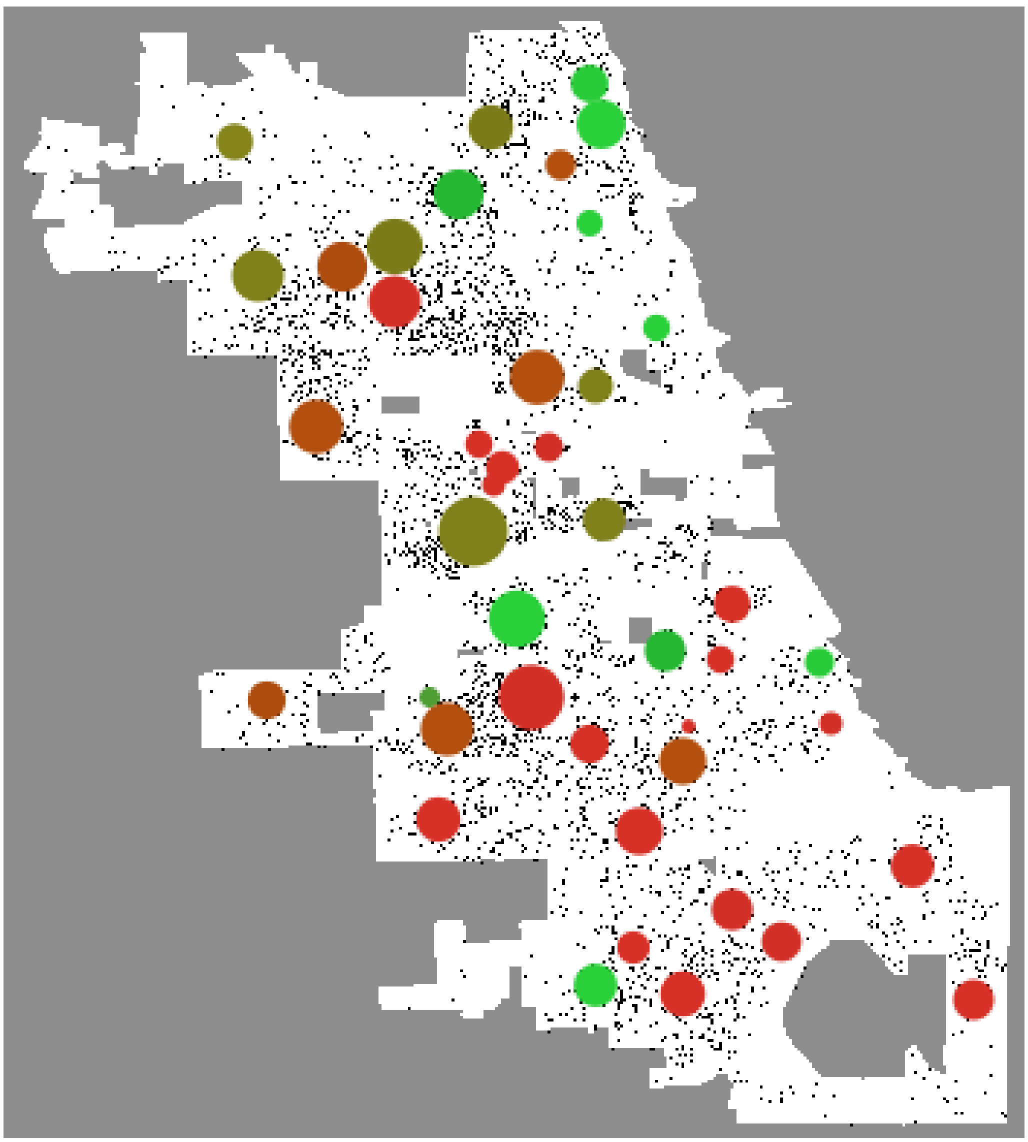

$r_{ij} \sim N(0, 0.3921)$

School choice. Incoming students select a school to attend and remain at the school for four years. Students who are not choosers attend their assigned neighborhood school. Students who are randomly designated as choosers, rank schools using a utility function that trades off between the mean achievement of the school, and its geographic proximity to the student, given by:

$$U_{ij}(ma_j, proximity_{ij}) \quad = ma_j^{\alpha} * proximity_{ij}^{(1-\alpha)}$$

where α is a tunable parameter that alters the weight given to the mean achievement of the school relative to the geographic proximity of the school to the student. Proximity is a function of $dist_{ij}$, the euclidian distance between a student and a school, and is normalized with respect to $dist_{max}$, the maximum distance between two points in the district:

$$proximity_{ij} \quad = \frac{dist_{max} - dist_{ij}}{dist_{max}}$$

After ranking their schools the choosers "apply" to their top choice school. Students apply to schools in random order each time period. If the school accepts them, it decrements the available spaces at that school. If the school does not accept them, students apply to the next school on their utility-ranked list. Regardless of availability, a student's assigned neighborhood school must accept them.

3.5 School Rules

School acceptance policy. Schools must accept students that are assigned to them. Students assigned to other neighborhood schools who apply are accepted on a first come, first served basis, as long as there is space available in the 9th grade. Space is limited by the design capacity of the school, dc, which can be scaled by $expansion$. Additionally, since there is only one round of application and acceptance in each time period, schools must forecast the number of students assigned to them that will actually attend their school. They do so based on the mean of the number of incoming assigned students that attended their school in the previous four years, $expassgn$. The available space at the beginning of any given time period is then given by:

$$availsp_j = \frac{expansion * dc_j}{4} - expassgn_j$$

No other selection criteria are used.

School closure. Schools permanently close if their total enrollment in all grades drops below exogenously determined fraction of their design capacity, ct. The only exception to this is for new schools, who are given a four year "incubation" period during which they do not have to meet the minimum enrollment threshold. Incoming students who are assigned to a closed incoming school, are reassigned randomly to one of the two geographically nearest schools.[10] Unless otherwise noted, for all the results presented in this paper, $ct = 0.05$.

3.6 Computational Experiments and Analysis

In the computational experiments presented in this section, we examined the "treatment effect" of a public choice program – i.e., the difference in achievement between choosers and non-choosers attributable to the being able to attend a school of their own choosing – varied depending on a) the school choice decision rule used by the choosers, and b) the percentage of students who take advantage of the ability to choose. Data were generated by running the Model II for twenty time periods under various combinations of α and $pctChoosers$. The model represents a stochastic process, so each unique combination of parameters was repeated twenty times to create a distribution of outcomes. Unlike Model I, the introduction of capacity constraints stabilizes the year-over-year enrollment levels. Consequently, the achievement comparisons at the end of twenty time periods represent steady-state comparisons.

Figure 6 plots the final mean achievement across all students versus the percent of them who

[10]School closure also creates a situation where a number of existing, upper-class students are left without a school to attend. One could create a rule for reassignment of those students, but in the results presented in this paper those students are assumed to leave the system. For the small closing threshold used in this simulation, this typically results in a very small number of students who leave the system, and subsequently does not alter the results. Additionally, losing the students does not impact the capacity of the system since each school determines available space only with respect to the 9th grade (i.e., only incoming, and not existing, students are used in the available space calculation.)

Figure 6: Mean Achievement vs. Percent Choosers

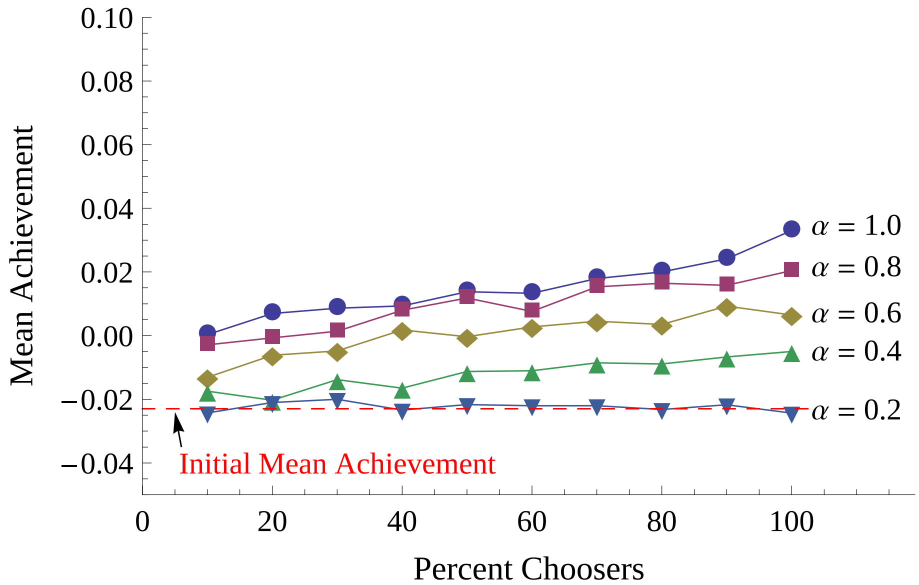

Figure 7: Treatment Effect vs. Percent Choosers

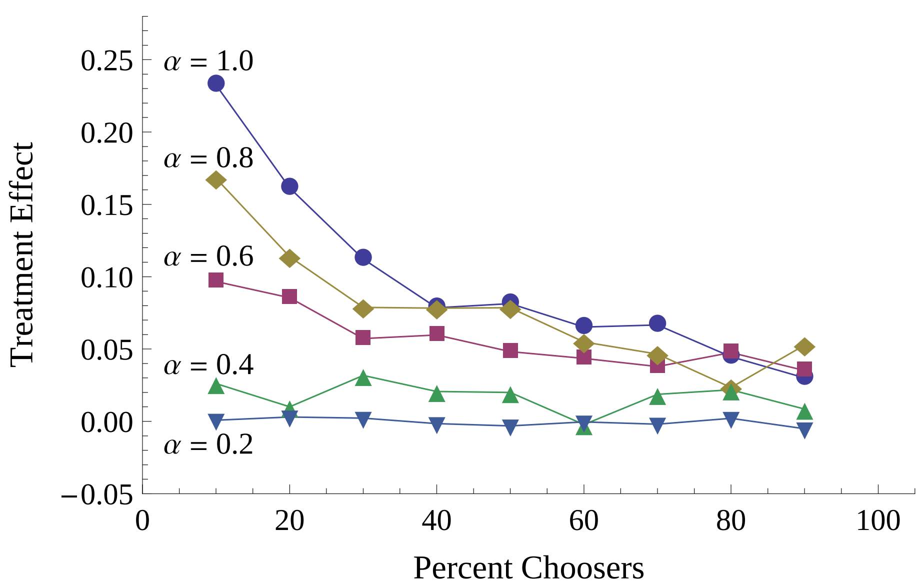

choose, at several different values of α. Figure 7 plots the treatment effect at the completion of the simulation given the same conditions. Since choosers are randomly selected, the treatment effect for any given run of the model can simply be calculated by taking the difference in mean achievement between choosers and non-choosers at the completions of the run. Each point represents the average across the twenty runs for that particular combination of parameters, and mean achievement is measured in standard deviations.[11]

The first thing to note is that the overall change across runs in Figure 6 is rather small. Even in the best case scenario, the difference between the initial mean achievement and the final mean achievement after giving students choice is not greater than 0.05 standard deviations. The reason for this is that for the sample of CPS students used to initialize the model, school effects are only responsible for a small fraction of the overall variance achievement (approximately 5% - 8%, see Appendix A). Stated differently, given the current distribution of value-added across schools – and in particularly the current values at the high end of the distribution – a school's value-added makes a relatively small contribution to a student achievement growth.

When comparing the mean achievement in each of the scenarios in Figure 6, there were no surprises. The more that individual students favored achievement relative to geographic proximity (i.e., high α), the higher the mean achievement. Since mean achievement and value-added are positively correlated, the gain comes straightforwardly from student attending higher value-added schools. Also, with the exception of the case where students heavily favor geographic proximity ($\alpha = 0.2$), the more people who choose, the higher the mean achievement.

However, when comparing the treatment effects across scenarios (Figure 7) notice that the treatment effect often goes down when there is greater participation – the exact opposite relationship than what is observed when one calculated the overall mean achievement of the students under each counterfactual condition.[12] There are two reasons for this, both of which highlight the importance of considering the underlying dynamics of a system when estimating causal effects. The first is that the experience of the control group – the non-choosers – changed. In the simulation, when there is

[11]The test scores where standardized across all students for a particular year of the test, and not only the students included in this sample. The fact that the mean across students in the simulation is slightly less than zero, reflects the fact that their incoming test scores are on average lower than students who attended selective schools, but were not included in the sample.

[12]We also ran the experiment using simulations where students choose schools based on the true value-added of the school, as opposed to the emergent mean achievement. The results presented here are even more pronounced under this condition.

high participation, a few of the low value-added schools close. This necessitates a reassignment of 8th grade students that would have attended the closed schools to new, higher value-added neighborhood schools, making the mean among non-choosers higher. The second reason is less often discussed, but perhaps a more likely real world situation – capacity constraint. More specifically, when there is very low participation and some excess capacity at the better schools, all participants looking for a new school find a spot at one of the highest value-added schools. As more and more people take advantage of the open enrollment option, and the top schools reach capacity, the choosers necessarily have to go to schools that do not have as high value-added (but likely still better than the schools from which they came). Consequently, the mean achievement value of the choosers is lower when more students choose.

To verify the latter reason, we run the Model II 200 times, each time randomly assigning the percentage of students who choose ($pctChoosers$), and the amount of initial excess capacity in the system (governed by the *expansion* variable). These two factors together determine the number of better "seats" available to participants in a choice program, which we capture by creating a measure called better available capacity, or *bac*. More specifically, *bac* is calculated by asking every student who is a chooser to calculate the quantity, $betterSpacesPerSchool$, the number of spaces per school available to them at a schools with a higher value-added. *bac* is the mean of $betterSpacesPerSchool$ across all choosers in the initial time period. As expected, Figure 8 illustrates a clear relationship between *bac* and treatment effects (as measured at the end of each run), with higher levels of better available capacity corresponding to larger differences in achievement between choosers and non-choosers.

4 Model III: Allowing New Entrants

Although Model II presents an environment useful for comparing choosers to non-choosers, such a comparison can only partially inform the larger question of whether increased choice would lead to better district-level outcomes. Two additional mechanisms are key to the causal story. The first is that choice may prompt existing schools implement new and improved ways of educating their students in response to competitive pressures. The second is that choice could enable the entry of new schools in the district with more effective ways of educating students. Indeed, some

Figure 8: Mean Achievement vs. Better Available Capacity

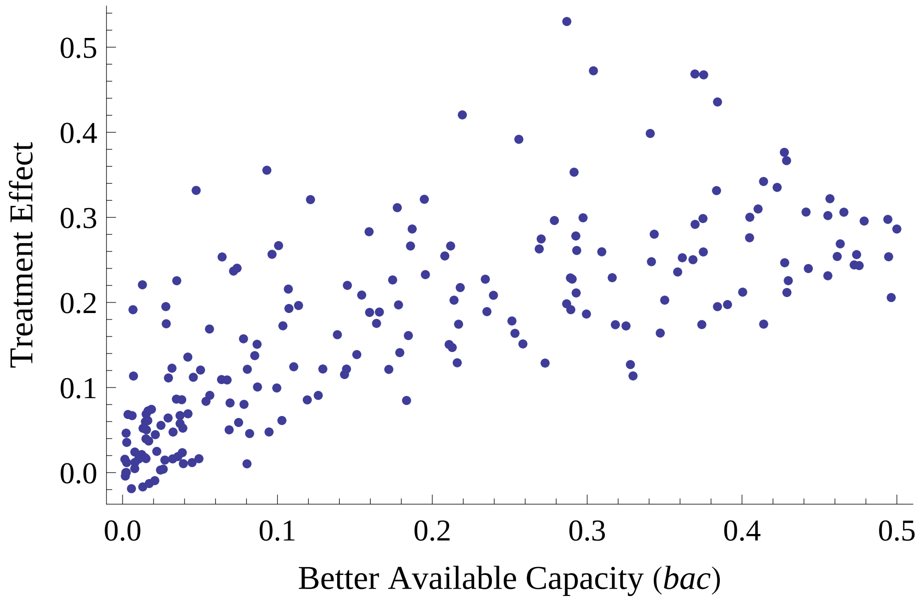

public choice programs – like the one in Chicago – are accompanied by a parallel effort to open charter schools with more freedom to operate their school in a manner the founders see fit. In Model III we introduced the latter mechanism by allowing for new, higher value-added schools to enter the district in each time period (leaving existing-school improvement for future work). Two outcomes are of particular interest. The first is market concentration, which often serves as a proxy for competition in educational markets (Borland & Howson, 1992); the second is the mean achievement of all students in the district.

4.1 Model Description

Model III built on Model II by introducing a "new schools" option. If the new schools option was specified, two new schools tried to open in randomly chosen locations at the end of each time period. Since these schools had no student achievement history, they were initialized with the median value of achievement of all open schools. In subsequent periods, the mean achievement equaled the mean of the students who chose to attend that school. In order to introduce a mechanism for improvement, the underlying value-added of the school was higher on average. More specifically, the value-added for each new school was a random draw from a normal distribution with a mean equal to the mean of the top decile of value-addeds from currently open schools, and a standard deviation equal to the standard deviation of value-addeds across the set of initial schools. One way to interpret this is that the new schools attempt to "copy" the best schools. Sometimes they "innovate" by doing even better than those they copy; sometimes they cannot do as good a job.

4.2 Computational Experiments and Analysis

As in the analysis of Model II, we were primarily interested in how district-level outcomes might vary depending on the school choice decision rule used by the choosers. More specifically, one might expect that an increasing emphasis by students on academic quality (increasing *alpha*) would have a positive impact on both the competitive structure and mean achievement of a school district. Additionally, as Model II illustrates, the amount of better available capacity should play an important role in determining outcomes. Therefore, data were then generated by running the model for thirty time periods under various combinations of α and *bac*. To increase the possibility that new schools will remain open, Model III was run using a relatively high percentage of choosers (80%). Each unique combination of parameters was repeated twenty times to create a distribution of outcomes.

Figure 9 plots the final market concentration in the district against α, across various levels of *bac*. Market concentration is measured using a Herfindahl Index of the sum of the squares of per-school enrollment proportions. Perhaps somewhat unexpectedly, for all levels of *bac* the more a student emphasizes achievement relative to geographic proximity, the higher the resulting market concentration of the district. In the case of plentiful better capacity ($bac = 0.5$), the concentration increases much more than in the case of very little better supply ($bac = 0.1$).

Figure 10 plots the mean achievement of all students in the district against α, across various levels of *bac*. For low levels of α the mean achievement behaves as might be expected – the more a student emphasizes achievement relative to geographic proximity, the higher the resulting mean achievement of the district. As in Model II, this improvement is in part due to students moving to higher value-added, existing schools. In this model, part of the improvement can also be attributed to students attending new schools. However note that again, an increasing emphasis on academic achievement at the individual-level leads to unexpected results at the macro-level. That is, for all levels of *bac*, after a certain point, the more the population of students emphasizes achievement in their decision-making process, the lower the mean achievement of the district.

Figure 9: Market Concentration vs. Emphasis on School Achievement

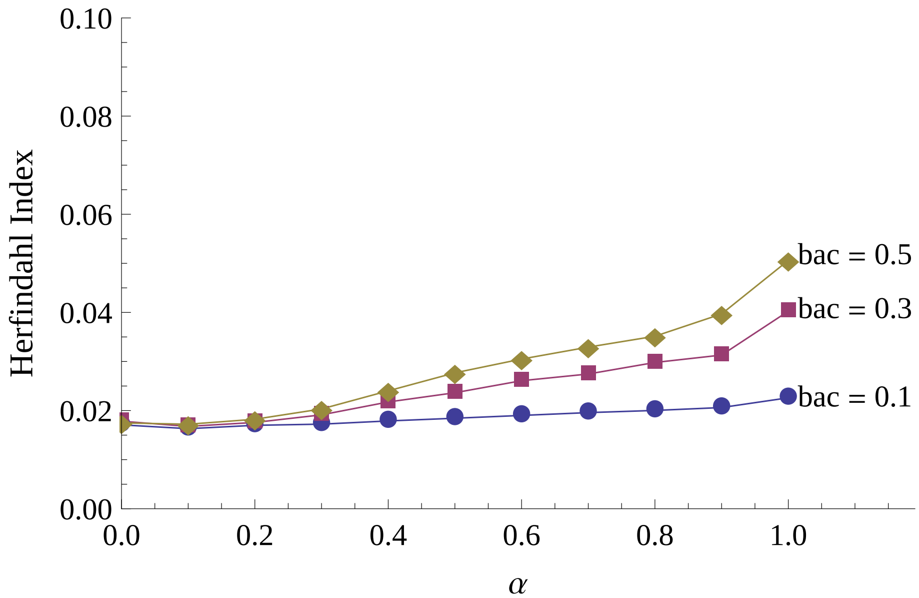

Figure 10: Mean Achievement vs. Emphasis on School Achievement

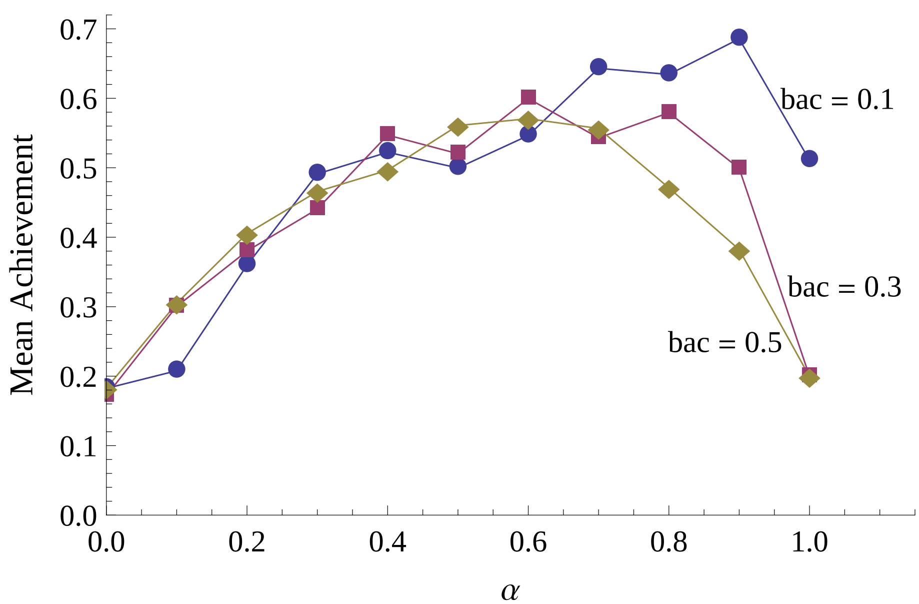

Figure 11 points to the reason for this somewhat surprising behavior in both outcomes: a very high α constrains the number of new schools that can survive. In the left panel, the number of new schools over time for a typical run of the simulation using a relatively low emphasis on school achievement ($\alpha = 0.2$), and the right panel using the maximum emphasis ($\alpha = 1$). For the low α condition, the number of new schools keeps climbing over time, with the opening of each new school bringing about a new opportunity for improvement, as decreasing the concentration of the market structure. However, for the high α condition, the number of new schools grows in the initial time periods, and then levels off.

The reason for this can be understood by contemplating the decision of an individual student, S, when a new school, N, is introduced. Let's suppose that the student's top available choice, School A, is currently a school with an above median mean achievement. Let's further suppose that School N is both geographically closer than School A, and has a much higher value-added (keeping in mind that since the school has no achievement history, in the first time period of its life, School N's mean achievement is perceived as equal to the median achievement of all schools). In the extreme case where $\alpha = 1$, there is no chance that Student S attends the School N over SchoolA even though it is closer and has a higher value-added – i.e., if all that matters in the utility function is the mean achievement of the target schools, Student S will only attend School B if her next best choice is below the median mean-achievement. In subsequent periods, it is possible that School B's mean achievement will rise (especially given the school's higher value-added), and that students with a similar decision may instead choose it over School A. But initially School N has a handicap when compared to incumbent, high mean achievement schools. Moreover, depending on the amount of better available capacity in the district, that handicap may very well lead to its extinction at the end of its initial incubation period. As α decreases, however, School N's attractiveness to Student S and others like her increases, along with the chances for School N's survival. In a sense, lower α makes the demand more "local," leading to more opportunities for the new, better schools. This dynamic occurs regardless of the amount of better available capacity of the district, with the better available capacity influencing exact point at which increasing α begins to lower district achievement (i.e., the peak of each of the curves is different for each level of *bac*).

Figure 11: Number of Schools over Time

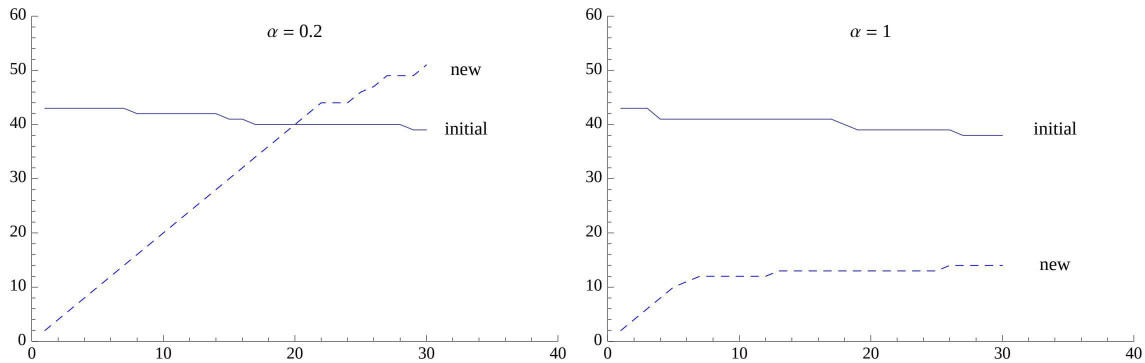

5 DISCUSSION

By developing an agent-based model of public school choice, this paper contributes to our understanding of the dynamic processes that lead to district-level outcomes, and illustrates how such models can help set expectations for typically measured outcomes in school choice research. More specifically, analysis of the model results in two key findings. First, treatment effects calculated by comparing choosers to non-choosers are highly dependent on both the household participation rates in the program and the distribution of available capacity across schools. In particular, as participation rates rise, the magnitude of the treatment effect *falls*, because capacity constraints increasingly limit the amount of choosers who are able to attend the highest value-added schools. The most direct implication of this finding is that one must account for the amount of "better" capacity available to choosers when using treatment effects estimated from existing programs to either a) project the impact of a larger scale program, or b) synthesize effect sizes estimated across programs. The measure we developed to characterize the better available capacity of the district in the model, *bac*, could be applied to program data to aid with both these purposes.

Second, to the extent one were willing to believe that improvement comes largely from new entrants (as opposed to incumbents), and that the quality of the new entrants is initially unknown, our model shows that too much emphasis on academic achievement at an individual-level can constrain district-level improvement by limiting the number of new schools that survive. Furthermore, this dynamic is moderated by the better available capacity in a district. One implication is that districts with a with rather limited amount of better available capacity and a population that places a low to moderate emphasis on academic achievement when choosing schools (relative to geographic proximity), will be helped the most by a public choice program that allows new entrants; and a district with a fair amount of better available capacity and a population that places a very high value on academic achievement, the least. A second implication is that the existence of diverse preferences among households in a district is potentially more important in getting public choice to "work" than a singular emphasis on academic achievement.

Although this study highlights several important implications of taking a more agent-based and dynamic perspective on school choice reform, several limitations of the model must be addressed in order to make more than suggestive statements about the potential impact of a choice program on a given district. First, a much more refined understanding of student preferences is required. More specifically, the current model only partially addresses heterogeneity in the decision-making rules of households. Additional heterogeneity could come in the form of categories of agents weighing elements of the existing preference function differently, or in the form of additional and varied criteria on which to judge schools that go beyond mean achievement and geographic proximity, some perhaps only offered by private schools (e.g., religious education). Second, the model does not yet address the possibility of within-school improvement; educational innovation in the current model comes only from new entrants. A mechanism for within-school improvement could be incorporated in the current model by introducing rule that enabled schools losing students to respond by increasing their value-added, potentially calibrated to be consistent with work attempting to estimate the magnitude of the competitive response (e.g., Figlio & Rouse, 2006). Third, one should be careful in extrapolating results to voucher-based systems in which schools are allowed to price differentially. One implication of incorporating prices is that the role currently served by capacity constraint in the model – namely, providing negative feedback to help equilibrate the system – would be shared by a mechanism that allowed for school-specific pricing. This would likely diminish the importance of the capacity constraints in the current findings. Future work should more carefully address these limitations.

APPENDIX A: ESTIMATING ACHIEVEMENT GROWTH BY SCHOOL

To obtain an estimate of achievement growth for a student attending a particular school in the simulation, I estimate a hierarchical linear model of student achievement that nests students inside of schools. More specifically, I estimated the following model that predicts 11th Prarie State Achievement Examination scores for the all students used in the simulation, using the 8th grade Iowa Test of Basic Skills scores and student-level demographics of those students as the independent variables:

$$achiev_{2ij} = \beta_{0j} + \beta_{1j} achiev_{1i} + \beta_{2j} white_i + \beta_{3j} male_i + \beta_{4j} poverty_i + va_j + r_{ij} \tag{A.1}$$

$$\beta_{0j} = \gamma_{00} + u_{0j} \tag{A.2}$$

$$where\ r_{ij} \sim N(0, \sigma^2)\ and\ u_{0j} \sim N(0, \tau_{00})$$

Table A-1 presents the results of the HLM estimate for both math and reading scores. Substituting Equation A.2 into Equation A.1, and replacing the βs with the estimated coefficients, yields the following equation used as the achievement growth rule in the simulation:

$$achiev_{2ij} = \\ -0.0956 + 0.6794 * achiev_{1i} + 0.1567 * white_i + 0.1151 * male_i - 0.0629 * poverty_i + va_j + r_{ij}$$

For each school j, the school-level residual u_{0j}, is used as an estimate of the value-added, va_j; r_{ij} is a random draw from N(0, 0.3921) every time the achievement growth equation is calculated.

Such an approach assumes that the value-added for each school is relatively stable year over year. To evaluate the stability of the value-added estimates, I also estimate the model separately for each incoming cohort of 8th grade students, and examine the year over year association between the school-level residuals (as opposed to Table A-1 which generates the estimate by using all the cohorts). Figure A-1 shows the that all the year-over-year correlations are strong and positive.

Table A-1: HLM Estimates of 11th Grade Test Scores
(17,131 students in 43 schools)

		Math		Reading	
		Estimate	S.E.	Estimate	S.E.
Fixed Effects					
Intercept	(γ_{00})	-0.0956	0.0301	-0.0025	0.0247
8th Grade Iowa Score	(β_{1j})	0.6794	0.0053	0.5951	0.0062
White	(β_{2j})	0.1567	0.0183	0.1798	0.0211
Male	(β_{3j})	0.1151	0.0097	-0.1180	0.0112
Poverty	(β_{4j})	-0.0629	0.0080	-0.0159	0.0091
Random Effects					
School-level residual variance	(τ_{00})	0.0363		0.0230	
Student-level residual variance	(σ^2)	0.3921		0.5252	

Figure A-1: Association Matrix of School Level Residuals. The upper half of the matrix contains Spearman rank correlation coefficients between the school-level residuals estimated for each cohort of incoming freshman; the lower half of the matrix shows the same association as a scatterplot; the diagonal contains the distribution of school-level residuals in a given year.

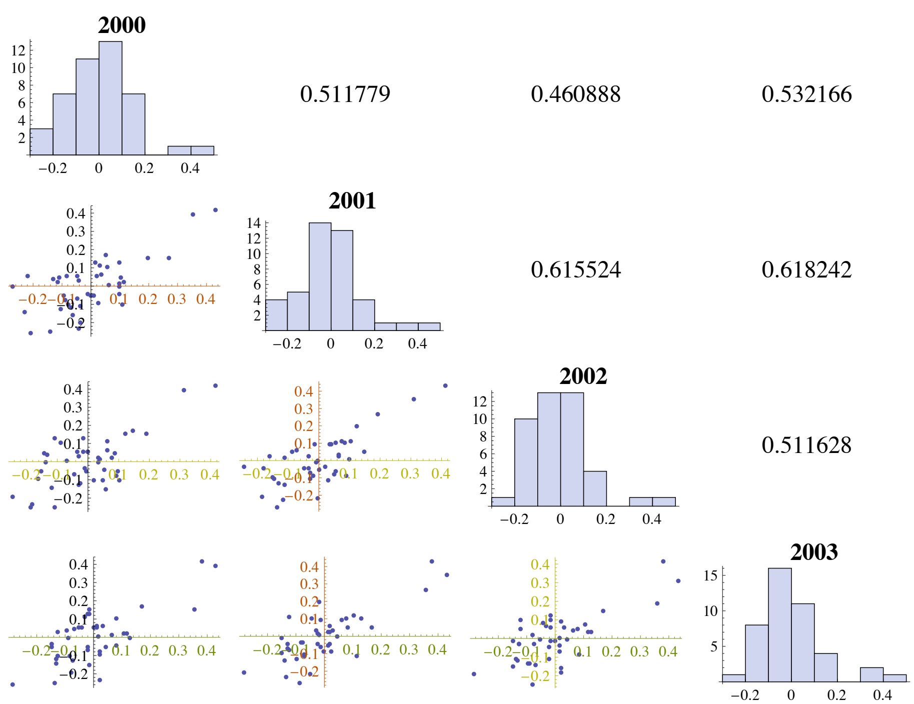

References

Alexander, K. L., & Pallas, A. M. (1983). School sector and cognitive performance: When is a little a little? *Sociology of Education*, *58*(2), 115-128.

Altonji, J., Elder, T., & Taber, C. (2002). *An evaluation of instrumental variable strategies for estimating the effect of catholic schools* (No. 9358). National Bureau of Economic Research.

Axelrod, R. M. (1984). *The evolution of cooperation.* Basic Books.

Bakshy, E., & Wilensky, U. (2007). Turtle histories and alternative universes: Exploratory modeling with netlogo and mathematica. In *Paper presented at agent2007, chicago, november 15-17.*

Belfield, C. R., & Levin, H. M. (2002). The effects of competition between schools on educational outcomes: A review for the united states. *Review of Educational Research*, *72*(2), 279-341.

Borland, M., & Howson, R. (1992). Students' academic achievement and the degree of market concentration in education. *Economics of Education Review*, *11*, 31-39.

Bryk, A. S., Lee, V. E., & Holland, P. B. (1993). *Catholic schools and the common good.* Cambridge: Harvard University Press.

Chubb, J., & Moe, T. (1990). *Politics, markets, and america's schools.* The Brookings Institution.

Coleman, J. S., Hoffer, T., & Kilgore, S. (1982). *High school achievement: Public, catholic, and private schools compared.* New York: Basic Books.

Cullen, J. B., Jacob, B. A., & Levitt, S. (2006). The effect of school choice on participants: Evidence from randomized lotteries. *Econometrica*, *74*(5), 1191-1230.

Epple, D., & Romano, R. (1998). Competition between private and public schools, vouchers, and peer-group effects. *American Economic Review*, *88*(1), 33-62.

Epple, D., & Romano, R. (2003). Neighborhood schools, choice, and the distribution of educational benefits. In C. M. Hoxby (Ed.), *Economics of school choice* (p. 227-286). Chicago University Press.

Epstein, J., & Axtell, R. (1996). *Growing artificial societies.* MIT Press.

Evans, W. N., & Schwab, R. M. (1995). Finishing high school and starting college: Do catholic schools make a difference? *The Quarterly Journal of Economics*, *110*, 941-974.

Fernandez, R., & Rogerson, R. (2003). Equity and resources: An analysis of education finance systems. *Journal of Political Economy*, *111*(4), 858-897.

Ferreyra, M. (2007). Estimating the effects of private school vouchers in multidistrict economies. *American Economic Review*, *93*(3), 789-817.

Figlio, D., & Rouse, C. E. (2006). Do accountability and voucher threats improve low-performing schools? *Journal of Public Economics*, *90*(1-2), 239-255.

Fiske, E. B., & Ladd, H. F. (2000). *When schools compete: A cautionary tale.* Washington, DC: Brookings Institution.

Gamoran, A. (1996). Student achievement in public magnet, public comprehensive, and private city high schools. *Educational Evaluation and Policy Analysis*, *14*(3), 185-204.

Goldhaber, D. D. (1996). Public and private high schools: Is school choice an answer to the productivity problem? *Economics of Education Review*, *15*(2), 93-109.

Greene, J. P., Peterson, P., & Du, J. (1997). *Effectiveness of school choice: The milwaukee experiment.*

Howell, W., Wolfe, P., Campbell, D., & Peterson, P. (2002). School vouchers and academic performance: Results from three randomized field trials. *Journal of Policy Analysis and Management*, *21*(2), 191-217.

Hoxby, C. M. (2003). School choice and school productivity: Could school choice be a tide that lifts all boats? In C. M. Hoxby (Ed.), *The economics of school choice.* Chicago: The University of Chicago Press.

Hoxby, C. M. (2005). *Competition among public schools: A reply to rothstein (2004)* (No. 11216). National Bureau of Economic Research.

Krueger, A., & Zhu, P. (2004). Another look at the new york city school voucher experiment. *American Behavioral Scientist*, *47*(5), 658-698.

Miller, J. H., & Page, S. E. (2007). *Complex adaptive systems: An introduction to computational models of social life.* Princeton, NJ: Princeton University Press.

Morgan, S. L. (2001). Counterfactuals, causal effect heterogeneity, and the catholic school effect on learning. *Sociology of Education*, *74*, 341-374.

Neal, D. (1997). The effects of catholic secondary schooling on educational achievement. *Journal of Labor Economics*, *15*(1), 98-123.

Nechyba, T. (2000). Mobility, targeting, and private-school vouchers. *The American Economic Review*, *90*(1), 130-146.

Nechyba, T. (2003). What can be (and what has been) learned from general equilibrium simulation models of school finance? *National Tax Journal*, *56*(2).

Rothstein, J. (2005). *Does competition among public schools benefit students and taxpayers? a comment on hoxby (2000)* (No. 11215). National Bureau of Economic Research.

Rouse, C. E. (1998). Private school vouchers and student achievement: An evaluation of the milwaukee parental choice program. *Quarterly Journal of Economics*, *113*, 553-602.

Tesfatsion, L., & Judd, K. L. (2006). *Handbook of computational economics* (Vol. 13). Amsterdam ; Oxford: Elsevier.

Wilensky, U. (1999). *NetLogo.* http://ccl.northwestern.edu/netlogo. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.

Wilensky, U. (2001). Modeling nature's emergent patterns with multi-agent languages. In *Eurologo 2001.* Linz, Austria.

Wilensky, U., & Rand, W. (2009). *Introduction to Agent-Based Modeling: Modeling Natural, Social, and Engineered Systems with NetLogo.* MIT Press.

Witte, J. F. (1997). Achievement effects of the milwaukee voucher program. In *American economics association annual meeting.* New Orleans.

Witte, J. F., Sterr, T., & Thorne, C. (1995). *Fifth-year report: Milwaukee parental choice program.*

Wolfram Research Inc. (2007). *Mathematica, version 6.0.*