Public School Choice: An Economic Analysis

Levon Barseghyan, Damon Clark and Stephen Coate*

November 2014

Abstract

Public school choice programs give households a free choice of public school and provide schools incentives to compete for students. Proponents of these programs argue that by the usual market logic, choice and competition will improve the quality of the education that schools provide. Critics counter that the usual market logic does not translate easily to schools, since households' perceptions of school quality depend not only on the efforts of school personnel but also on the composition of the student body (i.e., households have *peer preferences*). This paper advances this debate by developing and analyzing an economic model of public school choice. To capture the pro-choice argument, the model assumes that a neighborhood enrollment policy that provides schools with no incentives to exert effort is replaced by a prototypical public school choice policy in which households have a free choice of school and schools have incentives to compete for students. To capture the anti-choice argument the model assumes that households have peer preferences. The analysis of the equilibrium of this model generates three findings that highlight potential limitations of choice programs.

*For helpful comments and discussions, we thank Jean-Paul Carvalho, Jon Hamilton, Sergios Skaperdas, Steve Slutsky, David Sappington and various seminar participants. Barseghyan: Department of Economics, Cornell University, Ithaca NY 14853, lb247@cornell.edu. Clark: Department of Economics, UC Irvine, Irvine CA 92617, clarkd1@uci.edu. Coate: Department of Economics, Cornell University, Ithaca NY 14853, sc163@cornell.edu.

1 Introduction

Advocates of market-based school reform argue that when parents can choose schools, and when schools compete for students, then by the usual market logic, the forces of choice and competition will improve the quality of the education that schools provide.[1] This argument can be traced to Friedman (1953), and it has underpinned market-based school reforms - widely known as "school choice" policies - in countries around the world. Some choice policies enhance schooling options outside of the public sector, notably charter schools and private school vouchers. In most countries, including the US, the more common choice policy is public school choice. A variety of programs fall under the public school choice umbrella. The common theme is that parents are given more choice over the public school attended (e.g., by removing residency requirements) and schools are given incentives to compete for these students (e.g., by tying school funding to the number of students enrolled).[2]

Critics argue that public school choice programs will not be the panacea their advocates suggest. In the critics' view, "the economic model of markets does not translate easily into the provision of elementary and secondary education" (Ladd (2002), p.4). One issue is that school quality is multi-dimensional, such that parents' school preferences depend not only on school productivity but also on factors outside of a school's control (e.g., the neighborhood surrounding the school). As Rothstein (2006) notes, "Any factor that leads parents to choose any but the most effective available schools will tend to dilute the incentives for efficient management that choice might otherwise create." (p.1333).[3] A potentially deeper problem is that a school's peer composition constitutes one such factor (i.e., households have preferences over their children's peers). According to Ladd (2002), peer preferences mean that choice will exacerbate educational inequality, as more-advantaged households choose schools that enroll more-advantaged students. She also claims that peer preferences will blunt the incentive effects of choice, since both successful and unsuccessful schools will have limited incentives to improve quality. The concern for successful schools is that quality improvements will adversely affect the school's socioeconomic composition; the concern for unsuccessful schools is that despite quality improvements, they will be unable to attract parents.[4] She concludes that "any productivity

[1] By the "usual market logic", we mean the idea that if consumers can choose, firms have to compete for their business and will therefore take actions to make their products more attractive to consumers. In simple price setting models, this involves setting lower prices. But more generally, such actions can involve improving products via greater attention to detail and efforts to innovate.

[2] Programs differ by the level at which they operate and the regulations that apply if schools are over-subscribed. Programs include the No Child Left Behind (NCLB) provisions that give choice to students in "failing" schools, the choice programs operated by some US school districts (e.g., Charlotte-Mecklenberg) that use lotteries to assign places in over-subscribed schools, the national programs that use distance to assign places in over-subscribed schools (e.g., the UK system described by Chubb and Moe (1992)) and the national programs that allow over-subscribed schools to select students (e.g., the New Zealand program analyzed by Ladd and Fiske (2000)).

[3] Hastings et al. (2006) estimate households' school preferences in a public school choice setting and show that preferences over schools' racial mix weaken demand-side pressure.

[4] She claims that "successful schools will be reluctant to expand if doing so requires lowering the average socioeconomic

benefits are most likely to emerge when competition among schools occurs on a level playing field" (p. 11). Echoing this concern, Rothstein (2006) concludes that "It...seems unlikely that choice programs can produce substantial market pressure...without careful attention in their design to the role of peer groups in parental choices" (p.1348).

The purpose of this paper is to develop and analyze an economic model of public school choice. Thus far, both sides of this debate have based their arguments on informal reasoning. By modeling choice more formally, we aim to shed new light on this issue. Our model features one community divided into two equal-sized neighborhoods, each containing one school. Households differ by their socio-economic status and one neighborhood contains more advantaged students than the other. The model is dynamic and the two schools are infinitely-lived; a new cohort of students is enrolled in every period. Before choice, students attend their neighborhood school. After choice, households can enroll their students in either school, but face a cost of attending the non-neighborhood school. There are no capacity constraints and schools must admit any student that wishes to enroll. It follows that households will choose the non-neighborhood school if the gain in expected school quality exceeds the cost. We assume that expected quality is the quality observed in the previous period, so that households are backward-looking. We assume that quality depends on the efforts of school personnel and on the composition of the students enrolled, so that households have peer preferences. Schools obtain utility from the revenue that comes with enrolling more students. It follows that schools can increase revenue in the next period by exerting costly effort (and thereby increasing quality) in the current period. We study the Markov Perfect Equilibrium of this game. In particular, we derive the conditions that characterize the equilibrium and solve these conditions numerically.

The model incorporates peer preferences and thereby captures the concerns of choice critics. But it should also capture the full force of the pro-choice argument. First, our assumptions ensure that schools face no incentives to exert effort pre-choice. Second, our assumptions introduce a prototypical choice policy that is not compromised by capacity constraints or other limits on parents' ability to choose and schools' incentives to compete. Although no choice programs have all of these features, they generate arguably the best case for choice. As such, it is not surprising that in the special case of our model in which households do not have peer preferences, choice improves the quality of both schools and the welfare of all households. More surprising is that in the general case in which households have peer preferences, our analysis yields three findings that highlight potential limitations of choice programs.

Our first finding concerns the effects of choice on school quality. In particular, we show that choice is associated with a type of equity-efficiency trade-off: in settings in which choice generates large increases

level of their students" (p.7), and that, since "no...educational strategy can make a school with a large proportion of disadvantaged...students look effective" (p.7), these schools "have difficulty competing for students" (p.8).

in average school quality, it increases the gap between the quality of the advantaged and the disadvantaged school. This is because schools have no incentives to exert effort pre-choice and because post-choice incentives are driven by the fight for marginal students - those close to indifferent between the two schools. If these marginal students have higher socio-economic status than average (e.g., because advantaged students are more mobile), schools will fight harder for these students (because they will help schools to increase future enrollment) and hence choice will generate strong effort incentives. Since average effort is the sole determinant of average quality in our model (because composition effects wash out in the aggregate), it follows that choice will drive large increases in average school quality. But choice will also increase the gap between the quality of the advantaged and the disadvantaged schools in this setting, since the mobile advantaged students will leave the less mobile disadvantaged students behind. In settings in which choice will not increase the quality gap between schools (those in which the marginal students have lower socio-economic status than average), it will provide weaker effort incentives and will drive smaller increases in average school quality.

Our second finding is that choice can reduce household welfare. This is because the fight for marginal students is a fair one (i.e., the two schools face approximately the same incentives), such that the two schools exert similar levels of effort post-choice. This implies that in equilibrium, choice is driven by the preference for better peer groups. Peer-driven choice is optimal from the household's perspective but wasteful from the perspective of society, since choice is costly and one household's peer gain is another's peer loss. The waste from peer-driven choice can overwhelm the benefits stemming from increased school quality. This negative welfare conclusion is significant because, while the policy debate and the education literature do not consider welfare, from an economic perspective, it should be the ultimate concern.

Our third finding concerns policies commonly proposed to enhance choice. Choice advocates often acknowledge that choice could exacerbate inequality but claim that this problem can be solved by policies that help disadvantaged students to exercise choice (e.g., transport cost subsidies). We extend our model to analyze choice accompanied by this type of policy and show that it has unintended consequences. In particular, it runs afoul of the equity-efficiency trade-off and weakens the effort incentives provided by choice. We show that policy-makers can do better with a two-pronged approach that combines policies that help disadvantaged students to exercise choice with subsidies for schools to enroll disadvantaged students. Even when accompanied by this two-pronged policy however, choice still runs up against the equity-efficiency tradeoff and can still reduce welfare.

The paper builds on three strands of empirical literature. The first includes studies of how households choose schools in public school choice systems (e.g., Hastings et al. (2006); Burgess et al. (2014)). The main message to emerge from these studies is consistent with our framework: parents weigh the costs of exercising choice against the benefits of higher school quality.[5] The second includes evaluations of the impacts of public school choice on various measures of school quality (e.g., Ladd and Fiske (2000); Lavy (2010)). The evidence (discussed in more detail below) is mixed. This is consistent with the efficiency-equity trade-off that we identify, since this implies that choice effects will be setting-specific. The third concerns how the effects of public school choice are shaped by peer preferences. Although this has not been analyzed formally, Ladd (2002) invokes peer preferences to explain the apparently disappointing effects of public school choice in New Zealand. Similarly, Hsieh and Urquiola (2006) suggest that peer preferences might account for the disappointing effects of a large-scale private school voucher reform in Chile. Rothstein (2006) devises a test of whether parents choose among US school districts (i.e., Tiebout choice) on the basis of their effectiveness. He finds that they do not, and argues that the results are consistent with parents choosing districts on the basis of student composition.

To our knowledge, Epple and Romano (2003) is the only previous theoretical analysis of public school choice. They analyze the equity implications of public school choice in a model in which parents choose schools, choose neighborhoods and vote on taxes, and in which schools are passive. This complements Epple and Romano (1998), which analyzed how peer preferences shape the equity effects of private school vouchers in a model that features passive public schools and private schools that act as profit-maximizing clubs (Scotchmer (1985)). To focus more clearly on school behavior and the efficiency effects of choice, we model the behavior of schools, but abstract from the neighborhood and taxation choices of households. McMillan (2005) models school behavior but considers competition between a single public school (or district) and a competitive private school sector and does not consider peer preferences.[6] MacLeod and Urquiola (2013) consider a specific source of peer preference - labor market signaling. In their model, which is focused at the college level, choice drives non-productive colleges from the market, but only provides strong incentives for student effort when colleges enroll students of mixed ability. Otherwise, wages will depend on the signal sent by college attended, and there are lower returns to student effort. Like other analyses of choice at the K-12 level, we abstract from student behavior.

Our model is related to models in the industrial organization literature that deal with switching costs and network effects (Farrell and Klemperer (2007)). Whereas we consider the impacts of peer preferences,

[5] In these studies, costs are typically measured by the distance from home to school, while quality is typically measured by school-average test scores. In our setup, costs can include psychic as well as transport costs, while school quality is simply the utility associated with attending a particular school.

[6] Instead, he considers how the public school responds to voucher-driven changes in the private school options. He shows that when private school vouchers become more generous, public school effort can fall. The intuition is that a private school voucher can make the outside private school option especially attractive for higher-income students, thereby increasing the effort cost of keeping these students in the public school. As such, the public school can find it profitable to cater to higher-income students before but not after the private school vouchers become more generous. Manski (1992) simulates a model similar to that studied by McMillan but which does feature peer preferences. These are not the main focus of his paper however and so he does not examine how they shape outcomes.

papers in that literature consider the implications of size preferences (i.e., how do two firms compete when consumers' choices depend on each firm's market share). The idea is that consumers want to buy (e.g., telecommunications) products with larger market shares so that they can, for example, share the same network as friends and colleagues. An obvious and important difference is that there is no price-setting in our model. Hence while this literature finds that these size-dependent preferences can generate "fat cat" effects that soften competition between firms and increase price-cost mark-ups, it does not necessarily follow that peer preferences will blunt the incentive effects of public school choice.

The rest of the paper is organized as follows. Section 2 introduces the model. Section 3 defines, characterizes, and solves for equilibrium under choice. Section 4, the heart of the paper, explores how choice impacts school quality and household welfare. Section 5 examines policies that might improve outcomes under choice. Section 6 discusses our findings and Section 7 concludes.

2 The model

The model features a single community with a population of households of size 1. The community is divided into two neighborhoods, A and B, each containing 1/2 of the population. There are two schools serving the community, one in each neighborhood. The school in neighborhood $J \in \{A, B\}$ is referred to as school J. The time horizon is infinite with periods indexed by $t = 1, ..., \infty$.

Households differ in their socio-economic status. There are two types: rich (r) and poor (p). The aggregate fraction of rich households in the community is λ and the fraction in neighborhood A is μ. The parameter μ is greater than 1/2, implying that A is the richer neighborhood. The parameter λ measures *community heterogeneity* with λ equal to 1/2 representing maximal heterogeneity. The parameter μ represents the degree of *neighborhood inequality* with higher values of μ representing greater inequality.[7]

In each period, each household has a child which it must send to one of the two schools.[8] Households care about school quality (as they perceive it) but incur a cost if using the school not in their neighborhood. This cost captures the additional transaction costs arising from using the non-neighborhood school.[9] The cost is denoted c and varies across the households. In our base specification, for both rich and poor households, costs are uniformly distributed on the interval $[0, \overline{c}]$, so that the fraction of households with cost less than or equal to $c \in [0, \overline{c}]$ is $c/\overline{c}$. Later in the paper, we will allow costs to differ across income groups.

[7] Since each neighborhood contains half the population, $\lambda\mu$ must be less than 1/2.

[8] One could imagine a world with turnover in which new households with children enter the community each period or a more static setting in which fixed households have a sequence of children to educate.

[9] These include any additional time taken to travel to school, additional expenses arising from higher transport costs, psychic costs resulting from loss of community, etc.

Letting q_J denote the quality of school J, a household living in neighborhood A with cost c obtains a period payoff q_A from using school A and a payoff $q_B - c$ from using school B. Similarly, a household living in neighborhood B with cost c obtains a period payoff q_B from using school B and a payoff $q_A - c$ from using school A. The quality of school J in period t depends on the effort it exerts in period t and on the fraction of its children from rich households. Thus,

$$q_{Jt} = e_{Jt} + \alpha\lambda_{Jt}, \tag{1}$$

where e_{Jt} is school J's effort in period t, λ_{Jt} is the fraction of its students from rich households, and α is a parameter measuring the importance of peer composition.

The district provides schools with a per-student payment that exceeds the costs that an additional student creates. We normalize the per student surplus to one, so that school $J's$ payoff from period t is given by

$$E_{Jt} - \gamma e_{Jt}^2/2, \tag{2}$$

where E_{Jt} denotes enrollment and γ is a parameter measuring the marginal cost of effort. Schools discount future payoffs at rate β.

2.1 Choice

We are interested in how school choice impacts school quality and household welfare. To create a simple benchmark for comparison, we suppose that before the choice policy is introduced, households must enroll their children in their neighborhood school. With the choice policy in place, households can choose either school in the community, irrespective of the neighborhood in which they live.[10]

Before choice, each school's enrollment consists of the students in its neighborhood and thus is fixed at 1/2. Since enrollment is fixed and effort is costly, schools exert zero effort. Each school's quality is therefore determined by the fraction of its students who are from rich households (see (1)). Thus, in each period school A's quality is $\alpha 2\lambda\mu$ and school B's is $\alpha 2\lambda(1-\mu)$. It follows that a household living in neighborhood A obtains a payoff $\alpha 2\lambda\mu$ from enrolling their child in school A and a household living in neighborhood B obtains a payoff $\alpha 2\lambda(1-\mu)$ from enrolling their child in school B.

Under choice, households have to select their school at the beginning of each period. They cannot know what school quality will be since it depends on schools' effort choices during the period. We assume that households have *myopic expectations*, believing that school J's quality in period t will equal its quality in

[10] We ignore capacity constraints, assuming that both schools can accommodate additional students.

period $t-1$.[11] It follows that a household living in neighborhood A with cost c will choose school A in period t if q_{At-1} is at least as big as $q_{Bt-1}-c$ and school B otherwise. Similarly, a household living in neighborhood B with cost c will choose school B in period t if q_{Bt-1} is at least as big as $q_{At-1}-c$ and school A otherwise.

Letting $\Delta q_{t-1} = q_{At-1} - q_{Bt-1}$ denote school A's period $t-1$ *quality differential*, it follows that under choice, enrollment in school A in period t is given by[12]

$$E(\Delta q_{t-1}) = \frac{1+\frac{\Delta q_{t-1}}{\overline{c}}}{2}. \tag{3}$$

Enrollment in school B is given by $1 - E(\Delta q_{t-1})$. Rich enrollment in school A in period t is

$$E_r(\Delta q_{t-1}) = \begin{cases} \lambda\left(\mu + (1-\mu)\frac{\Delta q_{t-1}}{\overline{c}}\right) & \text{if } \Delta q_{t-1} \geq 0 \\ \lambda\mu\left(1+\frac{\Delta q_{t-1}}{\overline{c}}\right) & \text{if } \Delta q_{t-1} < 0 \end{cases}. \tag{4}$$

Given this notation, the fraction of rich children in school A in period t, $\lambda_A(\Delta q_{t-1})$, equals $E_r(\Delta q_{t-1})/E(\Delta q_{t-1})$ and the fraction in school B, $\lambda_B(\Delta q_{t-1})$, is $(\lambda - E_r(\Delta q_{t-1}))\,/\,(1-E(\Delta q_{t-1}))$. It follows that the quality differential in period t is

$$\Delta q_t = e_{At} - e_{Bt} + \alpha\left(\lambda_A(\Delta q_{t-1}) - \lambda_B(\Delta q_{t-1})\right). \tag{5}$$

It will be convenient to let Δq denote the expected quality differential at the beginning of a period (i.e., the quality differential realized in the previous period) and $\Delta\lambda(\Delta q)$ denote the peer differential between the two schools under choice in that period; i.e., $\lambda_A(\Delta q) - \lambda_B(\Delta q)$. It is straightforward to show that

$$\Delta\lambda(\Delta q) = \begin{cases} \frac{\overline{c}2\lambda(2\mu-1)}{\overline{c}+\Delta q} & \text{if } \Delta q \geq 0 \\ \frac{\overline{c}2\lambda(2\mu-1)}{\overline{c}-\Delta q} & \text{if } \Delta q < 0 \end{cases}. \tag{6}$$

This function, which plays a key role in the analysis, is graphed in Figure 1 for our benchmark parameter values $(\overline{c}, \lambda, \mu) = (1, 0.5, 0.75)$.

The peer differential is increasing at an increasing rate for negative values of Δq and decreasing at a decreasing rate for positive values. To understand its behavior, note that when Δq equals 0, all students attend their neighborhood school and thus the peer differential just equals the difference between the fraction of rich in each neighborhood (i.e., $2\lambda(2\mu-1) = 0.5$). As Δq increases from 0, all students from neighborhood

[11] Concretely, one might imagine households visiting and researching schools in the period before they enroll each child. Effectively, what we are assuming is that households believe that whatever quality they encounter is predictive of what their child would experience if enrolled. Introspection convinced us that this was the natural assumption to make. Initial expectations in period 1 (i.e., q_{A0} and q_{B0}) are taken to be the school qualities prevailing before choice.

[12] This expression assumes that $\Delta q_{t-1} \in [-\overline{c}, \overline{c}]$.

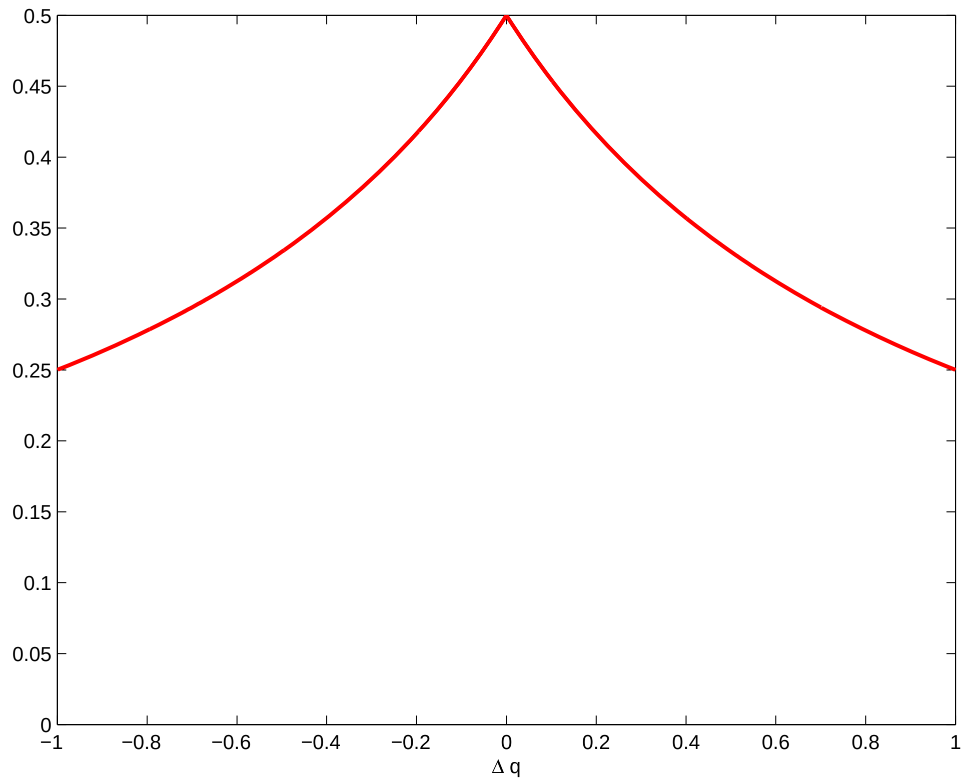

Figure 1: Peer differential

A are enrolled in school A and additional enrolling students come from neighborhood B. Given that neighborhood B is the disadvantaged neighborhood, this reduces the peer differential between the two schools. As Δq approaches $\overline{c}$, more and more students from neighborhood B enroll in school A and the fraction rich approaches that in the whole community. Meanwhile, the fraction rich in school B remains constant at the fraction in neighborhood B. Thus, the difference approaches $\lambda(2\mu - 1) = 0.25$.

2.2 Discussion of the model

In developing a model in which to analyze public school choice, we faced several modeling choices. To help readers understand the choices made, we briefly discuss the main ones here.

A first question was how to model the basic forces driving outcomes under choice; that is, the fundamental incentives facing households and schools. With respect to households, it seemed natural to assume that households must incur costs when choosing their non-neighborhood school and that these costs would vary across households. This is consistent with Epple and Romano (2003) and the empirical literature on how households choose schools (Hastings et al. (2006), Burgess et al. (2014)). It also seemed sensible to abstract from households' choice of neighborhood and of school financing because this would substantially complicate an already sophisticated model.

With respect to schools, it seemed clear we needed schools to value enrollment since, without this, choice would not be effective. To get enrollment to matter, we assume that funding depends on enrollment and the additional funds provided exceed the costs of higher enrollment.[13] We have in mind that surplus funds can be used to provide things that school personnel value. Examples might include bonus payments, ancillary staff, professional development courses, field trips, classroom equipment, sports programs, musical activities and so on. McMillan (2005) and Manski (1992) make similar assumptions when analyzing public school responses to private school vouchers. Similarly, it seemed clear that we should ignore school capacity constraints and the possibility that school personnel are intrinsically motivated to provide quality. These assumptions, along with the assumption that the cost to households of exercising choice are not too large, generate the best-case scenario for choice.

A second question was how to model household heterogeneity and households' preferences over schools. We see several advantages to our assumption that there are two types of households and that school quality, as seen by households, is just an additive function of school effort and the fraction of high types enrolled. First, the two-types assumption yields a simple representation of community heterogeneity and neighborhood inequality. Second, assuming that effort and composition are separable and that composition is "linear-in-means" ensures that average school quality is invariant to the distribution of students across schools and just equals the enrollment-weighted average of the effort exerted by the two schools. Although more flexible specifications may be seen as superior, it is not obvious how effort and composition interact in the production of school quality. Nor is it obvious that quality is a non-linear function of average composition, or that there is a more natural measure of composition than the average type. Similarly, we see our assumption that all households share the same view of school quality as reasonable because it is tractable and there is no clear-cut evidence on preference heterogeneity.[14]

It should be stressed that we are modelling school quality as perceived by households not school quality as measured by test scores (although in many studies school quality as measured by test scores is also assumed to depend additively on school effort and student composition). Even if there are no peer effects in the production of test scores (see Angrist (2013) for a recent review of the evidence) there are many channels through which student composition could affect perceived quality. In particular, there may be peer effects in

[13] An alternative mechanism would be that the salaries of school personnel are tied to enrollment, as they are in some settings (e.g., the UK).

[14] Burgess et al. (2014) do not find strong preference differences by socio-economic status in the UK. Hastings et al. (2006) find stronger preference differences by socio-economic status in the US, in addition to preferences over more than one dimension of household type (income and race).

the production of other outcomes that households care about (e.g., other measures of academic achievement, safety, future labor market networks, social networks and so on). The bottom line is that households will only be indifferent to composition if they value a set of school outcomes that they believe to be unaffected by composition. This is not supported by the evidence on how households choose schools.[15]

A third question was whether to employ a static or a dynamic framework. To study the same forces in a static model, we would have to assume that schools commit up front to providing particular levels of effort and that parents make their enrollment decisions anticipating those of other parents. This is possible, but we find it more natural to assume that parents are backwards-looking in their evaluation of schools and that schools adjust effort on a period-by-period basis anticipating its impact on future enrollment. An intermediate set-up would be a partially-dynamic model in which schools make a one-time and permanent effort choice in the first period and backwards-looking households make enrollment decisions period-by-period in response. This would be more analytically tractable than our fully-dynamic model. However, we felt that the one time effort choice assumption was sufficiently unnatural that we would, in any case, have to check that any conclusions emerging from the partially-dynamic model were robust to relaxing it. This suggested that we might as well tackle the fully-dynamic model from the outset.

3 Equilibrium under choice

Choice creates a dynamic game between the two schools. The state variable in this game is the expected quality differential Δq which determines each school's enrollment. If the schools make effort choices (e_A, e_B) this period, the expected quality differential in the next period $\Delta q'$ will equal $e_A - e_B + \alpha\Delta\lambda(\Delta q)$ which is the realized quality differential this period (see (5)). We will look for a Markov Perfect Equilibrium of this game. In such an equilibrium, schools' strategies depend only upon the state variable Δq. Let $V_J(\Delta q)$ denote school J's value function and $e_J(\Delta q)$ denote school J's strategy. Then, the two schools' strategies satisfy the requirements that

$$e_A(\Delta q) = \arg\max_{e_A} \left\{ E(\Delta q) - \gamma \frac{e_A^2}{2} + \beta V_A(e_A - e_B(\Delta q) + \alpha\Delta\lambda(\Delta q)) \right\}, \tag{7}$$

[15] Mizala et al. (2007) exploit quasi-random variation in the school value-added information (an approximation to school effort) observed by Chilean parents. They find that parents do not respond to better value-added performance, an apparent rejection of the hypothesis that parents care only about value added. Three further pieces of evidence also point away from this hypothesis. First, Hastings et al. (2006) find that test scores do a much better job of predicting preferences than do value-added scores. Second, in a study of school choice in Chicago, Cullen et al. (2006) find that school popularity (measured by applications or the probability that an offer is taken up) is strongly correlated with school test scores but at best weakly correlated with school value added. Third, the above-mentioned Rothstein (2006) study finds no evidence to suggest that parents residing in high-choice areas engage in more "effectiveness sorting". As Rothstein (2006) argues, this is consistent with them having stronger preferences for student composition than for school value added. Finally, to us at least, anecdotal evidence and introspection casts doubt on the proposition that parents would be indifferent between two schools that exerted the same level of effort but enrolled different types of students.

and that

$$e_B(\Delta q) = \arg\max_{e_B} \left\{ 1 - E(\Delta q) - \gamma \frac{e_B^2}{2} + \beta V_B(e_A(\Delta q) - e_B + \alpha \Delta\lambda(\Delta q)) \right\}. \tag{8}$$

Moreover, the two schools' value functions must satisfy the equations

$$V_A(\Delta q) = E(\Delta q) - \gamma \frac{e_A(\Delta q)^2}{2} + \beta V_A(e_A(\Delta q) - e_B(\Delta q) + \alpha \Delta\lambda(\Delta q)), \tag{9}$$

and

$$V_B(\Delta q) = 1 - E(\Delta q) - \gamma \frac{e_B(\Delta q)^2}{2} + \beta V_B(e_A(\Delta q) - e_B(\Delta q) + \alpha \Delta\lambda(\Delta q)). \tag{10}$$

An *equilibrium* consists of value functions $V_A(\Delta q)$ and $V_B(\Delta q)$ and strategies $e_A(\Delta q)$ and $e_B(\Delta q)$ satisfying these four equations.

3.1 Equilibrium characterization

Assuming that the value functions are differentiable, the first order conditions for the two schools' effort choices imply that:

$$\gamma e_A(\Delta q) = \beta \frac{dV_A(\Delta q')}{d\Delta q'}, \tag{11}$$

and that

$$\gamma e_B(\Delta q) = -\beta \frac{dV_B(\Delta q')}{d\Delta q'}. \tag{12}$$

Equation (11) says that for school A, the marginal cost of effort (which is $\gamma e_A(\Delta q)$) is equal to the marginal benefit (which is $\beta dV_A(\Delta q')/d\Delta q'$). This marginal benefit represents the impact of a *higher* expected quality differential next period. Equation (12) says that for school B, the marginal cost of effort (which is $\gamma e_B(\Delta q)$) is equal to the marginal benefit (which is $-\beta dV_B(\Delta q')/d\Delta q'$). The marginal benefit is that stemming from a *lower* expected quality differential next period.

To interpret the first order conditions, more information on the derivative of the value functions is required. Applying the *Envelope Theorem* to equation (9) and assuming that the equilibrium strategies are differentiable at Δq, we have that for school A:

$$\frac{dV_A(\Delta q)}{d\Delta q} = \frac{dE(\Delta q)}{d\Delta q} + \beta \frac{dV_A(\Delta q')}{d\Delta q'} \alpha \frac{d\Delta\lambda(\Delta q)}{d\Delta q} - \beta \frac{dV_A(\Delta q')}{d\Delta q'} \frac{de_B(\Delta q)}{d\Delta q}. \tag{13}$$

Equation (13) reveals that a marginal increase in the expected quality differential has three effects on school A's payoff. The first is to increase its enrollment, an effect which is always beneficial. The second is to change the peer differential. This effect influences school A's payoff via its impact on next period's expected quality differential. Figure 1 tells us that under the base cost specification, the change in the peer differential will be negative if Δq is positive and positive if Δq is negative. A negative change will reduce next period's expected quality differential and thereby reduce school A's payoff, while a positive change will raise it. The third effect is to change school B's effort, an effect that also impacts payoffs through its impact on next period's expected quality differential. The direction of change in school B's effort depends on equilibrium behavior and is unclear *a priori*. A reduction will increase next period's expected quality differential and thereby increase school A's payoff, while an increase will reduce it.

For school B, applying the *Envelope Theorem* to equation (10), we have that

$$-\frac{dV_B(\Delta q)}{d\Delta q} = \frac{dE(\Delta q)}{d\Delta q} - \beta\frac{dV_B(\Delta q')}{d\Delta q'}\alpha\frac{d\Delta\lambda(\Delta q)}{d\Delta q} - \beta\frac{dV_B(\Delta q')}{d\Delta q'}\frac{de_A(\Delta q)}{d\Delta q}. \tag{14}$$

This tells us that a marginal decrease in the expected quality differential has three effects on school B's payoff. The first is to increase its enrollment. The second is to change the peer differential. Again, under the base cost specification, this change will be positive if Δq is positive and negative if Δq is negative. A positive change will increase next period's expected quality differential and thereby reduce school B's payoff, while a negative change will raise it. The third effect is to change school A's effort. An increase will raise next period's expected quality differential and thereby reduce school B's payoff, while a decrease will increase it.

Combining these expressions for the value function derivatives with the first order conditions, we can write (11) and (12) as

$$\gamma e_A(\Delta q) = \beta\left[\frac{dE(\Delta q')}{d\Delta q'} + \gamma e_A(\Delta q')(\alpha\frac{d\Delta\lambda(\Delta q')}{d\Delta q'} - \frac{de_B(\Delta q')}{d\Delta q'})\right], \tag{15}$$

and

$$\gamma e_B(\Delta q) = \beta\left[\frac{dE(\Delta q')}{d\Delta q'} + \gamma e_B(\Delta q')(\alpha\frac{d\Delta\lambda(\Delta q')}{d\Delta q'} + \frac{de_A(\Delta q')}{d\Delta q'})\right]. \tag{16}$$

Moreover, we have that next period's quality differential is given by

$$\Delta q' = e_A(\Delta q) - e_B(\Delta q) + \alpha\Delta\lambda(\Delta q). \tag{17}$$

These are the three equations that characterize the equilibrium effort levels.

It is clear from these three conditions that equilibrium effort levels must in general depend on the state Δq. If they were constant, then $de_A(\Delta q')/d\Delta q'$ and $de_B(\Delta q')/d\Delta q'$ would equal zero, and the effort levels would be equal (i.e., $e_A = e_B$). But then the common effort level would depend on $\alpha d\Delta\lambda(\alpha\Delta\lambda(\Delta q))/d\Delta q'$ which depends on Δq - a contradiction. The only exception to this is when peer preferences are irrelevant, as happens when there are no peer preferences ($\alpha = 0$) or there is no neighborhood inequality ($\mu = 1/2$). In this case, there exists an equilibrium under choice in which the two schools' effort levels are always $\beta/2\gamma\bar{c}$. In this equilibrium, the steady state quality differential will be zero.

It is also apparent from (15) and (16) that the two schools must in general choose *different* effort levels in equilibrium. If $e_A(\Delta q) = e_B(\Delta q)$, then $de_A(\Delta q')/d\Delta q'$ would equal $de_B(\Delta q')/d\Delta q'$. But the derivative of the other school's effort enters with a different sign in each first order condition which is inconsistent with the hypothesis of symmetric effort levels.

Finally, note that the only place that Δq enters these equations is through the peer differential term $\alpha\Delta\lambda(\Delta q)$ in (17). Intuitively, this period's expected quality differential Δq determines this period's peer differential which, along with the two schools' effort levels, determines next period's expected quality differential $\Delta q'$. Observe from Figure 1 that the peer differential is symmetric in the sense that for any Δq in the interval $[-\bar{c}, \bar{c}]$, $\alpha\Delta\lambda(\Delta q)$ equals $\alpha\Delta\lambda(-\Delta q)$. It follows that, assuming there is a unique equilibrium, the equilibrium effort levels for each school J will be such that for any Δq in the interval $[-\bar{c}, \bar{c}]$, $e_J(\Delta q)$ equals $e_J(-\Delta q)$.

Beyond these three points, (15) and (16) do not reveal much about the nature of schools' equilibrium effort levels. The conditions are complex because each school's first order condition includes the derivative of its rival's future strategy. This complexity makes it hard to develop analytical results. Thus, we will solve for equilibrium numerically.

3.2 Steady states

The evolution of the quality differential in equilibrium is described by the difference equation (17). A quality differential Δq_s is thus a *steady state* if

$$\Delta q_s = e_A(\Delta q_s) - e_B(\Delta q_s) + \alpha\Delta\lambda(\Delta q_s). \tag{18}$$

Assuming that the system converges to a unique steady state, we will use steady state outcomes to assess system performance. To simplify notation, we will let e_{As} and e_{Bs} denote the steady state effort levels $e_A(\Delta q_s)$ and $e_B(\Delta q_s)$.

3.3 Solving for equilibrium

The solution procedure is as follows. As noted earlier, the equilibrium effort levels for each school J will be such that for any Δq in the interval $[-\overline{c}, \overline{c}]$, $e_J(\Delta q)$ equals $e_J(-\Delta q)$. Thus, we just need to solve for the strategies on the interval $[0, \overline{c}]$. We first conjecture (i) that the strategies are continuous on $[0, \overline{c}]$ and differentiable everywhere on $(0, \overline{c})$, (ii) that the equilibrium converges to a steady state in which the quality differential lies in the interval $(0, \overline{c})$, and (iii) that equilibrium strategies are such that $\Delta q'$ belongs to the interval $(0, \overline{c})$ for all Δq in the interval $[0, \overline{c}]$. We then substitute equation (17) into the two first order conditions (15) and (16) and note that this yields a pair of differential equations in the unknown effort functions $e_A(\Delta q)$ and $e_B(\Delta q)$ defined on the compact interval $[0, \overline{c}]$. To solve this pair of equations, we approximate the equilibrium effort functions with higher order Chebyshev polynomials. To find the coefficients for these polynomials we use both Galerkin and collocation methods[16] and confirm that the solutions under both methods coincide. Given these solutions, we verify our assumptions (ii) and (iii) ex post (i.e., that the equilibrium converges to a steady state in which the quality differential lies in $(0, \overline{c})$ and that $\Delta q'$ belongs to $(0, \overline{c})$ for all Δq in the interval $[0, \overline{c}]$).

This procedure provides us with conjectured equilibrium strategies on the interval $[0, \overline{c}]$. We then use the symmetry property to deduce the strategies on the interval $[-\overline{c}, 0]$. Given the symmetry property of the strategies and the fact that for any Δq in the interval $[-\overline{c}, 0]$, $\alpha\Delta\lambda(\Delta q)$ equals $\alpha\Delta\lambda(-\Delta q)$, the $\Delta q'$ associated with any Δq in the interval $[-\overline{c}, 0]$ just equals that associated with $-\Delta q$. Accordingly, $\Delta q'$ belongs to $(0, \overline{c})$ and the equilibrium converges to a steady state in which the quality differential lies in $(0, \overline{c})$. Finally, given the strategies, we construct the value functions to verify that (7) and (8) are concave problems and, hence, solving (15) and (16) indeed amounts to solving (7) and (8) for all $\Delta q \in [-\overline{c}, \overline{c}]$.

3.4 Equilibrium for a benchmark case

To illustrate the workings of the model, we now solve for the equilibrium for a benchmark parameterization which assumes that the parameters $(\overline{c}, \lambda, \mu, \alpha, \gamma, \beta)$ equal $(1, 0.5, 0.75, 0.75, 10, 0.95)$. The choice of 0.95 for the discount rate β is standard. The choices of 0.5 for the degree of community heterogeneity λ and 0.75 for the degree of neighborhood inequality μ have no particular significance and we will discuss the implications of different choices below. The choice of 0.75 for the importance of peer composition α is just a starting point and we will discuss the implications of varying it extensively. The remaining parameters, which are the upper bound of the cost distribution $\overline{c}$ and the school effort cost parameter γ will be held constant throughout the

[16] See Judd (1992).

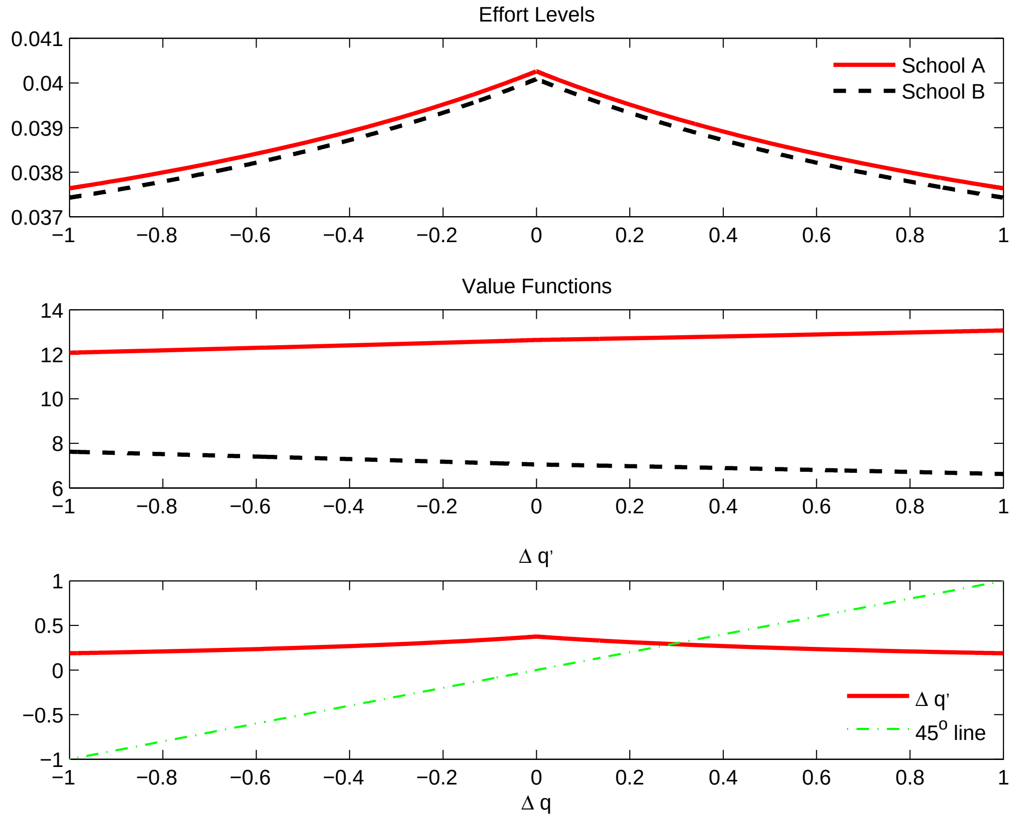

Figure 2: Solution for the benchmark case

analysis.[17]

Figure 2 describes the equilibrium solution. The top panel graphs the effort levels of the two schools as a function of the quality differential. Three points are noteworthy. First, while the equilibrium effort levels do not vary greatly, they are smaller the higher is the absolute value of the expected quality differential. Second, while the two schools' effort levels are similar, for any given quality differential, school A exerts marginally more effort than school B. Third, the two schools' effort levels are symmetric around $\Delta q = 0$.

The second panel of Figure 2 describes the value functions of the two schools. As expected, school A's value function is increasing in the quality differential and school B's is decreasing. While certainly not concave, the value functions are close to linear. When the convex cost of effort is factored in, this means that the optimization problems described in (7) and (8) are concave problems. This implies that the solutions to the first order conditions (11) and (12) are optimal effort levels.

The bottom panel describes the dynamic evolution of the system, graphing how next period's quality dif-

[17]The choice of $\overline{c}$ is a normalization and, once the results are understood, it will be obvious how varying γ will impact the equilibrium.

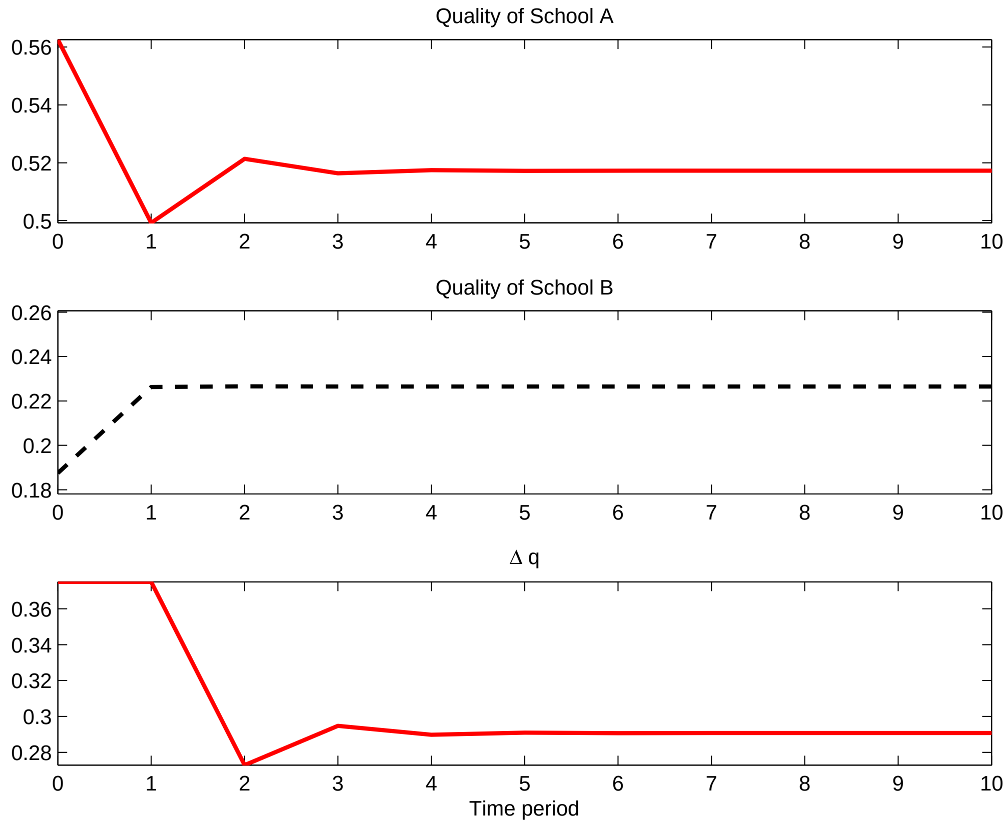

Figure 3: Dynamics for the benchmark case

ferential depends on this period's (i.e., equation (17)). Notice that whatever this period's quality differential, next period's will always be positive. This is because the peer differential is always positive (see Figure 1) and, in addition, school A exerts slightly more effort. As the top panel reveals, the difference in the two schools' effort levels is approximately constant over the entire range of quality differentials. As a consequence, the shape of the function describing next period's quality differential reflects the shape of the peer differential illustrated in Figure 1. The symmetry properties of the two schools' effort levels and the peer differential imply that $\Delta q'$ is symmetric around $\Delta q = 0$. This implies that the dynamic evolution of the system starting from initial condition Δq is exactly the same as that starting from initial condition $-\Delta q$.

The steady state quality differential can be found where the function in the bottom panel intersects the 45^o degree line. It is clear that there is a unique steady state and that the system converges monotonically to it for any given initial condition. The steady state quality differential and associated effort levels $(\Delta q_s, e_{As}, e_{Bs})$ turn out to equal $(0.2907, 0.0392, 0.039)$.

Figure 3 illustrates the dynamic evolution of quality in the two schools starting from the initial condition associated with no choice. Immediately after choice is introduced, quality in school B increases. Thereafter, it varies sufficiently little that it appears flat. This reflects two considerations. First, since equal proportions of rich and poor households are attracted to school A, the fraction of rich children in school B is constant. Second, since the variation in Δq is small along the transition path (as illustrated in the bottom panel), equilibrium effort provided by school B is relatively constant. Even though school A exerts slightly more effort in steady state, quality in school A falls. That is because the fraction of rich students falls as students from the disadvantaged neighborhood enroll. In the first period, quality falls by more than the steady state amount because myopic households overestimate the quality differential between the two schools and too many switch in. In the second period, households overreact to the fall in the quality differential and too few switch in. This adjustment process continues until about the fifth period, when the steady state is reached. As illustrated in the bottom panel, the quality differential is reduced relative to the pre-choice situation but still remains substantial.

To gain some intuition for the equilibrium effort levels described in the top panel of Figure 2, suppose that $\Delta\lambda(\Delta q)$ were piecewise linear rather than the curvy shape illustrated in Figure 1.[18] Then, conditions (15), (16), and (17) suggest the existence of an equilibrium with the following constant and symmetric equilibrium effort level

$$e^* = \frac{\beta}{\gamma 2\overline{c}} \left[\frac{1}{1 - \beta\alpha d\lambda/d\Delta q} \right], \tag{19}$$

where $d\lambda/d\Delta q$ is the constant negative slope of the peer differential on the interval $[0, \overline{c}]$. This candidate equilibrium effort level is increasing in the enrollment return to effort $(1/2\overline{c})$, decreasing in the cost of effort (γ), and decreasing in the absolute value of the composition return to effort $\alpha d\lambda/d\Delta q$.

The nonlinearity of the peer differential accounts for the difference between this constant-symmetric effort solution and the equilibrium effort levels described in the top panel of Figure 2. In particular, equilibrium effort is decreasing in $\triangle q$ on the interval $[0, \overline{c}]$ because next period's quality differential $\triangle q'$ is decreasing in $\triangle q$ and because the composition return to effort is larger when $\triangle q'$ is smaller (i.e., the absolute value of $\alpha d\lambda(\triangle q')/d\Delta q$ is larger as illustrated in Figure 1).[19] Because equilibrium effort is decreasing in $\triangle q$, effort increases by school A (which increase next period's quality differential) elicit a weaker rival reaction than effort increases by school B (which reduce next period's quality differential). This explains why school A exerts slightly more effort than school B. Since this a second-order phenomenon relative to the other forces shaping effort incentives (enrollment effects, effort costs and the first derivative of the composition return), it is not surprising that the equilibrium effort described in the top panel of Figure 2 is approximately constant and symmetric.[20] This is true in every case that we consider in the paper and it underpins many of our

[18] Imagine in Figure 1, two line segments connecting $(-1, 0.25)$ to $(0, 0.5)$ and $(0, 0.5)$ to $(1, 0.25)$.
[19] The first part follows provided effort levels are not "too asymmetric".
[20] Plugging in our benchmark parameters and assuming that $d\triangle\lambda/d\triangle q = -0.5$ (which is the slope of the line segment joining

findings.

Finally, note that this symmetry property explains why a substantial quality differential remains under choice. Because the disadvantaged school is located in the disadvantaged neighborhood and can never turn local students away, it can only neutralize its peer disadvantage by providing considerably more effort than the advantaged school. Such an asymmetry of effort choices is not compatible with the near-symmetry of effort incentives facing the two schools. These incentives imply that the disadvantaged school remains lower quality in equilibrium.

3.5 Equilibrium with cost differences between income groups

As noted in the introduction, it is often argued that more advantaged households are better able to take advantage of choice. In our model, the way to capture this is to assume that rich households face a more favorable cost distribution. The simplest and most tractable way of introducing asymmetries between the groups' cost distributions is to assume that some fraction of each group of households (i_r and i_p respectively) are "immobile". Intuitively, these households are those who will never choose the non-neighborhood school (e.g., because they are unaware of their options, or have prohibitively high costs). By assuming that i_p exceeds i_r, we can examine how equilibrium changes when the rich are more mobile than the poor.

This extension complicates the equations of the model but in a way that is straightforward to analyze. Enrollment in school A is now given by

$$E(\Delta q) = \begin{cases} \frac{1}{2} + \frac{\Delta q}{\bar{c}} \left[(1-i_r)\lambda(1-\mu) + (1-i_p)\left(1/2 - \lambda(1-\mu)\right) \right] & \text{if } \Delta q \geq 0 \\ \frac{1}{2} + \frac{\Delta q}{\bar{c}} \left[(1-i_r)\lambda\mu + (1-i_p)\left(1/2 - \lambda\mu\right) \right] & \text{if } \Delta q < 0 \end{cases}. \tag{20}$$

Note that in contrast to (3) this function is not differentiable at $\Delta q = 0$. Rich enrollment in school A is now given by

$$E_r(\Delta q) = \begin{cases} \lambda \left(\mu + \frac{\Delta q}{\bar{c}}(1-i_r)(1-\mu) \right) & \text{if } \Delta q \geq 0 \\ \lambda\mu \left(1 + \frac{\Delta q}{\bar{c}}(1-i_r) \right) & \text{if } \Delta q < 0 \end{cases} \tag{21}$$

and the peer differential is given by

$$\Delta\lambda(\Delta q) = \begin{cases} \frac{\lambda\left(\mu + \frac{\Delta q}{\bar{c}}(1-i_r)(1-\mu) - E(\Delta q)\right)}{E(\Delta q)(1-E(\Delta q))} & \text{if } \Delta q \geq 0 \\ \frac{\lambda\left(\mu\left(1 + \frac{\Delta q}{\bar{c}}(1-i_r)\right) - E(\Delta q)\right)}{E(\Delta q)(1-E(\Delta q))} & \text{if } \Delta q < 0 \end{cases}. \tag{22}$$

Equilibrium is defined in the same way and equations (15), (16), and (17) still characterize the equilibrium

$(0, 0.5)$ to $(1, 0.25)$), we obtain $e^* = 0.04$.

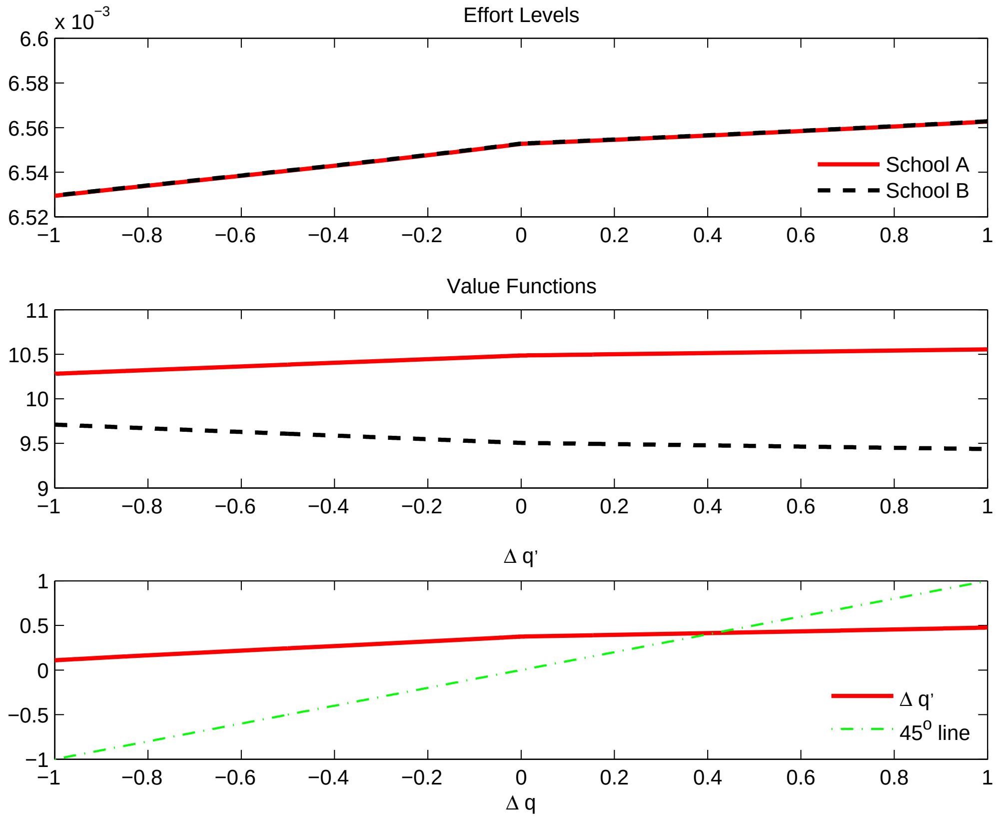

Figure 4: Solution for the assymetric cost case

effort levels. The difference is that the enrollment and peer differential functions are now given by (20) and (22).

To illustrate the implications of rich households being better able to exercise choice, we consider the case in which all the poor and half of the rich are immobile (i.e., $(i_p, i_r) = (1, 0.5)$).[21] The remaining parameters are as for the benchmark case. Figure 4 describes the equilibrium solution. The most striking difference between this and Figure 2 lies in the top panel which describes the equilibrium effort levels. These are *increasing* in the quality differential $\triangle q$ when this is positive as opposed to decreasing in the basic model. Moreover, they

[21] The solution procedure is slightly more complicated because the equilibrium effort levels are no longer symmetric around $\Delta q = 0$. We deal with this by first constructing the equilibrium strategies on the interval $[0, \bar{c}]$ and then using these to construct the strategies on the interval $[-\bar{c}, 0]$. It should be noted that equilibrium will not exist for all possible choices of i_p and i_r. This is because there may exist Δq such that the $\Delta q'$ associated with any candidate equilibrium equals 0. But at that point the functions $E(\Delta q')$ and $\Delta\lambda(\Delta q')$ are not differentiable and thus the first order conditions (15) and (16) are not well-defined. This problem can be circumvented by assuming that i_p and i_r are such that for all Δq it is the case that $\Delta\lambda(\Delta q)$ is positive. Then, provided that the effort levels of the two schools are similar in the candidate equilibrium, (17) implies that $\Delta q'$ is always positive. Intuitively, this assumption means that whatever the expected quality difference between the two schools, after households have made their choices, school A still has better peers. It rules out, for example, a situation in which all the rich are mobile and all the poor immobile (i.e., $(i_p, i_r) = (1, 0)$). In this case, if $\Delta q = -\bar{c}$, all the rich would attend school B and only neighborhood A's poor families would use school A. Clearly, then, $\Delta\lambda(-\bar{c})$ would be significantly negative. By continuity, with only a small difference between the two schools' effort levels, there would exist some $\Delta q \in (-\bar{c}, 0)$ at which $\Delta q'$ as described by (17) would equal 0 and the problem would arise.

are no longer symmetric around $\Delta q = 0$. Underlying this finding is the fact that the peer differential in this case is increasing when $\triangle q$ is positive as opposed to decreasing in the basic model. However, when $\triangle q$ is negative, the peer differential is increasing as in the basic model.[22] Intuitively, because only rich households in neighborhood B are mobile, a higher quality differential attracts more rich families and therefore increases the fraction of rich children in school A. Because the composition return is positive, schools have stronger incentives to exert effort. It is also the case that the effort levels of the two schools are even closer than in the baseline model and it is school B that puts in marginally more effort.[23]

A further difference is evident in the bottom panel which describes the dynamics of the equilibrium. In particular, next period's quality differential is now increasing in the current period's when the latter is positive. Again, since the two school's exert almost identical levels of effort, next period's quality differential is just the peer differential (see (17)). As just discussed, the peer differential is increasing in the quality differential in this case. The change in shape results in the steady state quality differential being higher than in the basic model. The steady state quality differential and associated effort levels $(\triangle q_s, e_{As}, e_{Bs})$ equal $(0.415, 0.006565, 0.006567)$. Effort levels are lower despite the positive effect of the quality differential on the peer differential because so many households are immobile and hence the overall response of enrollment to the quality differential is lower. This underscores the point that the strength of any school effort response resulting from choice should be judged *relative* to the pool of mobile households.

4 The impact of choice

We are now ready to study the impact of introducing choice. We study its impacts on school quality and household welfare. The policy debate and the empirical literature have focused on the quality impacts of choice, both in the aggregate and across schools and communities. The focus on household welfare is more in the spirit of traditional public economics.[24] Measures of household welfare include the additional costs that households incur when they choose their non-neighborhood school, which we view as a legitimate part of the social calculus.

We begin by defining the precise quantities of interest. Recall that without choice, school A's quality in

[22] The shape of the peer differential can be verified by plotting the function described in (22) with $(i_p, i_r) = (1, 0.5)$.

[23] Again, the non-linearity in the peer differential accounts for the difference between these equilibrium effort levels and the constant symmetric effort level. Equilibrium effort levels are slightly increasing in $\triangle q$ because the composition return to effort is positive and because the peer differential function is concave. This means that the composition return is largest when $\triangle q'$ is small (i.e., when $\triangle q$ is large). School B exerts more effort because an increase in school B's quality (which reduces the quality differential) elicits a weaker reaction from its rival.

[24] This said, we do not consider school payoffs in our welfare measure. This is because we see the policy problem to which choice is one possible answer as improving school performance for given levels of educational spending. Eliciting more effort from school personnel is not considered a social loss. In addition, our model of schools' payoffs is too reduced form to permit a satisfactory accounting of the surplus accruing to school personnel and stakeholders.

each period is $\alpha 2\lambda\mu$ and school B's is $\alpha 2\lambda(1-\mu)$. Let $(\Delta q_s, e_{As}, e_{Bs})$ be the steady state associated with the equilibrium under choice and assume that Δq_s is positive. Using (1), the changes in the two schools' qualities are

$$dq_A = e_{As} + \alpha \left(\frac{E_r(\Delta q_s)}{E(\Delta q_s)} - 2\lambda\mu \right), \tag{23}$$

and

$$dq_B = e_{Bs} + \alpha \left(\frac{\lambda - E_r(\Delta q_s)}{1 - E(\Delta q_s)} - 2\lambda(1-\mu) \right). \tag{24}$$

The enrollment-weighted average change in school quality, which we denote by dq, is[25]

$$dq = E(\Delta q_s)e_{As} + (1 - E(\Delta q_s))\, e_{Bs}. \tag{25}$$

The expression in (25) is so simple because changes in school qualities resulting from student composition are zero sum and hence wash out of the analysis.

Turning to welfare, three variables are of particular interest: the steady state average change in per period welfare of households in the two neighborhoods, which we denote dW_A and dW_B, and the average change in per-period welfare, which we denote by dW. The per period welfare change from choice for households in neighborhood A is just

$$dW_A = dq_A. \tag{26}$$

This reflects the fact that, since Δq_s is positive, all households in neighborhood A continue to send their children to school A and hence the only impact on their welfare is how the quality of their school changes. Households in neighborhood B are more complicated because some switch to school A and some do not. The non-switchers obtain a per period welfare change of dq_B. Those who do switch obtain a per period welfare change of $dq_A + \alpha 2\lambda(2\mu - 1) - c$. Averaging over switchers and non-switchers, we obtain

$$dW_B = \begin{array}{c} (1-2\lambda(1-\mu)) \left[i_p dq_B + (1-i_p) \left(\int_0^{\Delta q_s} (dq_A + \alpha 2\lambda(2\mu-1) - c) \frac{dc}{\bar{c}} + (1 - \frac{\Delta q_s}{\bar{c}}) dq_B \right) \right] \\ +2\lambda(1-\mu) \left[i_r dq_B + (1-i_r) \left(\int_0^{\Delta q_s} (dq_A + \alpha 2\lambda(2\mu-1) - c) \frac{dc}{\bar{c}} + (1 - \frac{\Delta q_s}{\bar{c}}) dq_B \right) \right] \end{array} \tag{27}$$

Using (26) and (27), it is straightforward to show that the average welfare gain is

$$dW = dq - (E(\Delta q_s) - \frac{1}{2}) \frac{\Delta q_s}{2}. \tag{28}$$

[25] The enrollment weighted average change in school quality is defined to be

$$dq = E(\Delta q_s)q_{As} + (1 - E(\Delta q_s))\, q_{Bs} - \frac{1}{2} \left(\alpha 2\lambda\mu + \alpha 2\lambda(1-\mu) \right),$$

where q_{Js} is the steady state quality of school J under choice.

This shows that the average welfare gain from choice depends on the difference between two terms. The first term is the change in school average quality, which we know is just the change in school average effort. The second term represents the additional costs incurred by households in neighborhood B who use school A. For choice to generate positive average welfare gains, increases in average quality must outweigh the additional costs incurred by switching households.

We are interested in the sign and magnitude of the six variables defined in (23)-(28) and in how they change with the importance of peer preferences and the mobility of the two groups. With respect to the former, we vary the parameter α from 0 to 1.5. With respect to the latter, we consider two scenarios: the benchmark case from the basic model in which all households are mobile ($(i_p, i_r) = (0, 0)$) and the asymmetric case in which only the rich are mobile ($(i_p, i_r) = (1, 0.5)$). Figure 5 presents our findings. Each of the six panels describes the level of one of the six variables as a function of α which is measured on the horizontal axis. The solid line in each panel describes what happens when all households are mobile, while the dashed line deals with the case in which only rich households are mobile. The parameters underlying the figure are those from the benchmark case.

4.1 School Quality

Our findings concerning the impacts of choice on average school quality are summarized in:

Finding 1 i) *Choice increases average school quality. When all households are mobile, stronger peer preferences reduce the gain in average school quality. When only rich households are mobile, stronger peer preferences increase the gain in average school quality.*

The first part of this finding is unsurprising: average school quality depends solely on average school effort and schools exert zero effort without choice. The second part is more interesting and the intuition is as follows. When all households are mobile, there is a negative "composition return" to higher effort. In other words, all else equal, school A is reluctant to exert effort because this would decrease the peer differential between the two schools; school B is reluctant to exert effort because this would increase the peer differential between the two schools. When peer preferences are stronger, the composition return is larger, hence stronger peer preferences reduce effort incentives and lessen the increase in average school quality. When only rich households are mobile, the composition return is positive and stronger peer preferences increase effort incentives and generate larger increases in average school quality.[26]

[26] Because effort levels depend on the quality differential, there is a second mechanism contributing to this finding. To see this, note that stronger peer preferences magnify the natural advantage of school A and induce some peer-driven choice. In the first case, this reduces the quality differential hence increases school effort (because effort is decreasing in the quality differential - see the top panel of Figure 2). In the second case, this increases the quality differential hence increases school effort (because effort is increasing the quality differential - see the top panel of Figure 4).

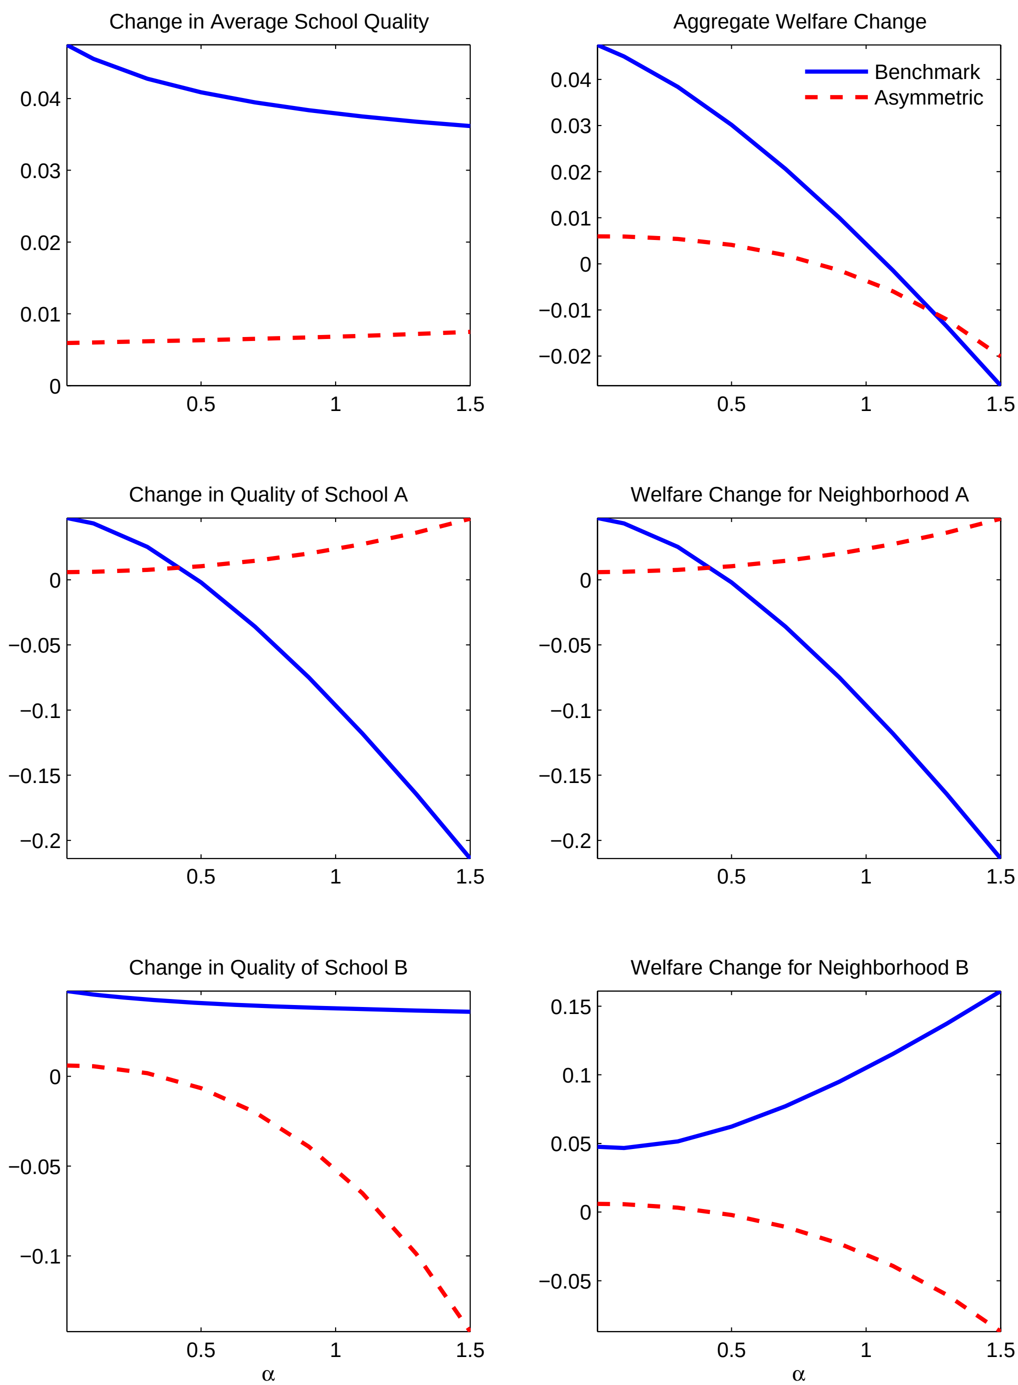

Figure 5: Impact of choice

Turning to the impact of choice on the two schools qualities, our findings are summarized in:

Finding 1 ii) *In the advantaged school, when all households are mobile, choice increases quality when peer preferences are weak but reduces it once peer preferences exceed a critical level. Moreover, stronger peer preferences increase the reduction. When only rich households are mobile, choice increases quality and stronger peer preferences increase this quality gain.*

Finding 1 iii) *In the disadvantaged school, when all households are mobile, choice increases quality but stronger peer preferences reduce this quality gain. When only rich households are mobile, choice increases quality when peer preferences are weak but reduces it once peer preferences exceed a critical level. Moreover, stronger peer preferences increase the reduction.*

Irrespective of which households are mobile, weak peer preferences mean that few households in the disadvantaged neighborhoods exercise choice. It follows that in both schools, choice increases effort without changing school peer groups. As such, choice increases quality in both schools. This is what we would expect to see in a model in which parents did not have peer preferences.

When peer preferences are stronger, households in the disadvantaged neighborhood exercise choice. When all households are mobile, household switching decreases peer quality in school A and leaves it unchanged in school B. Hence for school B, stronger peer preferences reduce the quality gain (because they reduce the effort increase) but the quality gain remains positive. For school A, stronger peer preferences generate a large reduction in the quality gain (because they reduce effort and peer quality) and the quality gain is negative for peer preferences above some critical level. When only the rich are mobile, household switching increases peer quality in school A and decreases it in school B. For school A, this increases its peer advantage and hence its quality advantage. For school B, the costs of worse peer quality can exceed the benefits of higher effort, resulting in a decrease in quality.

Findings 1i, 1ii and 1iii suggest that when households have strong peer preferences, the effects of choice exhibit an interesting type of equity-efficiency trade-off: when introduced into an environment in which rich and poor households are equally likely to take advantage of it, choice will generate modest incentives to provide effort but will decrease the peer quality difference between the advantaged and disadvantaged school; when introduced into an environment in which rich households are more likely to take advantage of it, choice will provide stronger effort incentives, but will increase the peer quality difference between the advantaged and disadvantaged school.[27]

[27] Panel A of Figure 5 suggests that when all households are mobile, choice impacts on average quality are larger than when only rich households are mobile. Although this appears to contradict the equity-efficiency trade-off described in the text, it is a consequence of the fact that when all households are mobile, there are four times as many mobile students as when only rich households are mobile. As such, effort effects on enrollment are four times as large. A proper comparison of the two cases must hold constant the number of mobile households. We make this comparison in Section 6. As seen from panel A of Figure 6, when

4.2 Household Welfare

Our findings concerning average household welfare are summarized in:

Finding 2 i) *Choice increases average welfare when peer preferences are weak and decreases average welfare when peer preferences are strong. The decrease in average welfare associated with stronger peer preferences is greater when all households are mobile than when only rich households are mobile.*

To gain intuition for this finding, recall that the average welfare gain from choice is just the difference between the average quality gain and the costs incurred by the switching households (equation (28)). In the context of Finding 1i, we explained how peer preferences and mobility considerations impact the average quality gain. Thus, we just need to understand the switching costs.

When peer preferences are weak, school A's natural peer advantage has little attraction. As a result, few households will exercise choice and hence these costs will be low. When peer preferences are stronger, school A's peer advantage will drive more households to exercise choice. These switching households obviously benefit from their decision, but their benefit comes at the expense of households in the advantaged school and thus there is no aggregate gain. Welfare is reduced by the costs that these households incur and, for high enough peer preferences, these costs overwhelm the benefits of higher average quality. The decrease in average welfare associated with stronger peer preferences is greater when all households are mobile because in this case average quality is also dropping.

Turning to the impact on welfare in the two neighborhoods, our findings are summarized in:

Finding 2 ii) *In the advantaged neighborhood, when all households are mobile, choice increases welfare when peer preferences are weak, but reduces it once peer preferences exceed a critical level. When only rich households are mobile, choice increases welfare and stronger peer preferences increase this welfare gain.*

Finding 2 iii) *In the disadvantaged neighborhood, when all households are mobile, choice increases welfare and stronger peer preferences increase this welfare gain. When only rich households are mobile, choice increases welfare when peer preferences are weak but reduces it once peer preferences exceed a critical level.*

Since no households in neighborhood A exercise choice, the change in welfare in neighborhood A equals the change in quality in school A. Thus, the intuition underlying the welfare change in the advantaged neighborhood is just that used to explain the quality change.

For neighborhood B, when all households are mobile, stronger peer preferences increase the welfare gain enjoyed by neighborhood B despite decreasing the quality of school B. This is because the benefits enjoyed

the mobile households include a larger fraction of rich students, choice impacts on average school quality are greater. As seen in panels B and C, this case also results in an increase in quality in school A and a decrease in quality in school B. This illustrates the trade-off described in the text.

by those households who switch to school A outweigh the costs (of lower school quality) experienced by those do not. The benefits enjoyed by switching households also explains why, when only rich households are mobile, the neighborhood B welfare reduction associated with stronger peer preferences is milder than the school B quality reduction.

4.3 The Role of Neighborhood Inequality

We have derived all these findings under the assumption that the fraction of rich in the community λ is equal to 0.5 and the fraction of rich in the advantaged neighborhood μ is equal to 0.75. It is worth discussing briefly how changing these values alters the impact of choice.

When all households are mobile, raising the fraction of rich in the community or the fraction of rich in the advantaged neighborhood has exactly the same implications for the impact of choice as increasing the strength of peer preferences. In other words, the impact of changing λ or μ is isomorphic to changing α. To see this, note that the equations characterizing the equilibrium (15), (16), and (17) depend on the peer differential (6) and its derivative multiplied by α, along with the derivative of the enrollment function (3). The enrollment function does not depend on λ and μ, while $\alpha\Delta\lambda(\Delta q)$ just depends on the product $\alpha\lambda(2\mu - 1)$. It follows that the equilibrium under choice and hence the steady state $(\Delta q_s, e_{As}, e_{Bs})$ just depends on $\alpha\lambda(2\mu-1)$. From (25) and (28), this implies that the average change in school quality dq and the average change in household welfare dW just depends on $\alpha\lambda(2\mu-1)$. Moreover, using this information along with (23), (24), (26), and (27), it is easy to verify that the other four variables only depend on the product $\alpha\lambda(2\mu - 1)$. This is the peer preference multiplied by the difference between the fraction of rich students in schools A and B in the pre-choice period. As such, it measures the advantage enjoyed by school A in the competition for students. It is interesting that in the baseline model in which all households are mobile, this is a sufficient statistic for the behavior of the model.

When only rich households are mobile, the equivalence between changing λ or μ and changing α no longer holds. To see this, note that the enrollment function depends on λ and μ and changing each of the three parameters has very different implications for $\alpha\Delta\lambda(\Delta q)$. Thus, to understand the implications of changing λ or μ, further analysis is required.[28] This analysis reveals that an increase in μ (i.e., a decrease in the fraction of the rich households that live in neighborhood B) leads to a reduction in all six variables. An increase in λ (i.e., an increase in the fraction of rich students in the community) has more complex effects, including increasing average quality and decreasing average welfare. None of these effects are especially informative however, since in the scenario where only rich households are mobile, changes to these parameters will

[28] This analysis is available from the authors upon request.

change the number of mobile households in neighborhood B. The resulting effort implications are somewhat mechanical, but will drive most of the effects of changing these parameters. The contrast between the model in which all households are mobile and the model in which only rich households are mobile is much better suited to the analysis of peer preferences.

5 Enhancing school choice

As noted in the Introduction, proponents of public school choice often claim that the equity effects of choice can be improved by policies that help disadvantaged students to exercise choice. In this Section, we consider choice combined with this policy. We also study choice combined with another policy that emerges naturally from our model: subsidizing schools for enrolling disadvantaged students. Since we show that first policy can weaken effort incentives, while the second can strengthen then, we then consider choice combined with a mix of these policies. Throughout, we take a balanced budget approach, requiring that the resources used to implement these policies come from within the school system.

5.1 Helping poor students to exercise choice

To permit a balanced budget analysis of policies that help poor students to exercise choice, we suppose that the school district has a fixed budget available for actions that encourage households in the disadvantaged neighborhood to take advantage of choice. In particular, we assume that these policies reduce the fractions of rich and poor students (i_r and i_p) that are immobile. The idea is that immobility is driven by high costs or ignorance of the choice program and is therefore responsive to policies such as subsidies and outreach. We further assume that by targeting its budget appropriately, the district can influence the extent of the reductions in i_r and i_p. For example, outreach targeted at the less-advantaged part of the neighborhood will cause a bigger reduction in i_p than i_r.

Formally, we assume that the combinations of i_p and i_r that the district can achieve with its budget satisfy an equation of the form

$$(\frac{1}{2} - \lambda(1-\mu))i_p + \lambda(1-\mu)i_r = \frac{k}{2}. \tag{29}$$

The parameter k measures the fraction of households in the disadvantaged neighborhood that are immobile. This will depend on the size of the district's subsidy and outreach budget and is thus exogenous for our purposes. Targeting more of the budget to helping poor students exercise choice will result in a decrease in i_p and a compensating increase in i_r to maintain (29). A key attraction of this formulation is that when we vary (i_r, i_p) in this way, we do not change the "direct" enrollment effect of a change in the quality differential. This can be seen from equation (20). The implicit assumption underlying this formulation is that a redistribution of outreach spending between rich and poor households will not change the overall fraction in neighborhood B that are immobile. Thus, the relationship between outreach dollars and immobility is linear.

We can now examine the implications of targeting resources to help poor students exercise choice. Given (29), such policies are equivalent to an increase in i_r and a compensating decrease in i_p. To illustrate the effect of such a policy, we assume that the district's budget is such that 50% of the households in neighborhood B are immobile after its outreach efforts (i.e., $k = 0.5$). In addition, we assume that under the initial allocation of the budget, each group is equally mobile so that $i_r = i_p = 0.5$. We then consider a retargeting policy that increases i_r to 0.6 and reduces i_p to 0.4666. All other parameters are set at their benchmark levels.

Figure 6 illustrates the effects of this policy. Again, the six panels describe the level of the six variables identified in Section 4 as a function of α. The solid line in each panel describes what happens when $i_r = i_p = 0.5$ and the dashed line deals with the case in which (i_r, i_p) is equal to $(0.6, 0.4666)$. The policy reduces average school quality and quality in school A but increases quality in school B. This reflects the fact that both schools put in lower effort under the policy because a higher fraction of the marginal students are poor. Those households from neighborhood B switching into school A now contain a higher fraction of poor students which hurts school A's peer quality and helps school B's. Mitigating the impact on school A is the fact that there are fewer households switching in because the quality differential between the schools is smaller. The policy has very little effect on aggregate welfare but does lead to a small increase for high peer preferences (i.e., decreases the average welfare loss). This reflects the fact that less switching is occurring. Welfare for the disadvantaged neighborhood is increased. We summarize these findings as:

Finding 3 i) *Relative to a situation in which rich and poor students are equally mobile, targeting resources to help poor rather than rich students exercise choice decreases average school quality and quality in the advantaged school, but increases quality in the disadvantaged school. When peer preferences are strong, the policy increases the welfare gain from choice in the disadvantaged neighborhood, decreases the welfare gain from choice in the advantaged neighborhood and decreases the average welfare loss from choice.*

5.2 Incentivizing schools to enroll poor students

In light of the drawbacks of policies to help poor students to exercise choice, it is natural to instead work on the supply side by encouraging schools to enroll poor students. The basic model assumes that the school district allocates surplus to schools according to the formula of 1 unit per enrolled student. Suppose instead that it provides a surplus of $b_p > 1$ per enrolled poor student and $b_r < 1$ per enrolled rich student. Maintaining

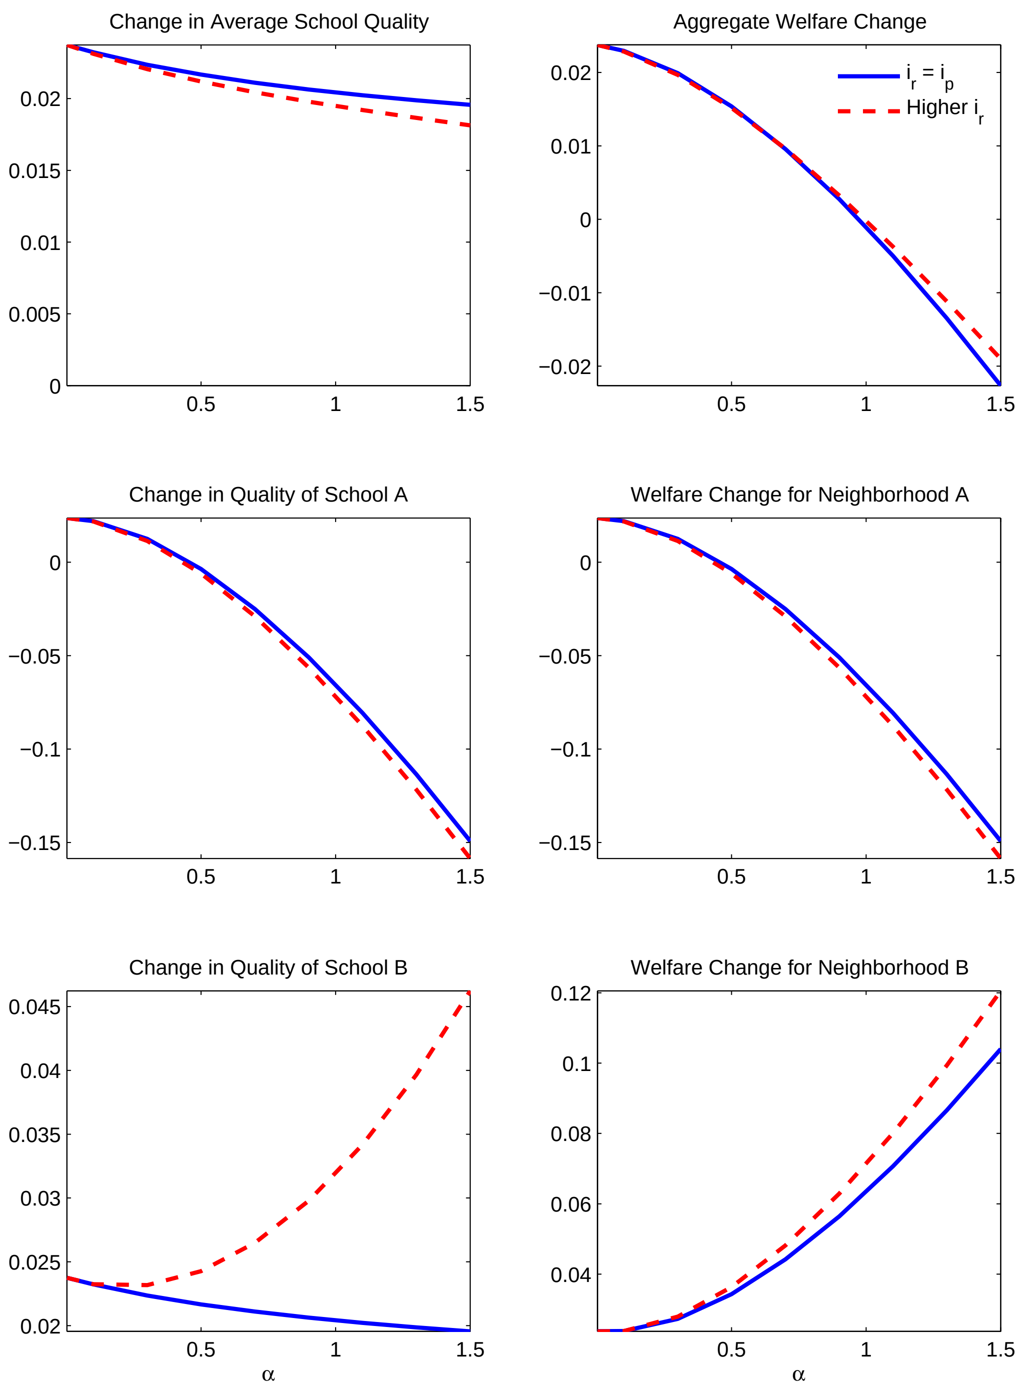

Figure 6: Helping the poor excercise choice

budget balance by holding constant the expenditure at 1 unit per student requires that

$$(1-\lambda)b_p + \lambda b_r = 1. \tag{30}$$

This equation implies that b_r equals $(1-(1-\lambda)b_p)/\lambda$ and that the difference $b_r - b_p$ equals $(1-b_p)/\lambda$. Ruling out negative surplus, the relevant range of b_p is therefore $[1, 1/(1-\lambda)]$. Setting $b_p = 1$ yields the benchmark case.

Under this policy, for any b_p in the range $[1, 1/(1-\lambda)]$, the two schools' strategies satisfy the requirements that

$$e_A(\Delta q) = \arg\max_{e_A} \left\{ b_p E(\Delta q) + (1-b_p)\frac{E_r(\Delta q)}{\lambda} - \gamma\frac{e_A^2}{2} + \beta V_A(e_A - e_B(\Delta q) + \alpha\Delta\lambda(\Delta q)) \right\}, \tag{31}$$

and that

$$e_B(\Delta q) = \arg\max_{e_B} \left\{ b_p (1-\lambda - E(\Delta q)) + (1-b_p)\left(1 - \frac{E_r(\Delta q)}{\lambda}\right) - \gamma\frac{e_B^2}{2} + \beta V_B(e_A(\Delta q) - e_B + \alpha\Delta\lambda(\Delta q)) \right\}. \tag{32}$$

Following the same steps as before, we can write the first order conditions describing the equilibrium effort levels as:

$$\gamma e_A(\Delta q) = \beta\left[b_p \frac{dE}{d\Delta q'} + \frac{(1-b_p)}{\lambda}\frac{dE_r}{d\Delta q'} + \gamma e_A(\Delta q')(\alpha\frac{d\Delta\lambda}{d\Delta q'} - \frac{de_B}{d\Delta q'}) \right], \tag{33}$$

and

$$\gamma e_B(\Delta q) = \beta\left[b_p \frac{dE}{d\Delta q'} + \frac{(1-b_p)}{\lambda}\frac{dE_r}{d\Delta q'} + \gamma e_B(\Delta q')(\alpha\frac{d\Delta\lambda}{d\Delta q'} + \frac{de_A}{d\Delta q'}) \right]. \tag{34}$$

The only difference between (33) and (34) and (15) and (16), is that the value of a change in enrollment is more complicated. This is the only place where the policy parameter b_p enters directly. Notice that when Δq is positive, then given the generalized cost structure introduced in Section 3.5, the enrollment value of a marginal increase in the quality difference is

$$b_p \frac{dE}{d\Delta q'} + \frac{(1-b_p)}{\lambda}\frac{dE_r}{d\Delta q'} = \frac{b_p}{\overline{c}}\left[(1-i_r)\lambda(1-\mu) + (1-i_p)\left(1/2 - \lambda(1-\mu)\right)\right] + \frac{(1-b_p)}{\overline{c}}(1-i_r)(1-\mu). \tag{35}$$

Note that this is increasing in b_p if

$$\mu - 1/2 + i_r(1-\mu)(1-\lambda) > i_p\left(1/2 - \lambda(1-\mu)\right). \tag{36}$$

Since μ exceeds $1/2$, this condition will be satisfied unless the fraction of immobile poor significantly exceeds the fraction of immobile rich. Assuming inequality (36) is satisfied, the maximum incentive for schools is obtained by setting b_p equal to its maximum level $1/(1-\lambda)$.

Figure 7 depicts the benefits of setting b_p equal to its maximum level as opposed to its benchmark level when all households are mobile.[29] The solid line in each panel describes what happens when b_p is equal to its benchmark value (1) and the dashed line deals with the case in which b_p is equal to its maximum value (2). The important point to note is that in all dimensions, the policy improves outcomes under choice. The reason is simple: the policy raises schools' equilibrium effort levels as they have greater incentive to compete for the marginal students in the disadvantaged neighborhood. It is still the case, however, that for sufficiently strong peer preferences, the aggregate welfare change under choice is negative. Moreover, quality in school A and welfare in neighborhood A is harmed by choice. We record this finding as:

Finding 3 ii) *When all households are mobile, then relative to the benchmark policy, a choice policy that provides schools with the maximal incentive to enroll poor students increases quality in both schools and welfare in both neighborhoods. However, when peer preferences are sufficiently strong, choice still reduces average welfare, welfare in the advantaged neighborhood and quality in the advantaged school.*

5.3 Combining the two policies

What would happen if we combined the two policies just analyzed? Figure 8 illustrates the effects of a policy which simultaneously increases i_r to 0.6, reduces i_p to 0.4666, and raises b_p to 2. Because this policy improves schools' effort incentives, it it seen to generate more desirable outcomes than a policy that only helps poor students to exercise choice. Moreover, when compared with a policy that only incentivizes schools to enroll poor students, it improves school quality and household welfare in the disadvantaged neighborhood. Nonetheless, even under this two-pronged policy, when peer preferences are sufficiently strong, average welfare gains remain negative. Again, this is because the two schools always exert similar levels of effort, so that choice is driven by peer differentials and hence is socially wasteful.

6 Discussion

We believe that the three findings discussed above have interesting implications for the choice debate. This Section summarizes the findings and discusses these implications.

[29] It goes without saying that if only rich households are mobile, paying schools for only enrolling poor students will completely backfire and provide no incentives for schools to exert effort.

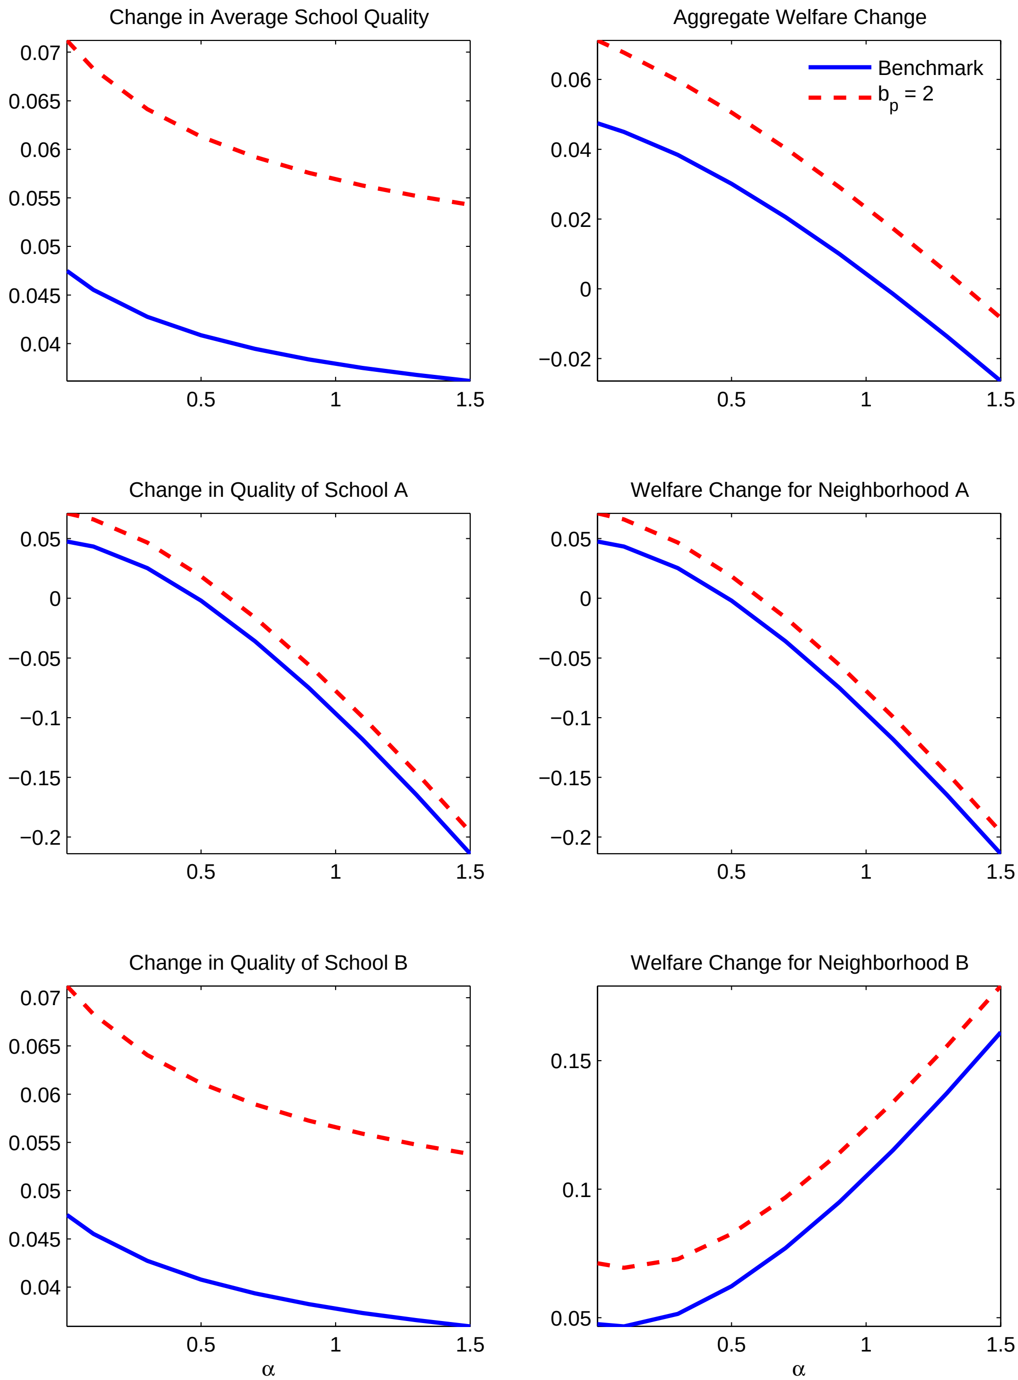

Figure 7: Incentivizing schools to enroll poor students

6.1 The effects of public school choice on school quality exhibit a type of equity-efficiency trade-off

When parents have strong peer preferences, the effects of choice on quality will exhibit a type of equity-efficiency trade-off. This trade-off stems from the dependence of effort incentives on the composition of the marginal students most likely to exercise choice. If the marginal students are richer than average (i.e., richer students are more mobile), schools will face strong incentives to recruit them and hence choice will have large impacts on average school quality. But choice will increase the gap between quality in the advantaged and the disadvantaged school, since the mobile rich students will leave the immobile poor students behind.

This trade-off implies that if public school choice is intended to impact school quality, then strong peer preferences maybe problematic for choice. Choice can increase average school quality by providing schools with strong incentives to compete for (advantaged) students. Alternatively, choice can reduce the gap between the quality of advantaged and disadvantaged schools by providing disadvantaged students with access to "good" schools. But choice may struggle to do both.

An implication of this trade-off is that the effects of public school choice will be setting-specific. This is consistent with results of choice evaluations, some of which conclude that choice improves measures of average school quality (e.g., the Lavy (2010) evaluation of an Israeli public school choice reform) and some of which conclude that it does not (e.g., the Ladd and Fiske (2000) evaluation of New Zealand public choice school reforms).[30] One explanation for the conflicting findings that is consistent with our model is that the households that exercised choice in New Zealand were less advantaged than those that exercised choice in Israel, so that New Zealand schools had weaker incentives to recruit them. Lavy (2010) is concerned exclusively with choice impacts on average outcomes and hence it is difficult to evaluate this possibility, but it is consistent with the Ladd and Fiske account of the New Zealand reforms. First, as our model would predict, Ladd and Fiske (2001) use school enrollment and composition data to show that post-choice, households gravitated from low- to high-quality schools (whether quality is measured by composition or achievement). Second, and also as our model would predict, Ladd and Fiske (2001) argue that when choosing schools, parents used "the mix of students in a school as a proxy for school quality" (p.49). Ladd's interpretation is that: "(S)uccessful schools in urban areas had no desire to expand their enrollment. To the contrary, they did everything they could to maintain the mix of students that made them attractive to parents and students

[30] Ladd and Fiske (2003) find no statistically significant correlations between the principal reports of the impact of choice on learning outcomes and principal reports of the level of local competition. Lavy (2010) evaluates an Israeli public school choice reform implemented in a single school district in Tel Aviv. He finds that relative to non-reforming control districts, a district that implemented a public school choice program enjoyed significant school productivity gains as measured by dropout rates, test scores and behavioral outcomes. Evidence from the UK is also mixed. Gibbons et al. (2008) exploit across-region variation in choice and find that choice is associated with few productivity gains. Bradley and Taylor (2002) find stronger productivity gains using a difference-in-difference approach.

in the first place..(S)chools with large concentrations of disadvantaged students have difficulty competing for students" (Ladd (2003); p.70). Although the New Zealand reform deviated somewhat from the assumptions underlying our model, and although it is not clear whether the students that exercised choice decreased the peer quality differential (a condition for the composition return to be negative in our model), these mechanisms bear a striking resemblance to those at work in our model.[31]

6.2 Public school choice can reduce average household welfare

When parents have strong preferences and when schools compete on an uneven playing field, choice may reduce welfare. How is this possible given that choice increases average school quality? The answer is that since the two schools exert similar levels of effort, the school quality differences that cause households to exercise choice will be driven purely by peer differences. But peer-driven choice is wasteful, because choice is costly and peer quality is a zero-sum construct, such that one household's peer gain is another's peer loss.[32] For the parameters that we consider, choice always reduces welfare when peer preferences are sufficiently strong (i.e., the welfare cost of wasteful choice outweighs the welfare benefit of higher effort). Notably, this is true even when only rich students are mobile so that stronger peer preferences increase schools' incentive to exert effort.

The underlying problem is that while choice is socially wasteful, the threat of choice is socially beneficial (since it elicits socially beneficial effort). When households have weak peer preferences, households threaten schools with choice, but do not exercise choice in equilibrium. That is because both schools exert similar levels of effort and hence provide similar levels of quality from the perspective of these peer-blind households. But when households have strong peer preferences, they threaten schools with choice and exercise choice in equilibrium.

One implication of this finding is that empirical evaluations of the impacts of choice on indicators of average school quality (e.g., test scores (Lavy (2010)) or the quality of teaching and learning (Ladd and Fiske (2003))) do not measure the full welfare effects of choice. That is because they exclude the costs incurred by households that exercise choice. These costs are as much a part of the welfare calculus as would be the costs of hiring superstar CEOs to run public schools. Hence even if choice increases average school quality as proxied by these measures, an economist should not recommend the policy if it reduces average welfare.

[31] One important difference is that many schools in New Zealand appear to have been capacity constrained. Note that in our model, the successful schools would like to expand enrollment and the unsuccessful schools would have no difficult competing for marginal students. However, both schools would have limited incentives to exert effort because of the negative composition returns that they faced.

[32] Put differently, we would like to subject schools to the threat of choice but would like to prevent households from actually exercising choice. The two goals are obviously incompatible.

6.3 Policies designed to enhance public school choice

As a result of the equity-efficiency trade-off that we have identified, policies that seek to improve the equity effects of choice by ensuring that choice can be exercised by disadvantaged households (e.g., by subsidizing the transport costs of poor students) will blunt schools' effort incentives and lessen the impact of choice on average school quality. This is important, since discussion of policies designed to improve outcomes under choice typically focuses on these types of policies. For example, in their "Proposal for Reform", Chubb and Moe (1990) argued for the creation of "Choice Offices" that will organize transportation "(F)or students that need it...especially the poor and those located in rural areas" (p.221) and provide "Parent Information Centers" to handle the applications process and provide information, advice and assistance to parents.[33] Our analysis implies that this type of activity could lessen the effects of choice on average school quality.

More constructively, our findings suggest that when parents have strong peer preferences, outcomes under choice can be improved by a policy not usually considered in school choice discussions: incentivizing schools for enrolling poor students. This policy seems especially promising because it is a version of the "weighted-student financing" (WSF) that already exists in many settings. Under these formulae, schools receive funds in proportion to the number of students taught with adjustments made for student composition. Some WSF formulas are fairly involved; some are more straightforward. In the UK for example, the school funding formula includes a "pupil premium" which attaches additional funds to each of a school's students eligible for a free school meal. The policy that we have in mind would compensate schools for the additional costs associated with educating disadvantaged students (as at present) and provide an additional premium to incentivize them to exert the effort necessary to attract less-advantaged students.

Of course schools cannot attract less-advantaged students if it is costly for those students to move. That is why our analysis found that a two-pronged approach offers the best hope for achieving desirable outcomes under choice. First, policies that help poor students to exercise choice, with a view to ensuring that the disadvantaged school and the disadvantaged households are not left behind. Second, policies that reward schools for enrolling poor students, thereby ensuring the largest possible efficiency gains from choice. Importantly though, even when accompanied by this two-pronged policy, choice can reduce welfare when peer preferences are sufficiently strong.

[33] Chubb and Moe (1990) note that their preferred policy would allow students to use vouchers at approved private schools, but stress that choice and competition among public schools would also be desirable.

7 Conclusion

To this point the choice debate has been characterized by largely informal arguments concerning whether or not market analysis can be applied to schools. Our aim was to shed light on this debate by providing a formal analysis of choice. To that end this paper has presented an economic model of household and school interaction in a system of public school choice and used it to assess the impact of choice on school quality and household welfare. The model assumed that households have peer preferences but was otherwise designed to capture the full force of the pro-choice argument.

With this in mind, it is perhaps surprising that our analysis ends up highlighting potential limitations of choice programs. In particular, while our model shows that choice will increase average school quality, it reveals two important caveats. First, when peer preferences are strong, choice could reduce average household welfare. Second, there will be a type of equity-efficiency trade-off associated with these effects on average quality: when peer preferences are strong, either the effects will be weak or the effects will be strong but will also be accompanied by an increase in the gap between quality in advantaged and disadvantaged schools. Choice advocates' typical response to this equity problem - the promotion of policies that help disadvantaged students to exercise choice - will have unintended consequences. Specifically, these policies will blunt schools' incentives to exert effort and thereby dampen choice's impact on average school quality. We show that some of these effects can be mitigated by also incentivizing schools to enroll disadvantaged students, but we show that choice can reduce welfare even under this two-pronged policy.

Of course our analysis is limited in many important ways. First, in focusing on public school choice, we have ignored other types of school choice, notably choice of school district (i.e., Tiebout choice) and choice policies designed to enhance options outside of the public sector (e.g., charter schools and private school vouchers). It seems reasonable to expect that peer preferences will shape the effects of these other types of choice, but the institutional environments are sufficiently different that new modeling is required.[34] The careful study of these other types of choice is an important topic for future work. Second, to shed light on the main forces driving outcomes under choice, we have made many strong assumptions (as discussed in Section 2.2). As a result, our model is too stylized to determine whether, given existing peer preferences, specific public school programs would improve on existing outcomes. Nonetheless, we hope that our paper lays some of the foundations necessary to answer this question.

[34] For example, a model of Tiebout choice must consider how districts make tax and spending decisions. An analysis of private school vouchers must consider how private schools admit students (e.g., whether or not they can turn voucher students away).

References

Angrist, J. (2013). The perils of peer effects. Working Paper 19774, National Bureau of Economic Research.

Bradley, S. and Taylor, J. (2002). The effect of the quasi-market on the efficiency-equity trade-off in the secondary sector. *Bulletin of Economic Research*, 54(3):295–314.

Burgess, S., Greaves, E., Vignoles, A., and Wilson, D. (2014). What parents want: School preferences and school choice. *Economic Journal*, Forthcoming.

Chubb, J. and Moe, T. (1990). *Politics, Markets, and America's Schools*. Brookings Institution Press.

Chubb, J. and Moe, T. (1992). *A Lesson in School Reform from Great Britain*. Brookings Institution Press.

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

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

Epple, D. and Romano, R. (2003). Neighborhood schools, choice and the distribution of educational benefits. In Hoxby, C. M., editor, *The Economics of School Choice*, chapter 7, pages 227–286. University of Chicago Press.

Farrell, J. and Klemperer, P. (2007). Coordination and lock-in: Competition with switching costs and network effects. In Armstrong, M. and Porter, R., editors, *The Handbook of Industrial Organization*, volume 3, chapter 31, pages 1967–2072. Elsevier.

Friedman, M. (1953). *Essays in Positive Economics*. Chicago University Press.

Gibbons, S., Machin, S., and Silva, O. (2008). Choice, competition, and pupil achievement. *Journal of the European Economic Association*, 6(4):912–947.

Hastings, J., Kane, T., and Staiger, D. (2006). Preferences and heterogeneous treatment effects in a public school choice lottery. NBER Working Papers 12145, National Bureau of Economic Research, Inc.

Hsieh, C.-T. and Urquiola, M. (2006). The effects of generalized school choice on achievement and stratification: Evidence from chile's voucher program. *Journal of Public Economics*, 90(8-9):1477–1503.

Judd, K. (1992). Projection methods for solving aggregate growth models. *Journal of Economic Theory*, 58(2):410–452.

Ladd, H. (2002). *Market-Based Reforms in Education*. Economic Policy Institiute, Washington, D.C.

Ladd, H. (2003). Comment on Caroline M. Hoxby: School choice and school competition: Evidence from the United States. *Swedish Economic Policy Review*, 10:67–76.

Ladd, H. and Fiske, E. (2000). *When Schools Compete: A Cautionary Tale*. Brookings Institution Press.

Ladd, H. and Fiske, E. (2001). The uneven playing field of school choice: Evidence from New Zealand. *Journal of Policy Analysis and Management*, 20(1):43–64.

Ladd, H. and Fiske, E. (2003). Does competition improve teaching and learning: Evidence from New Zealand. *Educational Evaluation and Policy Analysis*, 25(1):97–112.

Lavy, V. (2010). Effects of free choice among public schools. *Review of Economic Studies*, 77(3):1164–1191.

MacLeod, B. and Urquiola, M. (2013). Anti-lemons: Reputation and educational quality. Mimeo, version March 13, 2013.

Manski, C. (1992). Educational choice (vouchers) and social mobility. *Economics of Education Review*, 11(4):351–369.

McMillan, R. (2005). Competition, incentives, and public school productivity. *Journal of Public Economics*, 89:1131–1154.

Mizala, A., Romaguera, P., and Urquiola, M. (2007). Socioeconomic status or noise? Tradeoffs in the generation of school quality information. *Journal of Development Economics*, 84(1):61–75.

Rothstein, J. (2006). Good principals or good peers? Parental valuation of school characteristics, Tiebout equilibrium, and the incentive effects of competition among jurisdictions. *American Economic Review*, 96(4):1333–1350.

Scotchmer, S. (1985). Profit-maximizing clubs. *Journal of Public Economics*, 27(1):25–45.