Is There Any Cream to Skim? Sorting, Within-School Heterogeneity, and the Scope For Cream-Skimming

Patrick Walsh
University of Michigan

November 30, 2005

Abstract

Critics of school choice argue that cream-skimming will worsen outcomes for those left behind in public schools, a dynamic that relies on a substantial degree of within-school heterogeneity. Since "high quality" families may have *already* sorted themselves, or may represent a small fraction of the total, this paper will examine whether existing within-school heterogeneity leaves any scope for cream-skimming to operate. The first empirical section shows that the assumptions made by simulation studies over-estimate within-school heterogeneity by at least 20% to 40%, thus inflating the cream-skimming effect. The second empirical section asks, "given the current level of within-school heterogeneity, how strong would peer effects have to be to significantly worsen outcomes for those left behind?". In order for cream skimming to lower math test scores by a decile, the peer effect would have to be larger than the effect of converting both parents from college graduates to high-school dropouts. In order for cream skimming to substantially worsen dropout rates or college attendance rates, the peer effect would have to be two to three times larger than the strongest estimated predictor of these outcomes. The required peer effects would be smaller, but still unreasonably large, if family types started from a uniform distribution. These results indicate that current levels of within-school heterogeneity are so low that peer effects would have to be unrealistically strong to give cream skimming any bite.

Acknowledgements: My dissertation committee - Jim Hines (Chair), John Bound, and Dan Silverman - has provided valuable advice and feedback, as have Julie Cullen, Charlie Brown, and Lucas Davis. Fellow graduate students Mike Lovenheim and Jon Lanning have also provided valuable feedback and advice. All errors are of course my own.

1 Introduction

1.1 Overview

School-choice skeptics argue that cream-skimming will worsen outcomes for those left behind in public schools. This dynamic assumes a substantial degree of within-school heterogeneity. Actual within-school heterogeneity may be lower than assumed, for two reasons. First, ex-ante sorting among schools may result in relatively homogenous schools. Second, even if all types were spread uniformly, choice-induced sorting will have little effect on school composition if "high-quality" types represent a small share of the total.

This paper uses real-world data to examine within-school heterogeneity, and the corresponding scope for cream-skimming. It first lays out a mathematical framework for decomposing the cream-skimming effect into i) the difference between "stayers" and "leavers", ii) the share who leave, and iii) the strength of the peer effect. It shows that (i) and (ii) are increasing in within-school heterogeneity. The paper then evaluates a common assumption made by popular cream-skimming simulations: that the public sector consists of one large school. By artificially inflating within-school heterogeneity, this assumption exaggerates both the differences between leavers and stayers, and the share of students who would leave. This leads to an over-estimation of the cream-skimming effect. Moreover, the synthetic distributions of peer quality used in simulations will further exaggerate cream-skimming if their variances are over-estimated.

Rather than attempting to estimate peer effects directly, the second empirical section asks, "given existing levels of within-school heterogeneity, *how large would peer effects have to be to substantially impact those left behind?*" Families are synthetically sorted according to binary measures of their peer quality. Since many schools start with few high-quality families, even *perfect* out-sorting produces minor changes in school composition. As a result, the effect of each high-quality peer would have to be unrealistically large to produce substantial achievement effects. For example, even after sorting out *every* parent who had a BA or high income, just to change test scores by one or two points, the peer effect would have to be as large or larger than the effect of converting both parents from college graduates to high-school dropouts. In order for cream skimming to substantially worsen dropout rates or college attendance rates, the peer effect would have to be two to three times larger than the strongest estimated predictor of these outcomes.

The same analysis is repeated assuming that "high-quality" families are distributed uniformly across all schools. The required peer effects are smaller, but still generally larger than any other predictor of outcomes. This suggests that within-school heterogeneity is too small to facilitate cream-skimming, both because "high quality" families are unequally distributed, *and* because there are too few "high-quality" families to substantially affect school composition.

These results inform the ongoing policy debate over school choice and voucher plans. If within-school heterogeneity was large enough to facilitate a strong cream-skimming effect, choice plans would be harmful. It appears, however, that cream-skimming simulations have substantially over-estimated within-school heterogeneity. Moreover, since "high-quality" families make up a small share of the total, and since those families have already engaged in sorting, the cream-skimming downside to enhanced choice seems limited.

1.2 Background

Recently, the debate over school choice programs has focused on the general equilibrium effects - how would students remaining in their original schools be affected? To greatly oversimplify, the GE dynamics can be broken into two opposing effects: a competition effect (which improves the schools that face competition), and a cream-skimming effect (which leaves the "stayers" worse off).

When considering the cream skimming effect, two distinct branches of the literature provide conflicting predictions. The first strand of literature points to evidence for positive peer spillovers in education. In such an environment, the sorting of "good" families out of their existing schools (adverse sorting) reduces welfare for those left behind, and perhaps reduces welfare overall. Models in the adverse sorting literature tend to have similar characteristics. First, they assume a distribution of families over some set of attributes, such as income, ability, and possibly taste for education. Second, educational quality is modeled as a function of school inputs and average peer quality. Some models also address the housing market, and the effects of voter composition on public school funding. Epple and Romano (1998) is one of the best examples of the adverse sorting/cream-skimming literature. In addition to peer effects and a distribution over ability and income, their model allows private schools to price- and admission-discriminate based on ability. They find that those remaining in the public sector tend to be low-ability students, who are left significantly worse off by a voucher system. Caucutt (2002) uses another model with peer effects and private school discrimination to find roughly similar results. She shows that the largest welfare gains accrue to wealthy or high-ability families, and that aggregate welfare gains are decreasing in the strength of the peer effect. Other studies incorporating cream-skimming into a model of school choice include Manski (1992) and Adnett et al (2002).

The Tiebout (1956) theory, in which people sort themselves into local communities based on preferences for local public goods, suggests different implications for cream-skimming. To the extent that public-good preferences are correlated with factors like income and education, Tiebout sorting will lead to greater variance *across* communities than *within* communities. In this setting, much of the possible sorting has *already* taken place, even before a choice plan. Other studies have incorporated peer effects (critical to the cream-skimming argument) into this framework. In a model presented by De Bartolome (1991), peer effects actually increase within-community heterogeneity, nullifying the gains from Tiebout sorting. Fernandez (2001) shows that if parental human capital and community human capital are complements in producing educational quality, sorting will occur. This model builds on Fernandez and Rogerson (1996), in which sorting affects educational quality only through the composition of voters. They also show that families will sort by type, causing an inefficient equilibrium. Epple and Romano (2000) show that families will sort among schools in a multi-school district. Nechyba (1996 & 2000) models voucher plans directly, using models incorporating peer effects, competitive effects, housing markets, and a heterogenous public sector. With competitive effects and pre-sorting across public districts, these models show higher benefits and lower risks from vouchers relative to Epple and Romano's or Caucutt's models.

A large literature has examined the decision to attend a private school. Long and Toma (1988) find that this decision is influenced by family factors such as income, parental education, race, and religion, as well as local market factors such as availability of private schools, tuition, and public school quality. Lankford and Wyckoff (1992 and 1995) find generally similar results after matching student data with local or census data to evaluate school quality in the alterative sector. Figlio and Stone (2001) find that flight to private schools is influenced by public school quality, and by community factors such as crime. These *existing* patterns of private school enrollment have recently been used to inform the cream-skimming debate. Lankford and Wyckoff (2001) construct an empirical model of public/private choice, then use that model to simulate voucher-induced sorting. They find that "switchers" to the private sector have higher socio-economic status than the average public school student. However, the overall change in public-school composition is modest, particularly in terms of academic effort and achievement.

* * * *

This paper presents several innovations in the understanding of cream skimming. The accounting framework presented, although straightforward, provides a rigorous way to think about the scale of possible cream-skimming effects. While most simulations like Epple & Romano and Caucutt (as well as empirical studies like Lankford & Wyckoff 2001) implicitly assume that the public sector is made up of one large school, this paper examines the implications of a heterogenous public sector. Finally, the strength of the peer effect is critical to estimates of cream skimming, and simulations must either assume or estimate the strength of this externality. This paper simply asks, "given within-school heterogeneity, how strong would the peer effect have to be?"

2 Mathematical Framework

2.1 Components of the Cream-Skimming Effect

Peer effects are commonly modeled in the following way. Peer quality is a uni-dimensional measure. Achievement is a function of individual characteristics, school characteristics aside from peer quality, and the average peer quality at the school ($\overline{Q_s}$):

$$Achievement_{si} = \delta_i + \gamma_s + \alpha[\overline{Q_s}] \tag{1}$$

where i indexes individuals and s indexes students.

In line with most of the cream-skimming literature, α is here assumed to be positive. Some studies, however, suggest this may not always be the case[1]. Equation (1) specifies a linear relationship between $\overline{Q_s}$ and $Achievement_{si}$. This approach serves to simplify the de-composition of the cream-skimming effect. The linearity assumption will be relaxed in the empirical calculations, to allow for the possibility that the departure of the *last* high peer-quality families has more impact than the loss of infra-marginal families.

[1] Cullen, Jacob, & Levitt (2003)

This paper distinguishes between true peer effects - arising between students in the classroom - and parent effects, brought about by the presence of high-income or highly educated parents. However, it will henceforth refer to *both* effects as "peer" effects.

When some group of families leave due to increased choice, $\overline{Achievement_s}$ will change due to two effects: a composition effect, and a peer effect. By restricting attention to those who stay in the public sector, this paper will focus on the peer effect only. This effect can be modeled as:

$$\begin{aligned} \Delta Achievement_{stayers_{si}} &= [\delta_i + \gamma_s + \alpha[\overline{Q}_{s_{post-sort}}] \\ &- [\delta_i + \gamma_s + \alpha[\overline{Q}_{s_{pre-sort}}] \\ &= \alpha[\Delta\overline{Q_s}] \end{aligned} \tag{2}$$

After skimming, the change in the Achievement of those left behind is given by the change in average peer quality multiplied by α, which represents the strength of the peer effect.

We can further de-compose[2] $\Delta\overline{Q}$:

$$\Delta\overline{Q} = [share\ who\ leave][\overline{Q}_{stayers} - \overline{Q}_{leavers}] \tag{3}$$

yielding:

$$\Delta Achievement_{stayers} = \alpha[share\ who\ leave][\overline{Q}_{stayers} - \overline{Q}_{leavers}] \tag{4}$$

Equation (4), which shows that the cream skimming effect is a function of i) the difference between leavers and stayers, and ii) the share who leave, will be the key to analyzing claims about the strength of cream-skimming.

2.2 A Reduced-Form Model of Sorting

A framework is also needed to understand *who* sorts out of their current schools. Providing a utility-function based model of school choice would not only be beyond the scope of this paper, but would also be redundant in light of several existing models (Goldhaber (1996), Nechyba (1996 & 2000), Epple & Romano (1998), Lankford & Wyckoff (2001)). Instead, a more reduced-form model of sorting will be developed, which will both shed light on the distribution of "leavers" and be flexible enough to handle several different kinds of sorting. This model can be thought of as an application of the Roy model to the public-private school choice.

[2]See Appendix A for details.

Suppose that the families within a given school have the following distribution of peer quality:

$$Q_i \sim N(\mu_Q,\ \sigma_Q^2) \tag{5}$$

Every family also receives a benefit from attending private school, which is an increasing function of Q_i:

$$B_i = \lambda Q_i + \epsilon_i \tag{6}$$

where $\lambda \geq 0$ and $\epsilon \sim N(0, \sigma_\epsilon^2)$. This benefit can be thought of as a "taste for private schooling".

There is some cost C_0 associated with attending a private school. All families for whom $B_i > C_0$ attend private schools, while the rest attend public schools. A choice or voucher plan can be modeled as a change in cost to C_1, where $C_1 < C_0$. This dynamic is illustrated in Figure 1. The following can be shown[3]:

1. For any given C_0 and C_1, if $\lambda > 0$, then $[\overline{Q}_{stay} - \overline{Q}_{leave}]$ is increasing in σ_Q^2
2. If $[share\ who\ leave \mid C_1] < .5$ and $\lambda > 0$, then $[share\ who\ leave]$ is increasing in σ_Q^2.

$\sigma_\epsilon^2 = 0$ represents the "perfect sorting" case, where *all* families above a certain $\tilde{Q}$ will leave, while all families below $\tilde{Q}$ will stay. This is the worst-case scenario for cream-skimming.

3 So *Is* There Any Cream to Skim?

3.1 Within-School Heterogeneity and Simulation Assumptions

Simulations necessarily involve a number of assumptions, and fidelity to real-world parameters must often be sacrificed to improve the model's performance. However, some common simulation assumptions can reduce within-school heterogeneity, leading to over-estimates of the cream-skimming effect.

It has already been shown that within-school variance is critical to the cream-skimming effect, since it increases $[\overline{Q}_{stay} - \overline{Q}_{leave}]$ and $[share\ who\ leave]$. Within-school heterogeneity is *itself* increasing in the variance of the overall population, and decreasing in the degree of sorting among schools. By assuming that all students in the public sector attend one large school (or many identical ones), Epple & Romano and Caucutt zero out the sorting component of this relationship. This unambiguously inflates within-school heterogeneity, and by extension exaggerates the cream-skimming effect[4]. Second, although simulation studies have no choice but to derive a synthetic distribution for unobservable "true"ability/peer quality, within-school heterogeneity will be further inflated if the variance of this distribution is too high.

Both Epple & Romano and Caucutt model peer effects as operating entirely through student ability. The analysis presented by this paper treats student ability as an input to peer quality, but also considers family factors such as parental education and income.

[3]See Appendix B for details.

[4]Nechyba avoids this problem by introducing a heterogenous public sector.

3.1.1 Data

The data used for this study come from the National Educational Longitudinal Survey (NELS88). This dataset consists of a nationally representative sample of students who were 8th graders in 1988. Extensive survey data is collected from the students themselves, from their parents, and from the teachers & principal at their school. Follow-ups are conducted every two years thereafter.

The analysis is conducted with the base-year (8th grade) wave, which maximizes observations. Data on dropouts and college attendance are merged in from the 4th follow-up.

3.1.2 Sorting and Peer-Quality Variance

Both simulation studies assume that all students in the public sector attend one large school. In reality, at least some of the total variation in peer quality will be *between* schools, and therefore would not contribute to the variation within a given school. Table 1 presents the results of ANOVA decompositions of peer quality variance into between-school and within-school components. Because cream-skimming is thought to disproportionately hurt families with lower socio-economic status, this analysis is then repeated to show the overall breakdown if *all* schools looked like low-SES schools [5]. Low-SES schools tend to have lower within-school variance, suggesting that the bias introduced by assuming a monolithic public sector is even greater for exactly the population of most interest.

The results show that even for the full sample of schools, between-school variation accounts for 20-30% of total variation. If all schools looked like low-SES schools, between-school variation would account for 25% to 45% of total variation. Whatever the overall variance in peer quality assumed by the simulation studies, students at individual schools face 30-40% less variance due to sorting.

To get a better feel for the existing degree of sorting, Graphs 1-3 show the proportion of families with a given characteristic (parent has a BA or higher, for example) at each of the 900 schools in the NELS88 dataset. Each data point represents one school (unless several schools share the exact same proportion, as is the case at 0), and schools are ranked by their proportion. The horizontal distance between data points represents the number of students at a given school. Thus the curve is implicitly enrollment-weighted. In each graph, a horizontal line denotes the sample-wide proportion of the given characteristic. This line shows how the distribution would appear if "high-quality" families were distributed uniformly. These graphs show that many students *already* attend schools with extremely low proportions of "high-quality" peers.

Table 2 puts precise numbers on the story emerging from these graphs. For the same binary peer-quality measures, the sample-wide proportion of observations with that measure is reported in the first column. The next column reports the share of observations who attend schools with *less* than the sample-wide proportion of the given quality measure. For instance, 67% of observations attend schools where less than 30% of students have a math score above 61 (a uniform distribution of such students would yield 0% in this cell; a symmetric distribution would yield 50%). The next two columns report the share of

[5] Low-SES schools are defined as schools where more than 20% of the student body qualifies for free lunch. See Appendix C for details on the re-calculation based on this group.

observations who attend schools with less than 20% or less than 10% proportions of the given quality measure. For instance, 24% of observations attend schools where less than 10% of students have math scores above 61. Table 2 indicates that a significant proportion of observations already attend schools with extremely low proportions of high-quality peers.

These results indicate that, even if most high peer-quality families were to leave their current schools as the result of increased choice, many schools would see little change in composition.

3.1.3 Overall Variance in Peer Quality

Finally, consider the ability distribution of the overall population. Since true ability cannot be measured, and proxies such as test scores suffer from scaling problems, simulation studies must derive a synthetic distribution. Epple & Romano derive a lognormal ability distribution by using accepted inter-generational correlations between income and ability, assuming a steady state, and backing out ability from observed income. Caucutt's model uses a much simpler ability distribution: 50% of students are assumed to have an ability level of "1", while the other 50% are assumed to have an ability level of "4". Given that true ability cannot be measured, these approaches are not inappropriate. However, if the resulting ability variance is over-estimated, within-school heterogeneity will be artificially inflated.

There is no way to compare the simulated variances to an unobserved "true" variance. However, suppose that in addition to student ability, income and parental education are a large component of peer quality. Returning to the first column of Table 2[6], the sample-wide proportions of high income and highly-educated families is not large. This translates into a rather low horizontal line in Graphs 2 and 3: even if equally distributed, "high quality" peers make up (by these definitions) low fractions of the total. Caucutt's assumption that 50% of the population is a "high quality" peer seems especially questionable in this light. Epple & Romano's lognormal ability distribution seems much less likely to have "too many" high-quality peers, although a high variance could still arise if those peers are "too different" than the mean.

3.2 How Large Must Peer Effects Be?

3.2.1 Empirical Strategy

The analysis in the previous section shows that because of their assumptions about within-school heterogeneity, simulation studies may significantly overstate the cream-skimming effect. However, we have yet to put a number on *how much* smaller the actual effect will be. Recall that the cream-skimming effect is given by:

$$\Delta Achievement_{stayers} = \alpha[share\ who\ leave][\overline{Q}_{stayers} - \overline{Q}_{leavers}] \tag{7}$$

The (share who leave) term could be plausibly generated from real-world data and a sorting algorithm. However, the $[\overline{Q}_{stayers} - \overline{Q}_{leavers}]$ term would require a continuous measure of true peer quality. Likewise, α represents the strength of the peer effect, which is the focus of difficult research projects in its own right.

[6]Setting aside the sample-wide proportion of high-scoring students, which is here 30% by construction.

This paper does not attempt to estimate α, nor does it presume to appoint some family X as the "true" measure of peer quality. Instead, it uses the model,

$$\Delta Achievement_{stayers} = \beta[share\ who\ leave^{\theta}] \tag{8}$$

with the understanding that the leavers are somehow "different" from the stayers. Since, for example, a ten-percentage-point drop in the proportion of high peer-quality families may have a larger impact when such families make up 10% of the total than when they make up 50% of the total, θ allows for a non-linear relationship between $\Delta Achievement_{stayers}$ and $[share\ who\ leave]$. $0 < \theta \leq 1$. Calculations are made for $\theta = 0.25$, $\theta = 0.50$, and $\theta = 1.00$. Selected family characteristics are used to synthetically sort high-quality families out of their current schools. Several measures of $\Delta[Achievement]$ that seem "large" are then be selected. The final step calculates the peer-effect β that, given [share who leave], is required to produce $\Delta Achievement$. We can then ask whether this peer-effect β seems "reasonable" in comparison to other predictors of Achievement.

This approach deliberately avoids making claims about $[\overline{Q}_{stayers} - \overline{Q}_{leavers}]$. This is because we have no guarantee that a given measure of family quality (parental education, for example) is a one-for-one correlate for true peer quality. A practical complication is that most of the dataset's measures of family quality are categorical in nature, with ordinal but not cardinal values. This leaves $[\overline{Q}_{stayers} - \overline{Q}_{leavers}]$ undefined. Ignoring this term implicitly wraps $[\overline{Q}_{stayers} - \overline{Q}_{leavers}]$ into the β. Such an approach does *not* bias our calculated β's, but does affect their interpretation. In this framework, a β represents the peer effect on achievement of changing a school from 100% families with a given characteristic to 0% families with that characteristic. This effect can then be compared to the estimated effect of a 100%-0% change in some other peer quality measure.

It is possible to get the flavor of α's calculated using an explicit but questionable $[\overline{Q}_{stayers} - \overline{Q}_{leavers}]$. As near-continuous measures, student test score and parental education are at least technically viable candidates for "peer quality". When we sort families based on student test scores, for example, we can therefore calculate $[\overline{test\ scores}_{stayers} - \overline{test\ scores}_{leavers}]$. α is interpreted as the peer-effect impact of a one-point change in a school's average test score.

To calculate the required peer-effect β's, we first choose a discrete margin on which family peer-quality can be measured. Five binary measures and three interactions are used:

- Whether the parent has attained a BA or higher
- Whether the family's income is over $50,000
- Whether the student scored above the 70th %tile) on the standardized math test
- Whether the parents expect the student to eventually attain *higher* than a BA
- Whether the student spends more than 5 hours/week on homework
- The union of parental BA-or-higher and high income
- The union of BA-or-higher and test score above 70th %tile
- The union of BA-or-higher, high income, and test score above 70th %tile.

Binary family peer-quality measures are used to simplify the hypothetical sorting process; multi-category measures could also be used. The intuition for the interactions is as follows. The first four measures of peer quality are univariate, but are correlated with one another. Removing all parents with BA's or greater *also* removes many parents with incomes over $\$50K$. Binary sorting on any one measure may therefore under-estimate cream-skimming if we think that both the share of parental BA's *and* the share of high-income families provide peer effects. The interactions avoid this problem by very simply removing everyone from *both* groups, leaving only those with less than a BA *and* low income. Similarly, the second interaction removes parents with BAs *and* students with high test scores, while the third interaction removes all three groups.

It is commonly believed that involved parents, who contact the school and lobby for improved school resources, also provide positive externalities. If so, parents who contact the school should represent another binary peer-quality variable. However, if parental lobbying is endogenous, it will not necessarily signal high peer-quality. First, since parents react to perceived school quality, lobbying may simply signal low school quality. Second, if lobbying does provide a public good, skimming away the *current* active parents may induce previously inactive parents to contact the school, negating the cream-skimming effect. Moreover, as shown in Walsh (2005), active parents can create negative externalities on some margins, by capturing a larger share of fixed school resources for their own children. This implies that the net externality from involved parents is at best ambiguous.

Income $\geq \$50K$ will henceforth be used as an extended example of family peer-quality; the analysis was conducted for all other family peer-quality measures in an exactly analogous way.

Schools are sorted by their proportion of a given type - in our example, schools are sorted by their proportions of high-income families. Suppose that the distribution of high-income families among schools is represented by Figure 3. To focus on the worst possible case, further assume that a school choice plan results in perfect sorting. This would mean that a student either attends a school where *no* families have high income, or a school where *all* families have high income. To implement perfect sorting, the "cutoff" school is found, such that the number of high-income families below the cutoff (Group B in Figure 4) is equal to the number of *non*-high-income families above the cutoff (Group C in Figure 4). Each peer-quality measure will have its own unique Groups A, B, C, and D. Then, assume that Groups B and C swap places. Because of the univariate nature of the peer effect (see Equation 9 below), it does not matter which parents go to which new schools. After the sorting, the *only* effect that Group C has on Group A is to lower the proportion of high-income families to zero.

We wish to examine the impact of adverse sorting on those "left behind" in schools that high peer-quality families have fled. Thus we will restrict attention to Group A_{income} in Figure 4: the non-high-income parents who are now left in non-high-income schools.

Suppose that achievement is determined by:

$$Achievement_{mtsi} = \delta_i + \gamma_s + \beta_{mt}[\%\ high\ Q\ families^{\theta}_{ts}] \tag{9}$$

where m indexes achievement measures, t indexes family peer-quality measures, s indexes schools, and i indexes individuals.

For example,

$$Test\ Score_{income,\ si} = \delta_i + \gamma_s + \beta_{test\ score,\ income}[\%\ high\ income^{\theta}_{s}] \tag{10}$$

Abbreviating,

$$Test\ Score_{inc,\ si} = \delta_i + \gamma_s + \beta_{test,\ inc}[d^{\theta}_{s}] \tag{11}$$

where $d_s = \%\ high\ income_s$

After sorting, each school below the cutoff now has 0% high income, so the *drop* in % high income will be equal to the pre-sorting % high income. The change in test score for an individual in Group A_{income} will therefore be given by:

$$\begin{aligned} \Delta Test\ Score_{inc,\ si} &= [\delta_i + \gamma_s + \beta_{test,\ inc}(0)^{\theta}] - [\delta_i + \gamma_s + \beta_{test,\ inc}(d^{\theta}_{s,\ pre-sorting})] \\ &= -\beta_{test,\ inc}(d^{\theta}_{s,\ pre-sorting}) \end{aligned} \tag{12}$$

We want to produce a targeted change in the *average* test score for people in Group A_{income}. This gives:

$$\begin{aligned} \Delta(\overline{Test\ Score}_{inc,\ Group\ A}) &= \left[\frac{1}{N_{inc,\ Group\ A}}\right] \Sigma_{si} Test\ Score_{inc,\ si} \\ &- \left[\frac{1}{N_{inc,\ Group\ A}}\right] \Sigma_{si}[Test\ Score_{inc,\ si} - \beta_{test,\ inc}(d^{\theta}_{s,\ pre-sorting})] \\ &= \left[\frac{1}{N_{inc,\ Group\ A}}\right] \Sigma_{si}[-\beta_{test,\ inc}(d^{\theta}_{s,\ pre-sorting})] \end{aligned} \tag{13}$$

If the enrollment of school s is given by E_s, then equation (13) becomes:

$$\begin{aligned}\Delta(\overline{Test\ Score}_{inc,\ Group\ A}) &= \Sigma_s\left[(\frac{E_s}{N_{inc,\ Group\ A}})(\frac{1}{E_s})\Sigma_i[-\beta_{test,\ inc}(d^{\theta}_{si,\ pre-sorting})]\right] \\ &= -\beta_{test,\ inc}\Sigma_s[(\frac{E_s}{N_{inc,\ Group\ A}})(d^{\theta}_{s,\ pre-sorting})]\end{aligned} \tag{14}$$

That is, the change in average test scores for Group A is $-\beta$ times the enrollment-weighted average of $\%\ high\ income_{s|Group\ A,\ pre-sorting}$. Solving for $\beta_{test,\ inc}$ gives:

$$\beta_{test,\ inc} = -\frac{\Delta(\overline{Test\ Score}_{inc,\ Group\ A})}{\Sigma_s[(\frac{E_s}{N_{inc,\ Group\ A}})(d^{\theta}_{s,\ pre-sorting})]} \tag{15}$$

The "required peer-effect β's" are computed according to equation (15)[7].

For the test score measure of Achievement, four targets for $\Delta_{target}[Achievement]$ are used, two for math scores and two for reading scores: 1) the differences in test scores that would move a student from the 50th to the 40th %tiles of the individual test score distributions, and 2) the differences in test scores that would move a *school* from the 50th to the 40th %tiles of the *school average* test score distributions.

Since different readers will have different definitions of a "large" change, these targets are presented as guides, rather than definitive classifications of "large". Given the linear nature of the peer effect in Equation (1), the reader can simply scale the displayed results to his or her own definition of large: multiply the required β's by two to target 2 deciles, or by .5 to target a half-decile $\Delta_{target}[Achievement]$. The magnitudes of the decile $\Delta_{target}[Achievement]$ are given in Table 3.

A "large" increase in the dropout rate is defined as multiplying the existing average rate for Group A by 1.2. A dropout rate increasing from 15% to 18%, for example, would be considered large. A "large" drop in the college attendance rate is defined as multiplying the existing rate for Group A by 0.8. A college attendance rate falling from 60% to 48%, for example, would therefore be considered large.

3.2.2 Results

The current degree of within-school heterogeneity will determine the scope in which cream skimming may operate. If high peer-quality families are widely distributed across schools, perfect sorting will result in serious drops in % high income (for example) for many schools below the cutoff. In such an environment, a relatively small β can produce a substantial

[7] By analogous steps, the calculation of the α's where $(\overline{Q}_{stayers} - \overline{Q}_{leavers})$ is given by:

$$\alpha_{test,\ inc} = -\frac{\Delta(\overline{Test\ Score}_{inc,\ Group\ A})}{\Sigma_s[(\frac{E_s}{N_{inc,\ Group\ A}})][(d^{\theta}_{s,\ pre-sorting})(\overline{Q}_{stayers\ s} - \overline{Q}_{leavers\ s})]} \tag{16}$$

shift in Group A outcomes. On the other hand, if high peer-quality families are highly concentrated, schools below the cutoff will start with low proportions of high-income families. Perfect sorting will therefore result in only modest drops in % high income. In this case, a large peer-effect β will be required to move Group A outcomes substantially.

β, Test Scores Results for the β analysis, using math test scores as the outcome measures, are reported in Table 4. These test scores are normed to put the median at 50 points, with minimums in the low 30's and maximums in the high 70's. The first column in Table 4 reports the share of families below the cutoff who have a given quality characteristic. The next two columns report the "required peer-effect β's" for the two different $\Delta_{target}[Achievement]$ measures. The top panel displays results for the actual distribution of families across schools. The bottom panel calculates what the required peer-effect β's would be if high peer-quality families were uniformly distributed across schools.

An example of how the required peer-effect β's should be interpreted: How strong would the peer effect have to be for perfect sorting based on parental BA to drop Group A math test scores by one decile of the individual score distribution? Changing a school from 100% parental BA to 0% parental BA would have to drop math scores at that school by 27 points. Likewise, for perfect sorting on parental BA to drop Group A math test scores by a decile of the school-average score distribution, the same change would have to drop math scores by 11 points. These β's would compare to 14 and 6 points, respectively, if families were distributed uniformly. Note that required peer-effect β's increase with the size of the $\Delta_{target}[Achievement]$ measure, and decrease with the share who leave (given here by Group B/(Group A + B)).

The results for reading test scores are described in exactly the same manner in Table 5. These results closely follow the pattern of math-score required β's.

Are these β's believable estimates of the strength of the peer effect? If a school's proportion of a given peer-quality measure falls from 100% to 0%, these peer-effect β's indicate movement in average scores of 4 to 40 test points. This seems large, given that the test score distribution has a mean of 50 and a standard deviation of around 10. Then again, a shift from 100% parental BA to 0% parental BA (for example) is itself quite large. To put these coefficients into context, we can compare them to *other* predictors of student test scores. Table 8 shows the results of regressing student test scores on a variety of school- and family-level variables, for the entire sample. These regressions are *not* intended to be the definitive, last-word models of education production functions. They merely give us, for comparison purposes, an order-of-magnitude estimate of the strength of other test-score determinants.

The first impression from Table 8 suggests that our required peer-effect β's are larger, usually by one or more orders of magnitude, than any other predictors of test scores. A stronger test would be to use Table 8 to evaluate the test-score impact of some change that is expected to have a large impact on student achievement, and compare this to our peer-effect β's. For instance, we can evaluate the impact of both parents converting their education from BA to high school dropout. For math scores, the test score would fall by 1.2 point as the parent converted from BA to HS diploma, then by another 1.1 point as the parent converted from HS diploma to dropout. This gives a drop of 2.3 points per parent, or a 4.6 point drop overall. Nearly all the required peer-effect β's for math score are higher than this effect, and many are two or three times higher. The only peer-effect β's that come close to the 4.6 point drop are when we pick the smaller of our $\Delta_{target}[Achievement]$ measures, and assume that a choice plan will *perfectly* sort out any parents who have BAs *or* income above \$50K.

The results are quite similar for reading test scores. As shown in Table 8, converting both parents from BAs to dropouts results in a score drop of $2 * (.7 - (-.8)) = 3$ points. The only reading-score peer-effect β in that range comes with the smallest $\Delta_{target}[Achievement]$ measure and the assumption that a choice plan perfectly sorts out any parents with BAs *or* income above \$50K *or* attainment expectations above BA.

To conclude, for nearly all definitions of $\Delta_{target}[Achievement]$ and peer-quality sorting measures, the peer-effect β required to make a substantial impact on test scores exceeds the effect of converting both parents from BAs to high-school dropouts, even if we assume that a choice plan leads to *perfect* sorting. If high peer-quality families were instead *equally* distributed across schools, the required peer-effect β would roughly equal the effect of changing parental education, although only for the smallest definition of $\Delta_{target}[Achievement]$.

α, Test Scores Here, $[\overline{Q}_{stayers} - \overline{Q}_{leavers}]$ is calculated by using years of parental education as our measure of peer quality when sorting by parental BA, and math score as our measure of peer quality when sorting by math score > 61. This permits the calculation of the required peer-effect α's, which are given in Table 6. When families are sorted perfectly on parental BA, the enrollment-weighted average difference between leavers and stayers is 4.67 years of education. Likewise, when families are sorted perfectly on whether the student has a math score above 61, the enrollment-weighted average difference between leavers and stayers is 19.48 points.

An example of how the required peer-effect α's in Table 6 should be interpreted: in order for perfect sorting on parental BA to lower stayers' test scores by one decile of the individual distribution, a one-year change in a school's average parental education must change test scores by 5.7 points. Likewise, for perfect sorting on parental BA to lower stayers' test scores by one decile of the *school* distribution, the same one-year change in a school's average parental education would have to change test scores by about 2.4 points.

Are these results reasonable? Again, we turn to Table 8 to compare these effects to the other estimated predictors of test score. First consider sorting on parental BA, and a targeted change of one decile in the individual score distribution. The required peer-effect α on a one-year change in a school's average parental education (5.72) is over four times as large as the effect of converting the test-taker's parents from high-school graduates to college graduates (1.2). Since a one-year change in school average education is four times larger than an *four*-year change in own-parent education, this implies that a school's average education level must be, on a year-by-year basis, *16 times* more powerful than an individual's parents' education level. When we instead target a decile in the school distribution, the equivalent effects are (2.4) and (1.3). This implies that a school's average education level must be, on a year-by-year basis, about 7 more powerful than an individual's parents' education level.

Sorting on student test score yields a more intuitive interpretation. How strong would the peer effect have to be for perfect sorting on test score to lower stayers' test scores by one decile of the individual distribution? A one-point change in a school's average score must change an individual's score by about 1.2 points (through a peer effect, not composition). Likewise, for perfect sorting on student test score to lower stayers' test scores by one decile of the *school* distribution, a one-point change in a school's average score must change an individual's score by about 0.5 points.

β, **Dropout** The required peer-effects to raise Group A dropout rates by a factor of 1.2 are reported in Table 7. The same variety of peer-quality measures are used, and the β's are again reported for the actual distribution and for a hypothetical uniform distribution of high peer-quality families. To again interpret by example: in order for perfect sorting on parental BA to raise the dropout rate by a factor of 1.2, moving a school from 100% parental BA to 0% parental BA would have to raise the dropout rate by 28 percentage points.

For the dropout rate, most required peer-effect β's fall between 15 and 30 percentage points for a 100%-0% shift in the proportion of high peer-quality families. Again, to judge whether these are reasonable estimates of the peer effect, we look to other predictors of dropout rate. Table 9 shows the results of a regression of school dropout rate on a number of school-level variables. The required peer-effect β's are significantly larger than the largest predictor of dropout. They imply that the peer effect would have to be two to three times stronger than the effect of converting the school's % free lunch from 0% to 100%. These required peer-effect β's would come closer to (but still surpass) the free lunch effect if high peer-quality families were uniformly distributed.

β, **College Attendance** The required peer-effects to lower Group A college attendance rates by a factor of .8 are also reported in Table 7. To again interpret by example: in order for perfect sorting on family income to lower the college attendance rate by a factor of .8, moving a school from 100% high income to 0% high income would have to lower the college attendance rate by 85 percentage points.

For the college attendance rate, most required peer-effect β's fall between 30 and 100 percentage points for a 100%-0% shift in the proportion of high peer-quality families. We again look to other predictors of dropout rate in Table 9 to judge whether these are reasonable estimates of the peer effect. The largest coefficient is again on the school's % free lunch, but this coefficient is little more than 1/3 the magnitude of the smallest required peer-effect β.

These results suggest that, given the high degree of current sorting, peer effects would have to be unrealistically large for cream skimming to have a noticeable impact on college attendance rates.

4 Conclusions

This paper presents a simple framework demonstrating that the cream skimming effect is a function of the share who leave, the difference between leavers and stayers, and the strength of the peer effect. Both [*share who leave*] and [$\overline{Q}_{stayers} - \overline{Q}_{leavers}$] are increasing in within-school heterogeneity. Within-school heterogeneity is itself an increasing function of overall variance and a decreasing function of sorting among schools.

Empirically, this paper first shows that assuming a homogenous public sector increases within-school heterogeneity. This unambiguously leads to an over-estimate of the cream skimming effect. Moreover, if synthetic distributions of peer quality over-estimate variance, cream skimming will be further exaggerated.

This paper then asks, "given the current level of sorting, how strong would peer effects have to be to significantly worsen outcomes for those left behind?". For test scores of those left behind, I find that the required peer effect would have to be as strong or stronger than the effect of converting both parents from college graduates to high-school dropouts. The peer effects required to substantially raise the dropout rate or reduce the college attendance rate of those left behind are much larger than the largest estimated predictor of these outcomes. Overall, these results strongly suggest that existing within-school heterogeneity is too low for cream-skimming to have large effects. This low heterogeneity is driven by two factors - a low overall share of "high quality" families, and ex-ante sorting of families among schools.

References

[1] **Adnett, Nick, Spiros Bougheas, and Peter Davies**, 2002 "Market-Based Reforms of Public Schooling: Some Unpleasant Dynamics" *Economics of Education Review*, 21

[2] **Caucutt, Elizabeth M.**, 2002 "Educational Vouchers When There Are Peer Group Effects - Size Matters" *International Economic Review*, 43(1)

[3] **Cullen, Julie B., Brian Jacob and Steven Levitt**, 2003 "The Effect of School Choice on Student Outcomes: Evidence From Randomized Lotteries" NBER # 10113

[4] **DeBartolome, Charles A.M.**, 1990 "Equilibrium and Inefficiency in a Community Model With Peer Group Effects" *Journal of Political Economy*, 98(1)

[5] **Epple, Dennis and Richard E. Romano**, 1998 "Competition Between Private and Public Schools, Vouchers, and Peer-Group Effects" *American Economic Review*, 88(1)

[6] **Epple, Dennis and Richard E. Romano**, 2000 "Neighborhood Schools, Choice, and the Distribution of Educational Benefits" NBER # 7850

[7] **Fernandez, Raquel**, 2001 "Sorting, Education, and Inequality" NBER # 8101

[8] **Fernandez, Raquel and Richard Rogerson**, 1996 "Income Distribution, Communities, and the Quality of Public Education" *Quarterly Journal of Economics*, 111(1)

[9] **Figlio, David N. and Joe Stone**, 2001 "Can Public Policy Affect Private School Cream Skimming?" *Journal of Urban Economics*, 49

[10] **Lankford, Hamilton and James Wyckoff**, 1992 "Primary and Secondary School Choice Among Public and Religious Alternatives" *Economics of Education Review*, 11(4)

[11] **-** , 1995 "An Analysis of Elementary and Secondary School Choice" *Journal of Urban Economics*, 38

[12] **-** , 2001 "Who Would Be Left Behind By Enhanced Private School Choice?" *Journal of Urban Economics*, 50

[13] **Long, James E. and Eugenia Toma**, 1988 "The Determinants of Private School Attendance, 1970-1980" *Review of Economics and Statistics*, 70(2)

[14] **Manski, Charles F.**, 1992 "Educational Choice (Vouchers) and Social Mobility" *Economics of Education Review*, 11(4)

[15] **Nechyba, Thomas J.**, 1996 "Public School Finance in a General Equilibrium Tiebout World: Equalization Programs, Peer Effects, and Private School Vouchers" NBER # 5642

[16] **Nechyba, Thomas J.**, 2000 "Mobility, Targeting, and Private School Vouchers" *American Economic Review*, 90(1)

[17] **Tiebout, Charles M.**, 1956 "A Pure Theory of Local Expenditure" *Journal of Political Economy*, 64(October)

5 Results

Table 1: ANOVA Variance Decomposition

	Full Sample		If All Schools Had w/in School Variance of Low SES Schools	
Variance of:	**Between**	**Within**	**Between**	**Within**
Student Math Score	.300	.700	.344	.656
Student Reading Score	.228	.772	.258	.742
Parent has BA or Higher	.259	.741	.462	.538
Family Income $> \$50K$	.290	.710	.465	.535

Table 2: Distribution of "High Peer-Quality" Families

Tabular Representation of Graphs 1-3

Peer-Quality Measure	**Proportion of Overall Sample**	**Share at Schools with Proportion $<$ Overall**	**Share at Schools with Proportion $<$ 0.20**	**Share at Schools with Proportion $<$ 0.10**
Math Score $>$ 61 (70th %tile)	0.300	0.672	0.513	0.245
Parent has BA or Higher	0.226	0.649	0.592	0.345
Family Income $> \$50K$	0.274	0.631	0.502	0.256

Drop Schools with <15 Obs

Peer-Quality Measure	**Proportion of Overall Sample**	**Share at Schools with Proportion $<$ Overall**	**Share at Schools with Proportion $<$ 0.20**	**Share at Schools with Proportion $<$ 0.10**
Math Score $>$ 61 (70th %tile)	0.303	0.676	0.508	0.231
Parent has BA or Higher	0.227	0.651	0.589	0.337
Family Income $> \$50K$	0.277	0.630	0.496	0.246

Table 3: Targeted Test Score Changes

	Math Score	**Reading Score**
Decile, Individual Distrib.	3.18	3.19
Decile, School Distrib.	1.32	1.02

Table 4: Required β's For *Math* Test-Score Drops

Panel A: Actual
Within-School Heterogeneity

		Required β's					
	Average	Target Score Drop: Decile, Individual Distn'			Target Score Drop: Decile, School Average Distn'		
Peer-Quality Sorting Measure:	**Proportion At Schools Below Cutoff**	$\theta = 1$	$\theta = 0.50$	$\theta = 0.25$	$\theta = 1$	$\theta = 0.50$	$\theta = 0.25$
Parent Has BA or Higher	.119	26.71	10.55	6.44	11.77	4.65	2.84
Family Income > 50K$	.147	21.69	9.20	5.82	11.59	6.07	4.39
Student Math Score above 70th %tile	.138	22.97	9.43	5.90	10.65	5.82	2.60
Parent Expects Student > BA	.185	17.15	7.80	5.14	15.03	6.92	4.69
Student Does Homework > 5 Hrs/Week	.232	13.70	6.75	4.68	6.04	2.97	2.06
Parent BA *or* Income > 50K$	.222	14.35	7.26	5.02	9.59	5.52	4.19
Parent BA *or* Score above 70th %tile	.193	16.51	7.71	5.15	7.75	4.97	3.97
BA, Income, *or* Math Score	.246	12.92	6.71	4.74	6.30	4.48	3.77

Panel B: High Peer-Quality Families
Uniformly Distributed

		Required β's					
	Average	Target Score Drop: Decile, Individual Distn'			Target Score Drop: Decile, School Average Distn'		
Peer-Quality Sorting Measure:	**Proportion At Schools Below Cutoff**	$\theta = 1$	$\theta = 0.50$	$\theta = 0.25$	$\theta = 1$	$\theta = 0.50$	$\theta = 0.25$
Parent Has BA or Higher	.226	14.09	6.69	4.61	6.21	2.95	2.03
Family Income > 50K$	.274	11.59	6.07	4.39	5.11	2.68	1.94
Student Math Score above 70th %tile	.300	10.65	5.82	4.30	4.69	2.56	1.90
Parent Expects Student > BA	.212	15.03	6.92	4.69	6.63	3.05	2.07
Student Does Homework > 5 Hrs/Week	.384	8.29	5.13	4.04	3.65	2.26	1.78
Parent BA *or* Income > 50K$	.332	9.59	5.52	4.19	4.22	2.43	1.85
Parent BA *or* Score above 70th %tile	.410	7.75	4.97	3.97	3.42	2.19	1.75
BA, Income, *or* Math Score	.505	6.30	4.48	3.77	2.78	1.97	1.66

Table 5: Required β's For *Reading* Test-Score Drops

Panel A: Actual
Within-School Heterogeneity

		Required β's					
		Target Score Drop: Decile, Individual Distn'			Target Score Drop: Decile, School Average Distn'		
Peer-Quality Sorting Measure:	**Average Proportion At Schools Below Cutoff**	$\theta = 1$	$\theta = 0.50$	$\theta = 0.25$	$\theta = 1$	$\theta = 0.50$	$\theta = 0.25$
Parent Has BA or Higher	.119	26.76	10.57	6.45	8.54	3.37	2.06
Family Income > 50K$	.147	21.73	9.22	5.84	6.93	2.94	1.86
Student Math Score above 70th %tile	.138	23.02	9.45	5.91	7.34	3.01	1.88
Parent Expects Student > BA	.185	17.18	7.82	5.15	5.48	2.49	1.64
Student Does Homework > 5 Hrs/Week	.232	13.73	6.76	4.69	4.38	2.16	1.50
Parent BA *or* Income > 50K$	.222	14.38	7.27	5.03	4.59	2.32	1.60
Parent BA *or* Score above 70th %tile	.193	16.55	7.72	5.16	5.28	2.46	1.65
BA, Income, *or* Math Score	.246	12.95	6.72	4.75	4.13	2.14	1.51

Panel B: High Peer-Quality Families
Uniformly Distributed

		Required β's					
		Target Score Drop: Decile, Individual Distn'			Target Score Drop: Decile, School Average Distn'		
Peer-Quality Sorting Measure:	**Average Proportion At Schools Below Cutoff**	$\theta = 1$	$\theta = 0.50$	$\theta = 0.25$	$\theta = 1$	$\theta = 0.50$	$\theta = 0.25$
Parent Has BA or Higher	.226	14.12	6.71	4.62	4.50	2.14	1.47
Family Income > 50K$	.274	11.62	6.08	4.40	3.70	1.94	1.40
Student Math Score above 70th %tile	.300	10.67	5.83	4.31	3.40	1.86	1.37
Parent Expects Student > BA	.212	15.07	6.93	4.70	4.80	2.21	1.50
Student Does Homework > 5 Hrs/Week	.384	8.31	5.15	4.05	2.65	1.64	1.29
Parent BA *or* Income > 50K$	.332	9.61	5.53	4.20	3.06	1.76	1.34
Parent BA *or* Score above 70th %tile	.410	7.77	4.98	3.98	2.48	1.59	1.27
BA, Income, *or* Math Score	.504	6.32	4.49	3.78	2.01	1.43	1.21

Table 6: Required α's

Math Score Drops, using Actual Within-School Heterogeneity

Peer-Quality Sorting Measure:	**Average Proportion Below Cutoff** = Group B/(Group A + B)	$\overline{Q}_{stayers} - \overline{Q}_{leavers}$	**Decile, Individual Distribution**	**Decile, School Distribution**
Parent Has BA or Higher	.119	-4.67 yrs	5.72	2.37
Math Score > 61	.138	-19.48 points	1.18	0.49

Reading Score Drops, using Actual Within-School Heterogeneity

Peer-Quality Sorting Measure:	**Average Proportion Below Cutoff** = Group B/(Group A + B)	$\overline{Q}_{stayers} - \overline{Q}_{leavers}$	**Decile, Individual Distribution**	**Decile, School Distribution**
Parent Has BA or Higher	.119	-4.67 yrs	5.74	1.84
Math Score > 61	.138	-19.48 points	1.19	0.38

Table 7: Necessary β's to Increase Dropout Rate by a Factor of 1.2 Or Decrease College Attendance Rate by a Factor of .2

	Dropout Rate	
Peer-Quality Measure:	**Actual Within-School Heterogeneity**	**High Peer-Quality Families Uniformly Distributed**
Parent Has BA or Higher	.282	.148
Family Income > $50K	.224	.120
Math Score > 61	.237	.110
Parent Expects Student > BA	.167	.136
Parent BA *or* Income > $50K	.160	.107
Parent BA *or* Income > $50K *or* Parent Expects Student > BA	.124	.090

	College Attend Rate	
Peer-Quality Measure:	**Actual Within-School Heterogeneity**	**High Peer-Quality Families Uniformly Distributed**
Parent Has BA or Higher	1.040	.549
Family Income > $50K	.849	.454
Math Score > 61	.909	.421
Parent Expects Student > BA	.695	.565
Parent BA *or* Income > $50K	.535	.357
Parent BA *or* Income > $50K *or* Parent Expects Student > BA	.363	.263

Table 8: Comparison: Other Predictors of Test Score

	Math	Reading
Total School Enrollment	-.0008***	-.0007***
	(.0003)	(.0002)
Class Size	-.038***	-.024***
	(.015)	(.010)
Starting Teacher Salary (1,000's)	-.024***	-.011*
	(.008)	(.058)
School % Free Lunch	-3.368***	-1.547***
	(.510)	(.368)
School Urban	-.139	-.084
	(.234)	(.162)
School Private	-1.583***	.388
	(.333)	(.251)
School % Single Parents	.011	-.002
	(.012)	(.006)
Student's Parent has BA (relative to HS Diploma)	1.210***	.815***
	(.261)	(.205)
Student's Parent is HS dropout (relative to HS Diploma)	-1.104***	-.659***
	(.268)	(.225)
Student's Family has Income $\geq$ $50,000	.189	.213
	(.213)	(.180)
Student's Family has $\geq$ 50 books at home	1.101***	.908***
	(.215)	(.165)
Student's Family has a computer	1.345***	.366***
	(.163)	(.134)
Student's Parent Expects Student to get > BA	1.260***	1.008***
	(.232)	(.176)
N	13,961	13,978
R^2	.536	.670

standard errors in parentheses,
clustered at the school level
$*$ significant at 10% level
$**$ significant at 5% level
$***$ significant at 1% level

Table 9: Comparison: Other Predictors of Dropout, College Attend.
School Dropout/College Attend. Rate Regressed on School-Level Variables

	Dropout Rate	**College Attend Rate**
Total School Enrollment	.0000	.0000
	(.0000)	(.0000)
School % Single Parents	-.001***	.0025***
	(.0004)	(.0006)
School Student/Teacher Ratio	-.0018	.0018
	(.0012)	(.0012)
School % in Remedial Reading	-.0178	.0334
	(.0632)	(.0598)
School % Free Lunch	.0757**	-.1393***
	(.0385)	(.0424)
School Urban	.0029	-.0033
	(.0139)	(.0598)
School Private	.0192	.0303
	(.0225)	(.0253)
School avg. Math Standardized Score	-.0014	.0029
	(.0021)	(.0025)
School avg. Reading Standardized Score	-.0069*	.0040
	(.0039)	(.0043)
N	880	880
R^2	.227	.458

standard errors in parentheses,
clustered at the school level
$*$ significant at 10% level
$**$ significant at 5% level
$***$ significant at 1% level

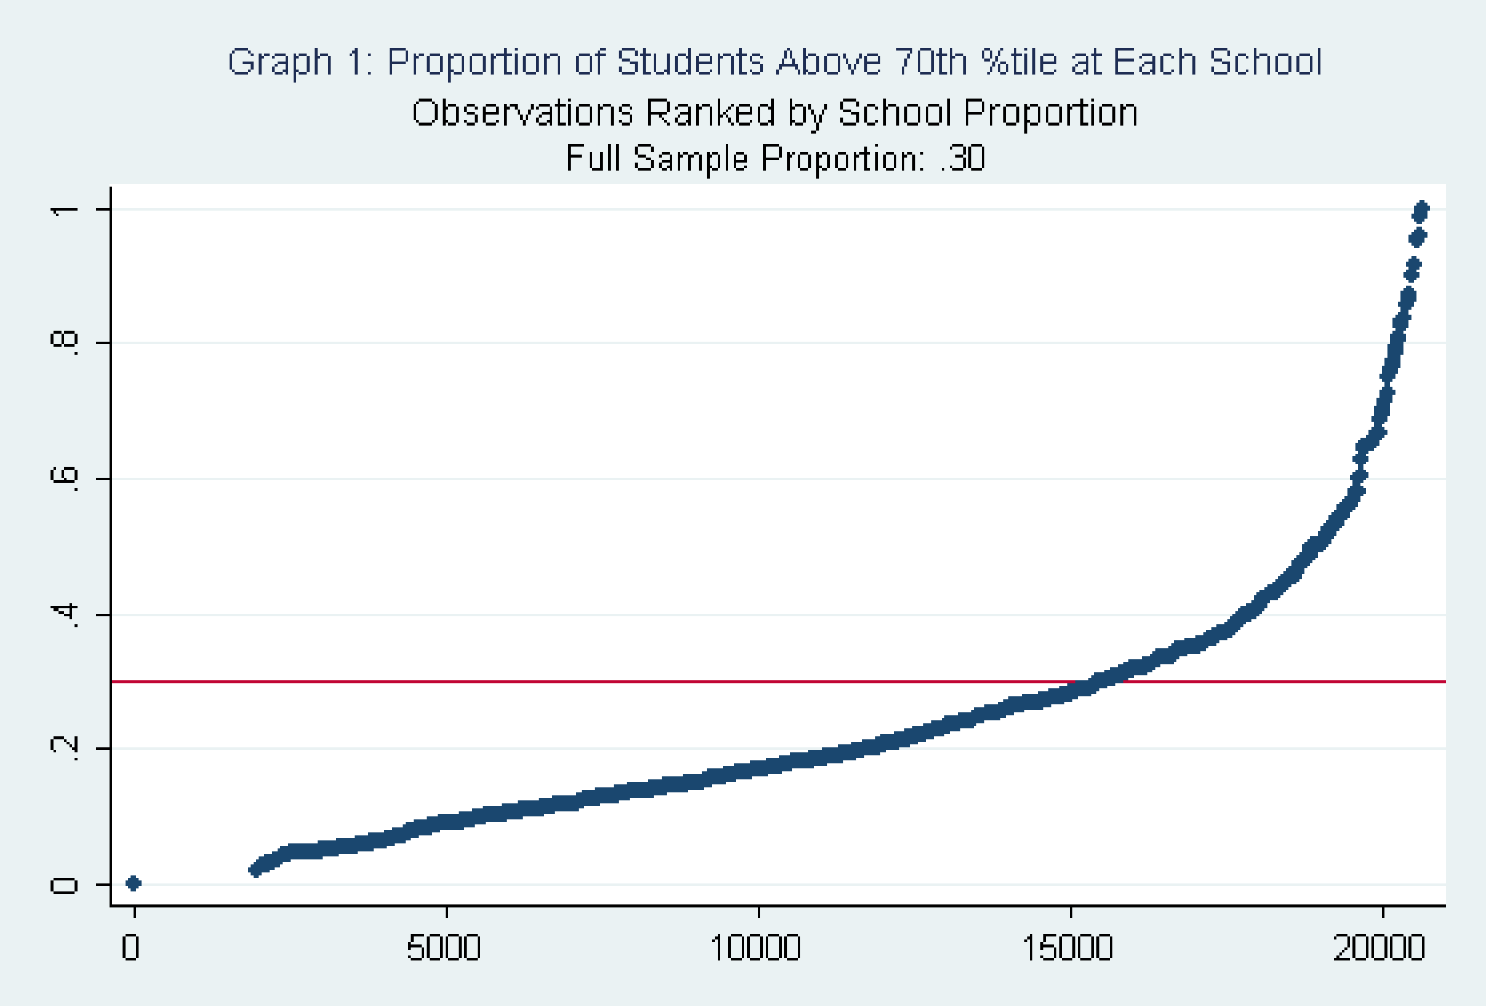
Graph 1: Proportion of Students Above 70th %tile at Each School
Observations Ranked by School Proportion
Full Sample Proportion: .30
0
.2
.4
.6
.8
1
0
5000
10000
15000
20000

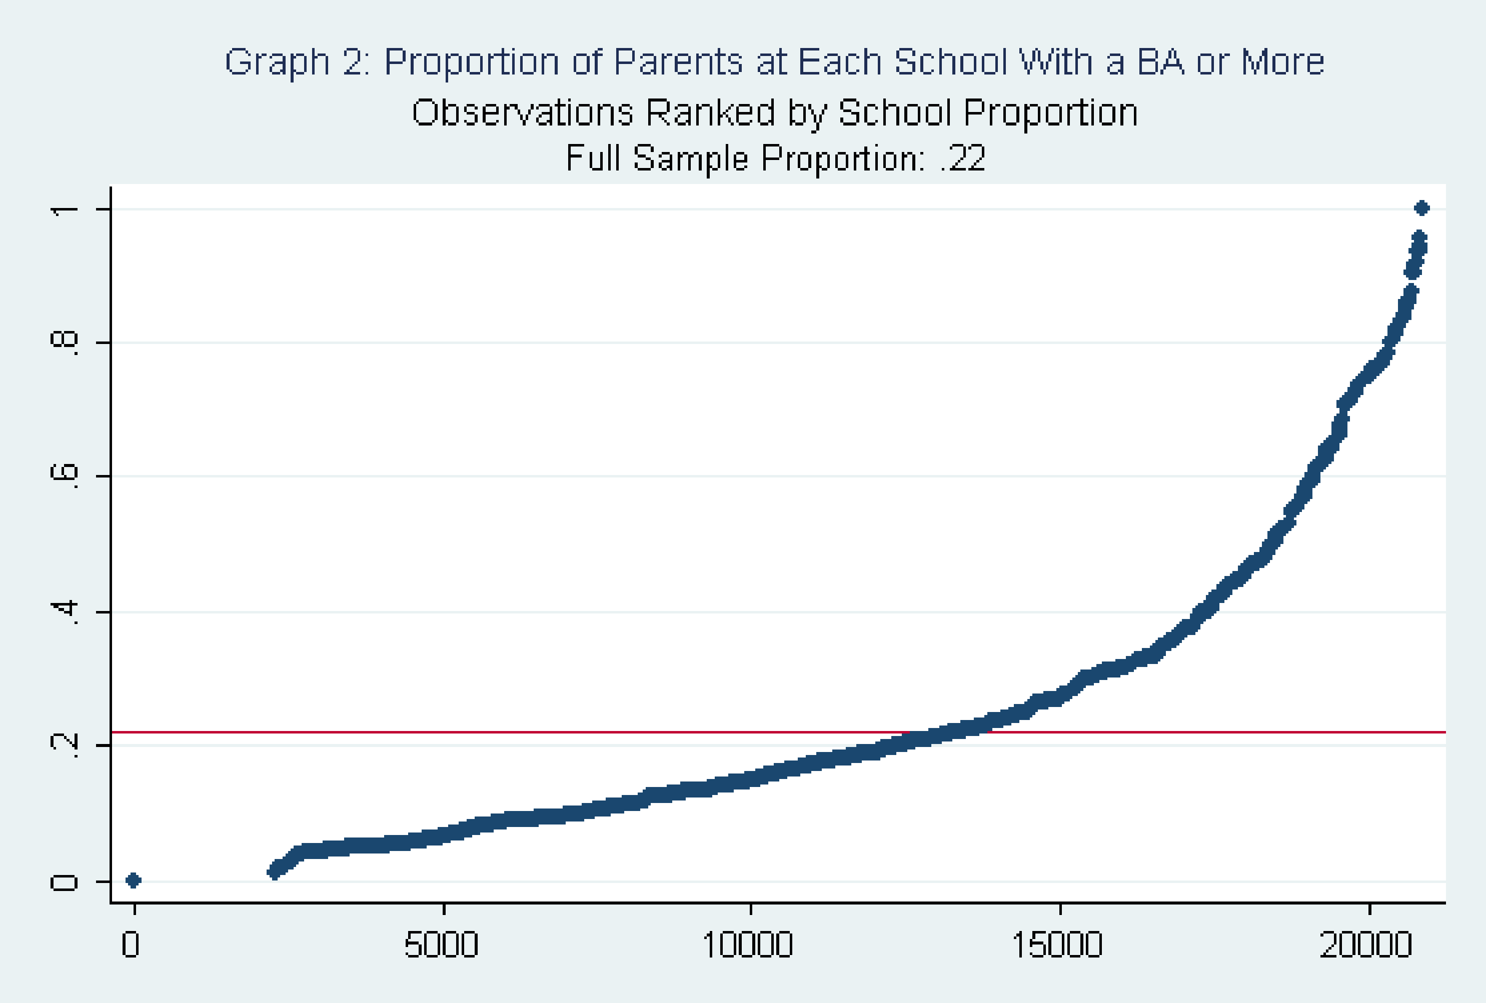
Graph 2: Proportion of Parents at Each School With a BA or More
Observations Ranked by School Proportion
Full Sample Proportion: .22
0
.2
.4
.6
.8
1
0
5000
10000
15000
20000

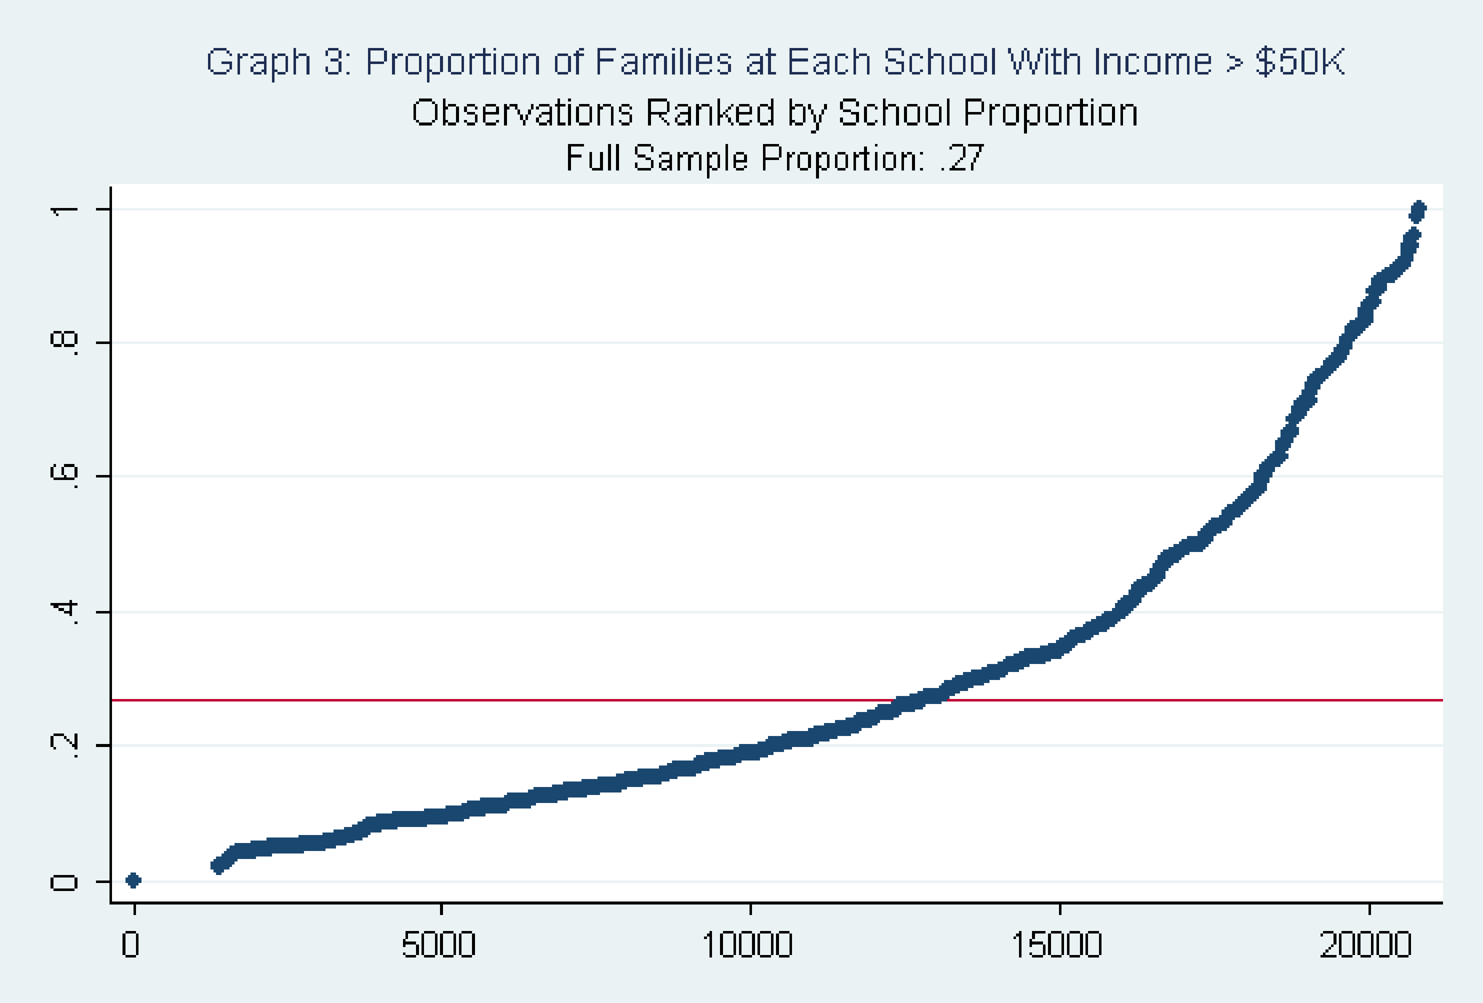
Graph 3: Proportion of Families at Each School With Income > $50K
Observations Ranked by School Proportion
Full Sample Proportion: .27
0
.2
.4
.6
.8
1
0
5000
10000
15000
20000

Figure 1: Reduced−Form Sorting Model

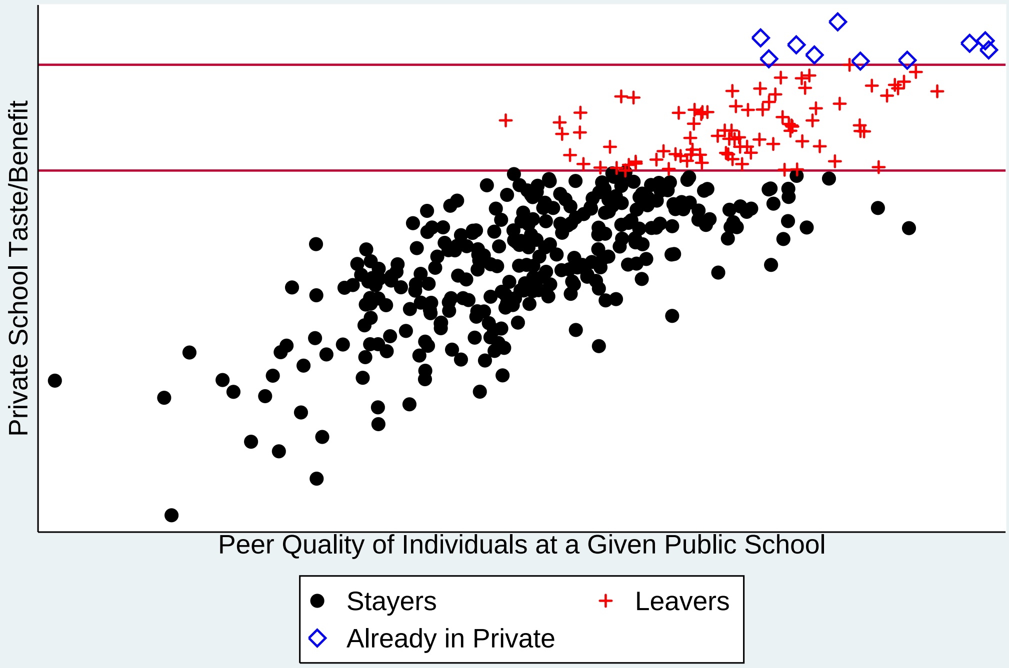

Figure 2: The Effect of Higher Variance

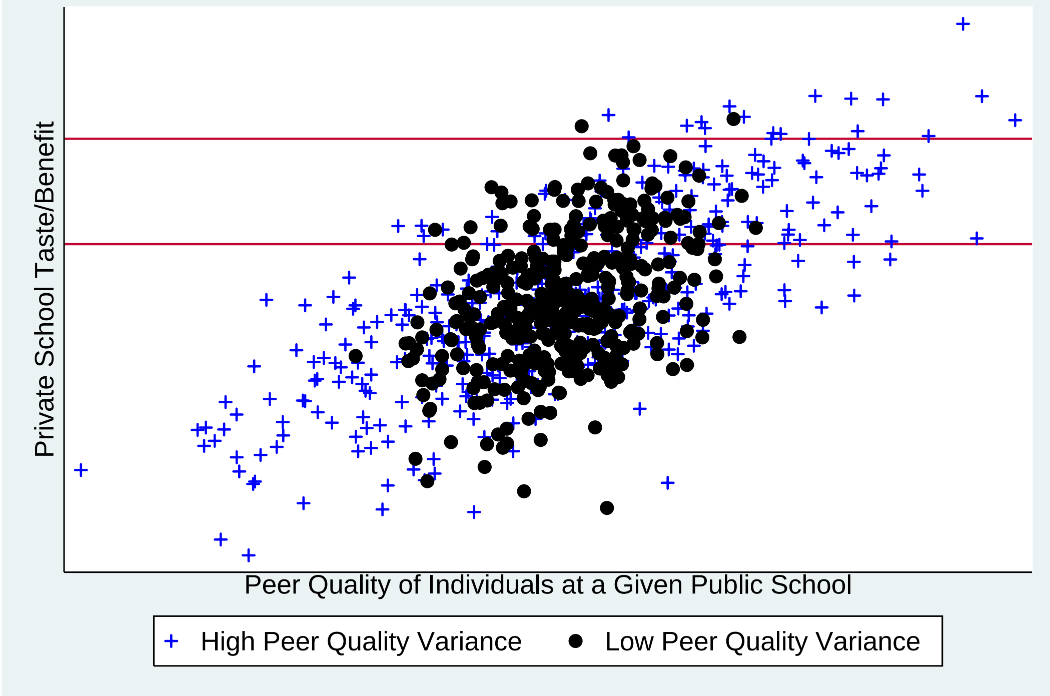

Figure 3

Schools Sorted by Share of Families With High Income

Representation; not from actual data

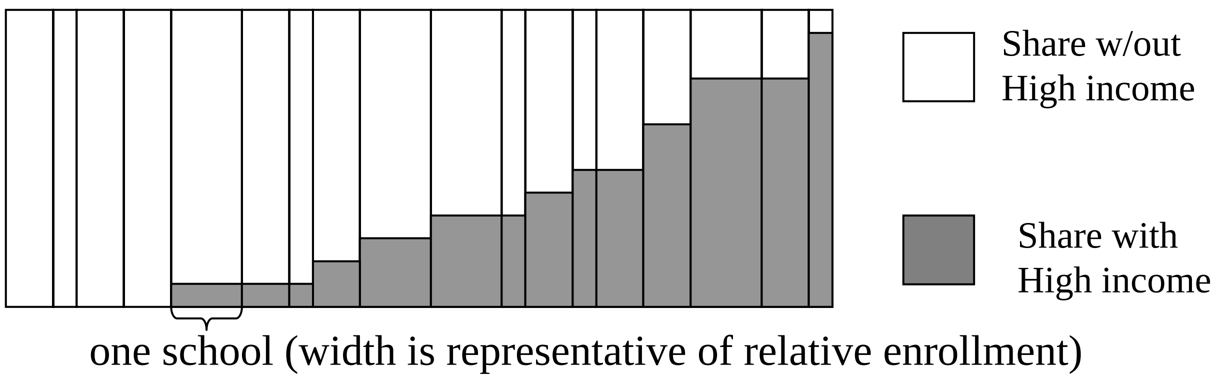

Figure 4

Synthetic “Perfect Sorting” Algorithm: Find “Cutoff” School such that Area B = Area C

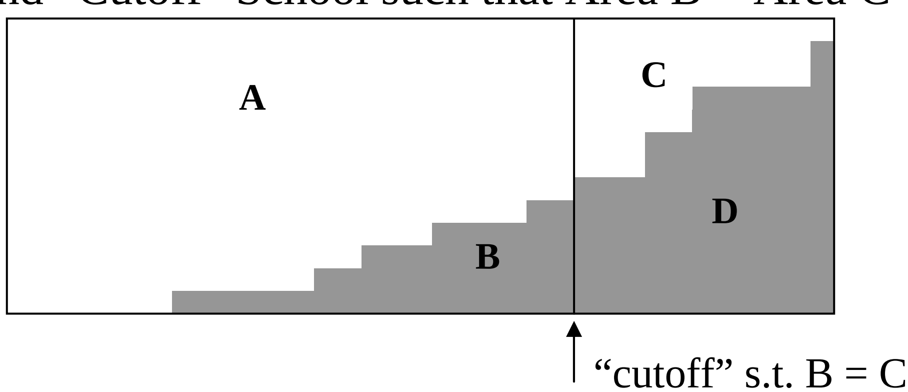

Figure 5

Perfect Sorting Implemented, Changes in School Composition Shown

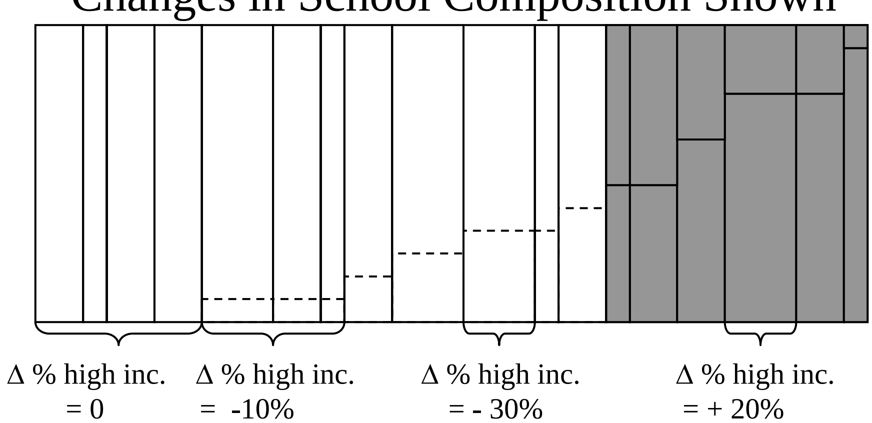

6 Appendices

6.1 Appendix A: Decomposition of $\Delta(\overline{Q_s})$

If N_t = total pre-sorting, N_v = *total leavers*, and N_s = *total stayers* = $N_t - N_v$, then:

$$\Delta(\overline{Q_s}) = \frac{1}{N_t-N_v}\sum_{i=1}^{N_t-N_v} Q_i - \frac{1}{N_t}\sum_{i=1}^{N_t} Q_i$$

$$\Delta(\overline{Q_s}) = \frac{1}{N_t-N_v}\sum_{i=1}^{N_t-N_v} Q_i - \left[\frac{N_v}{N_t}\frac{1}{N_v}\sum_{i=N_v}^{N_t} Q_i + \frac{N_t-N_v}{N_t}\frac{1}{N_t-N_v}\sum_{i=1}^{N_t-N_v} Q_i\right]$$

$$\Delta(\overline{Q_s}) = \frac{1}{N_t-N_v}\sum_{i=1}^{N_t-N_v} Q_i - \left[\frac{N_v}{N_t}\frac{1}{N_v}\sum_{i=N_v}^{N_t} Q_i + \frac{N_t}{N_t}\frac{1}{N_t-N_v}\sum_{i=1}^{N_t-N_v} Q_i - \frac{N_v}{N_t}\frac{1}{N_t-N_v}\sum_{i=1}^{N_v} Q_i\right]$$

$$\Delta(\overline{Q_s}) = \left[\frac{N_v}{N_t}\right]\left[\frac{1}{N_t-N_v}\sum_{i=1}^{N_t-N_v} Q_i - \frac{1}{N_v}\sum_{i=N_v}^{N_t} Q_i\right]$$

6.2 Appendix B: Proof *Sketch* For Sorting Results

Note that with $B_i = \lambda Q_i + \epsilon_i$, σ_B^2 is given by:

$$\sigma_B^2 = Var[\lambda Q + \epsilon] = \lambda^2\sigma_Q^2 + \sigma_\epsilon^2 \tag{1}$$

if ϵ is uncorrelated with Q. σ_B is therefore an increasing function of σ_Q.

Note that the correlation between Q and B is given by:

$$\rho = \frac{Cov[Q,\ B]}{\sigma_Q\sigma_B} = \frac{Cov[Q,\ \lambda Q + \epsilon]}{\sigma_Q\sqrt{\lambda^2\sigma_Q^2 + \sigma_\epsilon^2}} = \frac{\lambda\sigma_Q}{\sqrt{\lambda^2\sigma_Q^2 + \sigma_\epsilon^2}} \tag{2}$$

This implies that ρ is increasing in σ_Q.

1. See Figure 2 for illustration.
   - From (2), $\rho = 0$ implies that $\lambda = 0$. In this case, $[\overline{Q}_{stay} - \overline{Q}_{leave}] = 0$ because leavers are being drawn randomly.
   - $\rho = 1$ implies that $\sigma_\epsilon^2 = 0$. In this case, $[\overline{Q}_{stay} - \overline{Q}_{leave}]$ is maximized, since all families above a certain $\tilde{Q} = \frac{C}{\lambda}$ will leave, while all families below $\tilde{Q}$ will stay.
   - Given this relationship between ρ and $[\overline{Q}_{stay} - \overline{Q}_{leave}]$, since ρ is increasing in σ_Q, $[\overline{Q}_{stay} - \overline{Q}_{leave}]$ is also increasing in σ_Q.
2. See Figure 2 for illustration.

   The decision to leave for a private school is given entirely by the relationship between cost C and benefit B_i. Call $F_B(B)$ the CDF of B_i. [*share who leave*] is then given by:

$$[share\ who\ leave] = \frac{leavers}{leavers\ +\ stayers} = \frac{(1 - F_B(C_1)) - (1 - F_B(C_0))}{F_B(C_0)} \tag{3}$$

$F_B(C_0)$, the equivalent of (leavers + stayers), is decreasing in σ_B. This is because higher variance in B pushes more mass over the *initial* threshold C_0, reducing the total public sector enrollment at C_0. This gives the result that [*share who leave*] is rising in σ_B. Since σ_B is an increasing function of σ_Q, this implies that [*share who leave*] is rising in σ_Q.

6.3 Appendix C: Counterfactual ANOVA

For the full sample, the variance decomposition is calculated from the total sum of squares (TSS), model sum of squares (MSS) and residual sum of squares (RSS) of an ANOVA of a quality measure on school id:

$$Between\ school\ variance\ share = \frac{RSS}{TSS} \tag{4}$$

$$Within\ school\ variance\ share = \frac{MSS}{TSS} \tag{5}$$

"Low-SES" is defined as those schools where more than 20% of the students qualify for free lunches. When we extrapolate from the low-SES sample, we are asking "what would the decomposition look like if we retained the between-school variance from the full sample, but gave all schools the within-school variance of low-SES schools?" This is accomplished by calculating:

$$Between\ school\ variance\ share = \frac{[\frac{RSS_{low\ SES}}{RSS_{Total}}]}{[\frac{N_{low\ SES}}{N_{Total}}]} \frac{RSS_{Total}}{TSS_{Total}} \tag{6}$$

$$Within\ school\ variance\ share = (1 - Between\ school\ variance\ share) \tag{7}$$