Disentangling School and Peer Effects: A Non-linear Approach with School Lotteries

Mariana Laverde * Minseon Park† Suk Joon Son‡

September 2025

PRELIMINARY AND INCOMPLETE. PLEASE DO NOT CIRCULATE

Abstract

We develop a unified framework to disentangle school effects and peer effects on student achievement. Identification leverages exogenous tie-breaker lotteries in centralized school choice systems. To preserve statistical power while accounting for the potential non-linearity of peer effect, we summarize treatments and confounders parsimoniously: the full distribution of peers' baseline scores is captured by functional principal components (FPCA), and selection on observables is addressed by conditioning on expected peer distributions and school-specific admission probabilities; lottery–based deviations from these expectations provide instruments for realized peers. Applying the approach to New York City middle schools, we find a non-monotonic and non-linear peer effect: performance varies with the gap between classmates' and a student's own baseline scores, exhibiting both a rank effect and positive spillovers in parts of the distribution of peer baseline test scores. When we estimate school effects net of peers, the cross-school dispersion of value added rises by about 20% relative to conventional VA measures that omit peers, indicating that peer composition compresses VA estimates.

Keywords: Peer Effects, School Value-Added, Functional Principal Component Analysis, Educational Production Function

JEL Codes: I21, C31, I24, C26

*Boston College, laverdem@bc.edu.
†University of Michigan, minseonp@umich.edu.
‡University of Tokyo, sjson@e.u-tokyo.ac.jp.

1 Introduction

Disentangling school and peer effects is essential for evaluating policies that alter students' school assignments. Examples include shifts to open enrollment, expansions of the charter sector, changes in admission rules at selective schools, and district decisions to close schools. All such policies alter both the school students attend and the set of peers students encounter, and hence measuring their separate contributions is important for policy design and evaluation.

Relatedly, school accountability systems often rely on school value-added (VA) measures to assess performance and make high-stakes decisions such as closures. For such policies to be effective, VA must capture the causal contribution of the school itself rather than a mixture of school performance and peer effects. Despite its importance, much of the existing literature treats school effects and peer effects separately.

To fill this gap, we develop a unified framework that models test score production as a function of both the distribution of peer test scores and school effects. We show that both peer and school VA can be identified in our setting, and provide an estimation procedure that leverages assignment lotteries and can be applied to administrative data from centralized school choice systems. Crucially, our approach allows the full distribution of peer test scores—not only the mean—to influence outcomes. To summarize the distribution and make the estimation feasible, we apply functional principal component analysis (FPCA) and express the peer distribution as a combination of a finite number of principal components (Chang, Durlauf, Lee, and Park, 2023). To account for the endogeneity of the peer distribution, we use students' submitted preferences and priorities, summarized by the expected distribution of peer scores conditional on these confounders and the propensity scores for assignment into each school (Borusyak and Hull, 2023). Summarizing treatments and confounders in this way reduces the dimensionality of these objects and improves statistical efficiency.

We apply the method to data from New York City's centralized middle school choice system to study non-linear peer effects and school effects. Our results reveal that peer effects are non-linear, non-monotonic, and heterogeneous. Under the observed peer compositions, an incoming student who is in the bottom 10% of the achievement distribution experiences larger peer effect gains at the most and least selective schools in the city, while for a student in the top 10%, peer effects are lower at the most selective schools. We find that after netting out peer effects, value-added estimates are more dispersed than conven-

tional measures, suggesting a correlation between peer and school effects.

Empirical models of peer effects often focus on the influence of the average peer, rather than examining the impact of peers across different positions in the peer distribution. This approach is primarily chosen for convenience because assessing the full range of peer effects is complex. It involves handling the dimensionality of distribution functions and the corresponding lack of statistical power for accurate estimation (Hoxby, 2006). Moreover, the need for methods that leverage exogenous variation in peer composition further reduces power, making the estimation of flexible peer effect models prohibitive. On a substantive level, it is crucial to consider non-linear models of peer effects to address questions about student reassignments effectively. Under a linear model, student reassignments cannot affect aggregate outcomes by construction.

Our test score production function explicitly models both peer and school effects. The production function incorporates three key inputs: the school a student attends, the distribution of peers' ability (proxied by their past test scores), and the student's own ability. Our specification is general and can nest the often considered linear-in-means model. For added flexibility, we allow the peer and school effects to be heterogeneous across students along observable dimensions.

To address the endogeneity of observed peer composition and school enrollment, we leverage the conditionally random assignments generated by tie-breaking lotteries embedded in the student assignment algorithm. We use administrative data from New York City (NYC), focusing on public middle school students. NYC uses Deferred Acceptance (DA) to determine school assignments, taking students' submitted preferences and schools' admission priorities. Priorities are often coarse. In the case of ties, tie-breaking lotteries are used, creating sizable exogenous variation in school assignments and peer composition.

Our estimation strategy addresses two key econometric challenges that arise from the high-dimensional nature of both treatments and confounders. First, we summarize one component of the treatment—the peer distribution—using functional principal component analysis (FPCA), which captures the full distribution of peer test scores with only a few (three to five) basis functions while retaining most variation (93-97%). With this methodology, each student's peer density function can be approximated by the weighted sum of basis functions. These weights provide a finite-dimensional representation of the distribution that we use as regressors in our production function model.

Second, we summarize confounders using the recentering approach of Borusyak and Hull (2023). This procedure isolates exogenous lottery-based variation by taking the differ-

ence between realized and expected values of (1) school assignments and (2) co-assigned peer distributions. Realized values reflect both endogenous variation (from student preferences and school priorities) and lottery variation. The expected values—computed by averaging many runs of the centralized algorithm under simulated lotteries—capture only the endogenous component, so their difference isolates the exogenous lottery variation.

Our identification results clarify that certain features of our empirical analysis—such as non-linearity of peer effects or time-constant school effects with multiple cohorts—are not actually driving the identification. The results also show that we need two types of exogenous variation due to the lottery: (1) variation in school assignments and (2) variation in peer composition within schools. We show that such variation is available in our data: on average, 9.8% of students (over 21,000 students) switch schools across two sets of lottery draws, and among non-switchers, 7.7% (over 13,500 students) experience a peer-distribution change exceeding 5% in total variation distance.[1] This highlights that even students who do not face uncertainty in their school assignments may be subject to randomness in their peer composition.

Finally, because enrollments may diverge from assignments, we face a compliance issue. Since instrumenting hundreds of school indicators lacks sufficient statistical power (Angrist, Hull, Pathak, and Walters, 2024), we estimate the model using either (1) ordinary least squares (OLS), controlling for assignment propensity scores and expected peer distributions (our preferred specification), or (2) two-stage least squares (2SLS) instrumenting co-enrolled peer distribution with recentered co-assigned peer distribution while assuming school enrollment is conditionally independent.

Our empirical analysis reveals significant non-linearities and non-monotonicities in peer effects that vary systematically across student characteristics. For low-achieving students, we find a non-monotonic pattern: replacing 10% of same-scoring peers with peers scoring one standard deviation higher actually decreases their own test scores by 0.02 standard deviations, consistent with rank or relative position effects documented in prior literature (Carneiro, Aguayo, Salvati, and Schady, 2022; Carrell, Sacerdote, and West, 2013; Murphy and Weinhardt, 2020). However, the same students benefit when exposed to much higher-achieving peers (2.5 standard deviations above), gaining 0.02 standard deviations. In contrast, high-achieving students experience monotonically negative peer effects

[1] Total variation distance (TVD) equals one half of the L^1 distance between two discrete distributions and can be interpreted as the minimum fraction of mass that must be relocated to transform one distribution into the other. The TVD is calculated using 20 percentile-based bins of the peer score distribution.

when exposed to higher-achieving peers. Standard linear-in-means peer effect models often miss these patterns or fail to detect any significant peer effects.

These non-linear peer effects have implications for school VA measurement. When we estimate school effects net of peer effects rather than the conventional VA measures, we find more variation in school effects. This dispersion occurs because high-scoring students sort into schools with high conventional VA, but their presence creates negative peer effects for other high-achieving students, thereby compressing the observed VA measures.

Relation to the Literature Our work contributes to the literature on peer effects by examining the entire distribution of peer test scores, rather than focusing solely on the average peer characteristic or a limited set of moments. Most studies in this area have analyzed the effects of the mean peer test score or other summary statistics of the peer distribution. A few notable exceptions include Hoxby and Weingarth Salyer (2006), Gibbons and Telhaj (2008), and Lavy, Silva, and Weinhardt (2009). Unlike these approaches, which often represent non-linear effects by including covariates corresponding to groups of students along the peer score distribution, we employ Functional Principal Component Analysis (FPCA). This approach allows us to analyze peer effects flexibly while maintaining sufficient statistical power. Substantively, our analysis uncovers some non-linear patterns of peer influence that previous studies did not reveal. Specifically, we find consistent evidence of both rank effects and learning effects, aligning with recent experimental findings (Carneiro, Aguayo, Salvati, and Schady, 2022).

While random variation in school choice algorithms has been used to estimate school effectiveness (Abdulkadiroğlu, Angrist, and Pathak, 2014; Abdulkadiroğlu, Angrist, Narita, and Pathak, 2017, 2022; Angrist, Hull, Pathak, and Walters, 2024), our paper presents identification results that clarify when lottery-induced variation can be used to recover *both* peer and school effects. Our empirical strategy shows alternative assumptions that may be used when, in practice, ideal conditions are not satisfied. Also, our strategy builds on (Borusyak and Hull, 2023; Abdulkadiroğlu, Angrist, Narita, and Pathak, 2017) and develops a method to leverage the conditionally random assignments generated by centralized algorithms with known rules to identify peer effects. As other papers in the peer effects literature do, we leverage quasi-random peer assignments for identification of peer effects. This follows work by Boozer and Cacciola (2001) and Duflo, Dupas, and Kremer (2008) for primary and secondary education; and Sacerdote (2001); Zimmerman (2003); Carrell, Fullerton, and West (2009); Lyle (2009) at higher education levels.

Finally, our work is among a few studies that jointly identify and estimate school and peer effects. To our knowledge, only a small number of papers (Allende, 2019; Woodcock, Friesen, and Meilman Lomaz Cohn, 2025) are taking steps in this direction.

2 Model and Methods

2.1 Test Score Production Function

We consider the following production function of student i's test score:

$$y_i = \alpha^g + \int \beta^g(r) f_i(r) dr + \mu^g_{j(i)} + \varepsilon_i \tag{1}$$

where y_i is the outcome of interest for student i, $g \equiv g(i)$ is the demographic group that i belongs to, $r = \tau - \tau_i$ is the difference between a peer's baseline test score τ and the student i's own baseline test score τ_i, $f_i(r)$ is the probability density function of relative test scores r among i's co-enrolled peers leaving i out, $j(i)$ is the school that i enrolls in, μ_j is the school j's effect, and ε_i is the student's ability, explained both by observables and unobservables. While the parameters are indexed by the demographic group g, we drop the indices for the sake of exposition unless necessary.

$\beta(r)$ is interpretable as the effect per unit mass of peers at r. That is, for any two points r and r' in the support of f_i and probability mass m, $(\beta(r') - \beta(r))m$ is the effect on y_i of removing a mass m of peers at approximately r and adding the same mass m of peers at r'.[2] Location of $\beta(\cdot)$ is normalized so that $\int \beta(r) dr = 0$.[3]

The key flexibility in our model is that we allow the entire distribution of peer scores $f_i(r)$ to affect student outcomes. Our model nests many commonly studied specifications. The most prominent among them is the linear-in-means model:

$$y_i = \alpha + b \int r f_i(r) dr + \mu_{j(i)} + \varepsilon_i \tag{2}$$

[2] See Appendix A.1 for the proof. Approximation error arises from the finite sum approximation of an integral.

[3] The location of $\beta(\cdot)$ has no inherent interpretation and is not separately identified from the intercept. That is, we can always write $\alpha + k + \int (\beta(r) - k) f_i(r) dr$ for any constant k, as any density $f_i(\cdot)$ integrates to 1.

where b is a constant.[4] This is a special case of our model with $\beta(\tau) = b\tau$. More generally, any linear combinations of (centered) moments, which can be written as $\sum_{l=1}^{k} b_l \int \tau^l f(\tau) d\tau$, can be represented in our model with $\beta(\tau) = \sum_{l=1}^{k} b_l \cdot \tau^l$.[5]

We conceptually distinguish between peer and school effects as follows. School effect μ_j includes the effect of teachers, principals, and school infrastructure, among others. Peer effect captures, first, the direct interaction among students. But also, any response by the school based on the student composition of that year is also part of the peer effect. For example, teachers might change curriculum slightly from one year to another after knowing the distribution of academic preparedness of students the school enrolls.

Whether $f_i(\cdot)$ represents the distribution of peers' *relative* test scores $r = \tau - \tau_i$ or *absolute* test scores τ is not crucial. We estimate peer effects separately for each ability group, defined by a grouping of each student's own baseline test score τ_i. Within these groups, all students are evaluated against the similar baseline, which neutralizes the differences between absolute and relative scores.

2.2 Random Variation and Summarizing Treatment and Confounders

Admission Lotteries in Centralized School Assignment To deal with the endogeneity of observed peer composition *and* school enrollment, we leverage the conditionally random assignments generated by random tie-breaking embedded in the student assignment algorithm (Abdulkadiroğlu, Angrist, Narita, and Pathak, 2017). Our primary dataset from NYC allocates students using a centralized DA algorithm. To determine the assignments, DA takes in applicants' reported rank-ordered lists (ROLs) $(\succ_i)_i$ and priority groups $(\rho_{ij})_{ij}$, which defines the schools' preferences over applicants. However, often the priority groups are coarse, i.e., many students are tied within the same priority group. When the schools are oversubscribed and ties arise, the algorithm typically uses single-draw random lotteries $(\ell_i)_i$ to break the ties. While students' ROLs or priorities may be correlated with unobserved ability, the tie-breaking lotteries aren't, providing source of exogenous variation.

[4] More precisely, denoting τ to be the peer's "absolute" test score, the literature's specification is $y_i = \alpha + b \int \tau f_i(\tau) d\tau + \varepsilon_i$. However, noting that $r = \tau - \tau_i$ and denoting the density of r and τ by f_i^r and f_i^τ, we have $f_i^r(x) = f_i^\tau(x + \tau_i)$. Thus, $b \int x f_i^\tau(x) dx = b \int (x + \tau_i) f_i^r(x + \tau_i) dx = b \int r f_i^r(r) dr + b\tau_i$, meaning that as long as we include τ_i as a covariate, it is not important whether we specify the model in terms of r or τ under the linear-in-means framework.

[5] The centering may have to depend on groups. For example, within-group variance is captured by $\tau = r - \bar{r}_{g(i)}$ where $\bar{r}_{g(i)}$ denotes student i's (leave-one-out) group mean and $k = 2$.

While the DA mechanism clearly provides conditionally random variation in school assignments, there remains at least three key econometric issues. The first issue is the statistical efficiency. In our setting, both confounders and treatments are extremely high-dimensional. Regarding the high dimensionality of confounders, note that a single student's assignment and peers can, in principle, depend on the rank-ordered lists and priorities of *all* students. Letting $\theta \equiv (\theta_i)_i \equiv (\succ_i, (\rho_{ij})_j)_i$ denote the vector of all students' submitted preferences and priority groups, this creates a potentially enormous set of confounding variables that must be controlled for to isolate the exogenous variation from lotteries. Note also that one component of the treatment, the peer distribution, is a *functional* treatment, which is nearly[6] infinite-dimensional.

Below, we discuss how we summarize the high-dimensional peer distribution treatment and the confounders.

Summarizing Peer Distribution with Functional Principal Components We have an infinite-dimensional input—the distribution of peers' test scores—in our structural equation 2. To reduce the dimensionality, we use a finite number of basis functions to approximate the distribution f_i. Representation of a function with a finite-length vector allows us to perform regression analyses in the usual way.

To achieve statistical efficiency, we adopt the approach of Chang, Durlauf, Lee, and Park (2023), utilizing the leading functional principal components of $(f_i)_i$ as our basis. This method minimizes the mean squared approximation error in the L^2 sense[7] and ensures that such errors do not cause omitted variable bias. The orthogonality property of principal components eliminates multicollinearity among regressors, enhancing efficiency. In our application, FPCA finds a succinct set of basis functions that best capture the variation in the distribution of peers' baseline test scores.

We follow Chang, Durlauf, Lee, and Park (2023) to consider the estimation of the benchmark model 1 by FPCA. Most of the discussion below is adapted from their work. To estimate the benchmark equation 1, we represent the functional covariate $(f_i)_i$ by finite-dimensional vectors.

Define inner product $\langle \beta, f_i \rangle := \int \beta(r) f_i(r) ds$. Using any orthonormal basis $(v_j) \equiv$

[6] We say *nearly* because the empirical distribution of test scores are nearly continuously distributed.

[7] The mean squared approximation error is minimized over all possible m-dimensional linear subspaces of the function space, where m is the number of principal components used.

$(v_j)_{j=1}^{\infty}$ of the functional space H[8] for which $f \in H$, f can be written as

$$f = \sum_{j=1}^{\infty} \langle v_j, f \rangle \, v_j.$$

Due to orthonormality, $\langle v_j, f \rangle$ is the projection of f onto v_j, or equivalently it is the *coordinate* of f in the direction of v_j. Because we use FPCA to obtain the orthonormal basis, our (v_j) are *principal components* of $(f_i)_i$. The ordering of (v_j) is such that $m-$leading principal components are $(v_j)_{j=1}^{m}$. Using the leading m functional principal components $(v_j)_{j=1}^{m}$, we may approximate $(f_i)_i$ as

$$f_i \approx \sum_{j=1}^{m} \langle v_j, f_i \rangle \, v_j$$

with some m.

To represent the functional covariate $(f_i)_i$ by m-dimensional vectors, we use the mapping π on H defined as

$$\pi(f) := \begin{pmatrix} \langle v_1, f \rangle \\ \vdots \\ \langle v_m, f \rangle \end{pmatrix} \in R^m.$$

That is, π is the mapping that takes in a function f and returns its m coordinates $\pi(f)$ in the directions of $(v_j)_{j=1}^{m}$ as their $m-$dimensional representation. Note that the peer effects term in our test score production function can then be approximated as

$$\int \beta(r) f_i(r) ds \equiv \langle \beta, f_i \rangle \approx \left\langle \beta, \sum_{j=1}^{m} \langle v_j, f_i \rangle \, v_j \right\rangle = \sum_{j=1}^{m} \langle v_j, \beta \rangle \, \langle v_j, f_i \rangle = \pi(\beta)' \pi(f_i).$$

Therefore, the structural equation 2 can be approximated as

$$y_i \approx \alpha + \pi(\beta)' \pi(f_i) + \mu_{j(i)} + \varepsilon_i \tag{3}$$

which suggests that we may use $\pi(f_i)$ as regressors to estimate $\pi(\beta)$. Denoting the estimate

[8] As in their paper, we only consider the functions β and $(f_i)_i$ that are square integrable and belong to a Hilbert space H with inner product $\langle \beta, f_i \rangle = \int \beta(r) f_i(r) ds$.

of $\pi(\beta)$ by $\widehat{\pi(\beta)}$, we can obtain an estimate for the functional parameter $\beta(\cdot)$ by

$$\hat{\beta}(\cdot) := \sum_{j=1}^{m} \widehat{\pi(\beta)}_j v_j(\cdot) \simeq \sum_{j=1}^{m} \pi(\beta)_j v_j(\cdot) \simeq \beta(\cdot),$$

where the first approximation sign is due to the estimation error and the second is due to the projection error. Chang, Durlauf, Lee, and Park (2023) shows that $\hat{\beta}(\cdot)$ is a consistent estimator of $\beta(\cdot)$ projected onto $(v_j)_{j=1}^m$, under regularity conditions and if $\widehat{\pi(\beta)}$ is consistent for $\pi(\beta)$. While the basis functions $(v_j)_j$ are built to best approximate $(f_i)_i$ rather than $\beta(\cdot)$, the Monte Carlo analysis in Figure A.1 demonstrates that $\beta(\cdot)$ is also well approximated by these basis functions.

Summarizing High-Dimensional Confounders High-dimensional confounders is a well-recognized challenge in statistics and econometrics. One standard approach is to construct a summary of the confounders, such as the propensity score, that preserves desirable statistical properties including conditional independence (Rosenbaum and Rubin, 1983). In centralized school choice settings, Abdulkadiroğlu, Angrist, Narita, and Pathak (2017) apply similar ideas by computing school assignment probabilities to control for non-random student preferences and school priorities. Our setting presents additional complexity as we must handle functional (or continuous) treatments representing peer distributions (or their finite-vector representation) in addition to high-dimensional binary treatments (school enrollments). While propensity scores are well-suited for binary treatments, they do not readily extend to continuous treatments.

We resolve this challenge by utilizing the recentering approach developed by Borusyak and Hull (2023) that allow for continuous treatments. Their approach works with "formula treatments"—treatments computed from multiple sources of variation according to a known rule. Our setting is a textbook example: both school assignments and peer compositions are deterministic functions of students' preferences, school priorities, and random lottery draws. The key insight is to isolate the purely exogenous lottery-based variation by "recentering" the assigned treatments around their expected values. The expecation is taken over the exogenous variation—the lottery draws in our case—and hence serves as a control function that captures purely endogenous variation.

More specifically, let $s_i^a \in \{0,1\}^J$ denote the vector of binary indicators representing student i's *assigned* (offered) school, with J being the number of schools. Let f_i^a denote the

corresponding co-assigned peer distribution. Let s_i and f_i represent the *enrolled* school and co-enrolled peer distribution. The DA assignments can be written as known functions of all students' preferences and priorities $\theta = (\theta_{i'})_{i'}$ and the vector of random lottery draws $\ell = (\ell_{i'})_{i'}$:

$$\begin{pmatrix} s_i^a \\ f_i^a \end{pmatrix} = \begin{pmatrix} g_i(\theta, \ell) \\ h_i(\theta, \ell) \end{pmatrix} \tag{4}$$

This known functional form is crucial because it allows us to simulate counterfactual assignments by drawing alternative lottery realizations ℓ' while holding the observed θ fixed, allowing us to calculate the expected assignment vector (propensity scores) and peer distributions $E[s_i^a|\theta] =: p_i$ and $E[f_i^a|\theta]$. In practice, we approximate these expectations by simulating the DA algorithm across 1,000 independent lottery draws and computing sample averages of the resulting assignments and peer distributions. Intuively, these expectations cancel out all the lottery-induced random variation, thereby reflecting purely endogenous variation.

The recentered assignments $s_i^a - E[s_i^a|\theta]$ and recentered assigned peer distributions $f_i^a - E[f_i^a|\theta]$—which reflect purely exogenous variation—can then serve as instruments for the endogenous enrolled school and peer distribution (s_i, f_i). Provided relevance of the recentered instruments, Borusyak and Hull (2023) show that the resulting instrumental-variables estimator achives consistency for true parameters provided relevance and regularity conditions.

A question that may arise is whether lotteries provide *separate* variation for peer and school effects. We turn to this issue next.

2.3 Identification

Here we discuss the separate identification of peer and school effects. We abstract away from statistical power considerations for now and return to this in the next subsection. The goal of this subsection is twofold. The first is to demonstrate that the variation generated by the lotteries are not "collinear", in the sense that peer and school effects can be separated out. The second is to clarify that certain features of our empirical analysis—such as non-linearity of peer effects or time-constant school effects with multiple cohorts—are not driving the identification. While they help, they are not necessary.

To clarify the sources of identification, we consider a more flexible model than what

we use for empirical analysis:

$$y_{ij} = \mu_j^{\tau_i, g(i)} + \int \beta^{\tau_i, g(i)}(\tau) f_j(\tau) d\tau + \epsilon_i.$$

This model is more flexible compared to our empirical analysis in several respects. First, note that instead of $\beta^{g(i)}(r)$ with $r = \tau - \tau_i$ representing the difference between a peer's baseline test score τ and the student i's own test score τ_i, we have $\beta^{\tau_i, g(i)}(\tau)$. $\beta^{\tau_i, g(i)}(\tau)$ is a more general form than $\beta^{g(i)}(r)$. Relatedly, $f_j(\tau)$ now denotes the (absolute) peer test score distribution within school j. To make the identification even more (ostensibly) difficult, we do not use the leave-one-out distribution for the peer distribution. Also, school effects $\mu_j^{\tau_i, g(i)}(i)$ are allowed to vary across $(\tau_i, g(i))$ rather than just $g(i)$.

Note that we only observe y_{ij} for i's enrolled school $j(i)$:

$$y_{ij(i)} = \sum_j s_{ij} \underbrace{\left(\mu_j^{\tau_i, g(i)} + \int \beta^{\tau_i, g(i)}(\tau) f_j(\tau) d\tau \right)}_{\delta_j^{\tau_i, g(i)}} + \epsilon_i.$$

Condition on individuals with $\tau_i = \tau', g(i) = g$. For all schools j, consider $s_{ij}^a - E_l[s_{ij}^a|\theta]$ as an instrument for s_{ij}. Exogeneity of the instrument with respect to the errors is satisfied following the arguments of Borusyak and Hull (2023). Assume that rank condition is satisfied, so that there is nondegenerate variation in the instrument for all j. Then, the usual IV argument of instrumenting s_{ij} with $s_{ij}^a - E_l[s_{ij}^a|\theta]$ allows us to identify $\delta_j^{\tau', g}$ for all (τ', g).

Having identified $\delta_j^{\tau', g}$, we turn to separating the two terms—peer and school effects—within it. Fix (τ', g) as before. Consider now a school-level equation:

$$\delta_j^{\tau', g} = \int \beta^{\tau', g}(\tau) f_j(\tau) d\tau + \mu_j^{\tau', g}$$

Consider instrumenting $f_j(\tau)$ with $f_j^a(\tau) - E_l[f_j^a(\tau)|\theta]$ for all τ. Assuming that rank condition holds, the terms in this equation can be identified by the usual IV argument. Instrument exogneity is satisfied following the arguments of Borusyak and Hull (2023).

Note that same argument may be applied to the linear-in-means specification, regardless of whether one leaves out the student i.

Taking stock, student-level and school-level lottery-induced variation allow us to disentangle peer and school effects. In particular, the two effects being "subsumed" by the

school fixed effects in the student-level IV regression can be resolved by the second-stage IV regression at the school level. While two rank conditions are assumed, they are less about the separability of the two effects, and more about the power of the instruments to provide variation for each "element" *within* each of the two effects (j for school effects and τ for peer effects).

The identification is achieved even without leave-one-out distributions, non-linearity of peer effects, availability of multiple cohorts data (and assuming time-constant school effects), or assuming constant school effects within groups g. On the other hand, we do have these additional sources of identification, which only adds to the identification strength.

2.4 Empirical Strategy

One practical issue of identifying our production function is that many schools in centralized assignment systems are undersubscribed, meaning they admit all applicants and are subject to little or no lottery-based variation; in such a case, $s^a_{ij} - E[s^a_{ij}|\theta]$ is nearly constant for such j. This is a well-recognized limitation in the literature (Angrist, Hull, Pathak, and Walters, 2024). Therefore, we use the Conditional Independence Assumption (CIA) as in the literature:

Assumption 1 (Risk Controlled Conditional Independence Assumption (RC-CIA)).

$$\varepsilon_i \perp\!\!\!\perp s_i \mid (x_i, p_i)$$

This assumption states that, conditional on observed characteristics and *propensity scores*, a student's unobserved ability is independent of their school enrollment. Propensity scores are admissions probability for each school, which can be calculated given that school assignment procedure is known. Put differently, RC-CIA imposes that observable characteristics x_i derives noncompliance with the assignment. Our x_i includes students' 5th grade standardized test scores (lagged test scores), gender, ethnicity, Free-or-Reduced lunch eligibility, special education status, and English Learner status. Note that, under perfect compliance, this assumption is satisfied automatically (Angrist, Hull, Pathak, and Walters, 2024), and this is what makes RC-CIA different from CIA in other value-added models (Rosenbaum and Rubin, 1983; Chetty, Friedman, and Rockoff, 2014). In that sense, RC-CIA is a statement about compliance with the assignment.[9] Angrist, Hull, Pathak, and

[9] Angrist, Hull, Pathak, and Walters (2024) also shows that CIA is implied by $\varepsilon_i \perp\!\!\!\perp s_i \mid (x_i, p_i, s^a_i)$.

Walters (2024) validates the assumption by comparing the results under RC-CIA with IV estimates based on randomized school offers.

In a similar spirit to RC-CIA, we also consider the following assumption:

Assumption 2 (Augmented Risk Controlled Conditional Independence Assumption (aRC-CIA)).

$$\varepsilon_i \perp\!\!\!\perp (s_i, f_i) \mid (x_i, p_i, E[f_i^a|\theta])$$

Relative to RC-CIA, the main content of the assumption is that the a student's co-enrolled peer distribution is independent of their unobserved ability, conditional on the observed characteristics, propensity scores, expected co-assigned peer distribution, and enrolled school.[10] There would be a violation of aRC-CIA if there is coordinated compliance decision on unobserved ability not explained by $x_i, p_i, s_i, E[f_i^a|\theta]$. Like the RC-CIA, this assumption is essentially a compliance condition: under perfect compliance, the covariance-zero version is automatically satisfied. The aRC-CIA assumption motivates the following OLS specification with:

$$y_i = \sum_{k=1}^{m} b_k c_k^{ij} + \sum_j \mu_j s_{ij} + \sum_{k=1}^{m} \rho_k E[z_k^{ij}] + \sum_j \gamma_j p_{ij} + X_i'\alpha + \epsilon_i \tag{5}$$

where c_k^{ij} are FPC coordinates of realized peer distribution, $E[z_k^{ij}]$ are expected coordinates of the co-assigned peer distribution, s_{ij} are school indicators, and p_{ij} are propensity scores. This represents our main specification for the empirical analyses. As discussed in section 2.2, given the estimates $(\hat{b}_k)_k$ of $(b_k)_k$, we can derive an estimate of the functional parameter $\beta(\cdot)$ by $\hat{\beta}(\cdot) := \sum_{j=1}^{m} \hat{b}_k v_k(\cdot)$.

Alternatively, the RC-CIA assumption motivates the following 2SLS procedure that instruments realized distribution with co-assigned distribution, controlling for expected

[10] More specifically, aRC-CIA can be decomposed into the following two assumptions:

$$\varepsilon_i \perp\!\!\!\perp s_i \mid (x_i, p_i, E[f_i^a|\theta])$$
$$\varepsilon_i \perp\!\!\!\perp f_i \mid (x_i, p_i, s_i, E[f_i^a|\theta])$$

co-assigned distribution (Borusyak and Hull, 2023):

$$\text{Second stage: } y_i = \sum_{k=1}^{m} b_k c_k^{ij} + \sum_j \mu_j s_{ij} + \sum_{k=1}^{m} \rho_k E[z_k^{ij}] + \sum_j \gamma_j p_{ij} + X_i'\alpha + \epsilon_i \tag{6}$$

$$\text{First stage: } c_k^{ij} = \delta_k z_k^{ij} + \text{controls} + e_{ik} \tag{7}$$

where z_k^{ij} are the FPCA coordinates of the assigned peer distribution. Given the relatively limited exogenous variation compared to OLS, estimating this equation separately for different demographic groups is challenging. Accordingly, we employ 2SLS as a robustness specification in our analyses.

3 Estimates of Peer and School Effects

3.1 Data

We use middle school student-level data from the New York City Department of Education (DOE) as the main data set for the analysis. Our sample consists of the 2014-15, 2015-16, and 2016-17 middle school application cohort. Across three cohorts, we have 138,318 students in the analysis sample after some data cleaning. We have data on students' test scores on the New York State Math and English Language Arts (ELA) tests. We also observe students' submitted rank-ordered lists, final school assignments, admissions priorities at each school program, and demographic information. The demographic information includes student race/ethnicity, English Language Learner status, and subsidized lunch status.

Panel A of table 1 shows summary statistics of student characteristics. The mean of baseline test score is 0.046 and 0.923, showing that our analysis sample is slightly higher achieving than the whole population of NYC students who took the standardized tests. 50% of students in our analysis sample are male, 21%, 41%, and 18% are Black, Hispanic, and Asian, respectively.

We estimate peer effect at each school by cohort level. There are 507 unique schools, and 1,464 school by cohort. Table 1 shows summary statistics at school- and student-level. As shown in panel B, an average school has 94 schools in their 6th grade, and the standard deviation of the cohort size is 89. Comparing summary statistics of student characteristics at the student- and school-level, the pattern shows that schools that larger schools have

higher achieving, and less minority students.

Table 1: Summary Statistics

	Mean (1)	SD (2)
Panel A: Student Characteristics (N=138318)		
5th Grade Score	0.046	0.923
Male	0.504	0.500
Black	0.213	0.410
Hispanic	0.413	0.492
Asian	0.188	0.391
FRL	0.790	0.408
Special Education	0.208	0.406
ELL	0.097	0.296
Panel B: School Characteristics (N=1464)		
6th Grade Size	94.480	89.577
Mean 5th Grade Score	-0.144	0.508
Share Male	0.503	0.116
Share Black	0.311	0.281
Share Hispanic	0.437	0.269
Share Asian	0.115	0.169
Share FRL	0.841	0.198
Share Special Education	0.242	0.100
Share ELL	0.099	0.098

Note: This table reports summary statistics for the student- and school-by-cohort-level variables used in the analysis. The 5th grade score is the average of students' Math and English scores on the New York State standardized test. FRL indicates Free- or Reduced-Price Lunch eligibility, and ELL indicates English-Language Learner status.

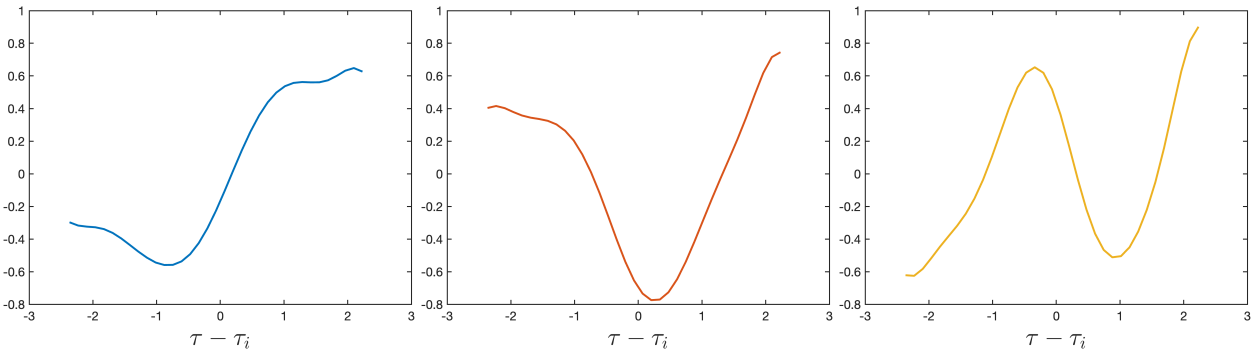
3.2 Principal Components, and Variation for Identification

In this section, we show that we have still sufficient variation to obtain peer effects even after controlling for a rich set of characteristics in our main specification 5, and such a variation is as good as random by testing the identification assumptions.

Principal Components Figure 1 displays the first three functional principal components (FPCs) of the co-enrolled peer score distribution, which summarize the main dimensions

of variation in peers' baseline test scores. The first component captures overall shifts in peer quality, essentially reflecting differences in mean peer achievement. The second component emphasizes dispersion around the mean, distinguishing between peer groups that are more homogeneous versus those with greater spread. The third component highlights asymmetry, capturing whether the peer distribution is skewed toward higher- or lower-achieving students. Together, these three components account for more than ninety percent of the total variation in peer distributions, providing a parsimonious yet rich representation of the underlying functional treatment as explained in section 2.2.

Figure 1: The First Three Principal Components



Notes: The three panels display the first three functional principal components (FPCs) (v_j) used to approximate the peer–effect function $f_j(\tau)$. Each curve is an eigenfunction of the estimated covariance operator from our FPCA step, centered and L^2-normalized (sign is arbitrary). These orthonormal basis functions span the space in which we estimate $\beta(\cdot)$.

Identifying Variation In section 2.3, we established that lottery-induced variation provides separate identifying variation for peer and school effects, *provided that* the two rank conditions at the student and school level, respectively.

Here, we investigate the feasibility of this identification approach with finite data. We ask three questions: whether there exists sufficient variation to (1) separate the effects of peers and schools, (2) identify the effects of *all* schools, and (3) identify the effects of peer *distribution*.

Before proceeding to empirical results, we first examine why DA lotteries may provide more identifying variation than initially apparent. As emphasized by Abdulkadiroğlu, Angrist, Narita, and Pathak (2017), even schools with strict priority orderings (i.e., no within-priority ties and thus no direct use of lottery tie-breakers) can experience effectively random fluctuation in enrolled composition through rejection chains. Applicants who are

randomly rejected at oversubscribed, lottery-using schools are redirected to lower-ranked options, propagating randomness to schools that otherwise would not use tie-breakers.

Notably, our focus on peer effects provides supplementary sources of variation: we leverage randomness not only in students' own assignments but also in the assignments of their peers. Even when student i's own school assignment is effectively deterministic, the composition of i's peers can remain random because the students co-enrolling with i are subject to lotteries at the schools they rank.

Now we turn to empirical evidence regarding the three questions. The first question asks whether there exists sufficient separate variation for both peer and school effects. This would be the case if there are sufficiently many students who, due to the lotteries, (1) switch schools, and (2) among non-switchers, experience nonnegligible changes in their peer distribution. Across two lottery draws, 9.8% ($\approx$ 21,000 students) of students switch schools, and among the non-switchers, 20.3% experience a distributional change with total variation distance (TVD[11]) ≥ 0.05.

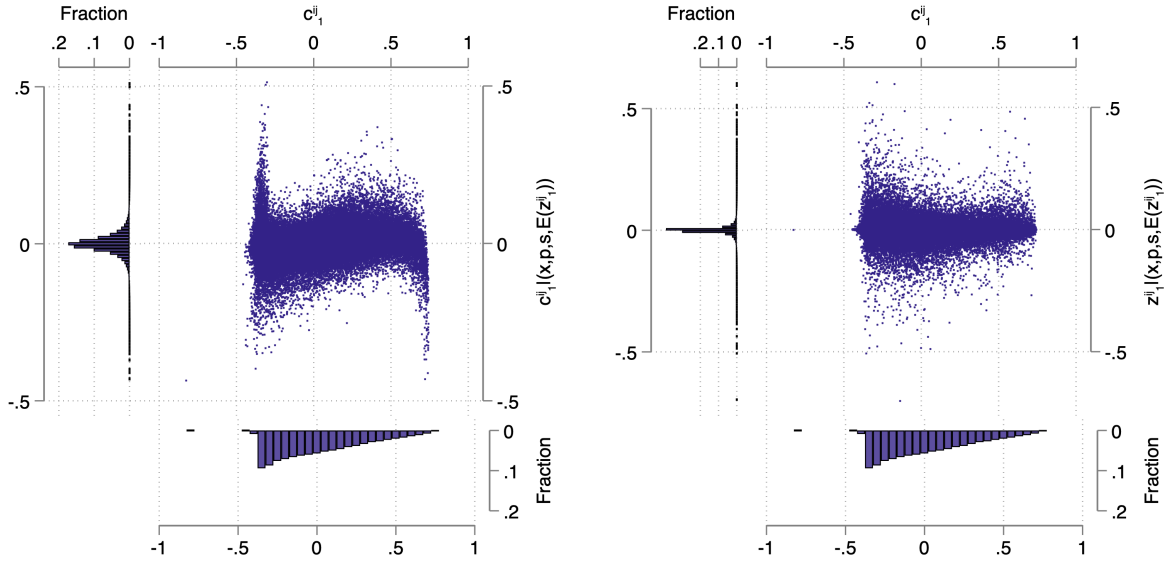
The second question concerns whether there exists sufficient variation to identify the effects of all individual schools. We address this challenge using established solutions discussed in subsection 2.4.

When it comes to the third question, figure 2 illustrates the sources of variation that underpin our identification strategy for peer effect, focusing on the first component for the sake of space. In both panels, the horizontal axis shows the raw coordinates of the first functional principal component of co-enrolled peer's baseline test score distribution. The vertical axis differs across panels: in the left panel we plot the residual variation in this component after controlling for expected peer distributions, enrolled school dummies, assignment propensity scores, and rich student covariates x_i. This is precisely the variation we exploit in our OLS specification, net of predictable or endogenous components. In the right panel, we plot the recentered coordinates of distribution of co-assigned peer's baseline test scores that serve as instruments in our 2SLS specification, capturing the exogenous variation induced by the admission lotteries. As the figure shows, the usable variation is naturally smaller than the raw dispersion of peer characteristics, but there remains substantial heterogeneity across students. In particular, the OLS residuals continue to span a wide range, and even the IV-based variation retains nontrivial spread, confirming that we

[11]TVD equals one half of the L^1 distance between two discrete distributions and can be interpreted as the minimum fraction of mass that must be relocated to transform one distribution into the other. See Appendix Table **??** for a full comparison across raw, binned, and smoothed representations.

have sufficient statistical leverage for identification.

Figure 2: Raw vs. Identifying Variation



Next, we turn to ask if variations shown in figure 2 is as good as random, by testing our identification assumptions. From aRC-CIA (Assumption 2), we can derive a weaker condition as follows.

$$\varepsilon_i \perp\!\!\!\perp (s_i, f_i) \mid (x_i, p_i, E[f_i^a|\theta]) \quad \Rightarrow \quad \varepsilon_i \perp\!\!\!\perp f_i \mid (x_i, p_i, E[f_i^a|\theta], s_i) \tag{8}$$

This condition is infeasible to test because first, ε_i is unobserved, and second, f_i is high-dimensional.

Instead, we test the following condition.

$$w_i \perp\!\!\!\perp c_k^{ij} \mid (x_i, p_i, E[c_k^{ij}|\theta], s_i) \tag{9}$$

First, although ε_i is unobserved, we have a rich set of observable characteristics, w_i, that are not included in x_i. We include students' 5th grade standardized test scores (mean of Math and English), their ethnicity, English-learner status, Free-or-Reduced lunch eligibility, gender, and special education status in x_i School district fixed effects are further included.

Importantly, our identification concern mainly arises from non-compliance. Thus, we include a set of characteristics that are not part of x_i, but one might be suspicious to cause

different behavior in compliance — students' attendance record, the length of submitted rank ordered list, and the log median income of their residential census tract. Furthermore, we replace distribution with FPCA coordinates as in our main specification.

Table 2 reports test results. Columns (1)-(3) present test results for equation 9. In each column, we regress the share of present days in students' 5th grade, the length of submitted middle school application list, and the log of median income of students' residential census tract on $\{c_k^{ij}\}$, controlling for expected peer distribution coordinates, propensity scores, school enrollment dummies, and baseline observable characteristics included in x_i. We present p-values from the test that coefficients in front of $\{c_k^{ij}\}$ are jointly null. The joint p-values indicate that balance is well satisfied, except for that that the test is marginally rejected in column (1). At the same time, a few individual coefficients on the FPCA coordinates show statistical significance in some specifications. We interpret this not as evidence against validity but as a natural byproduct of testing multiple high-dimensional instruments, where occasional significant coefficients can arise by chance even when the joint test shows no systematic imbalance.

Columns (4)-(6) further checks if our instrumental variable, coordinates from re-centered distribution of co-assigned peers, are correlated with variables in w_i. These results provide a benchmark given that the discrepancy between the co-assigned peers and the expected co-assigned peers distribution is only generated by random tie-breaker, as explained in section 3.2. Overall, the joint tests provide reassurance that our identifying variation is plausibly exogenous and consistent with the assumptions underlying our empirical strategy.

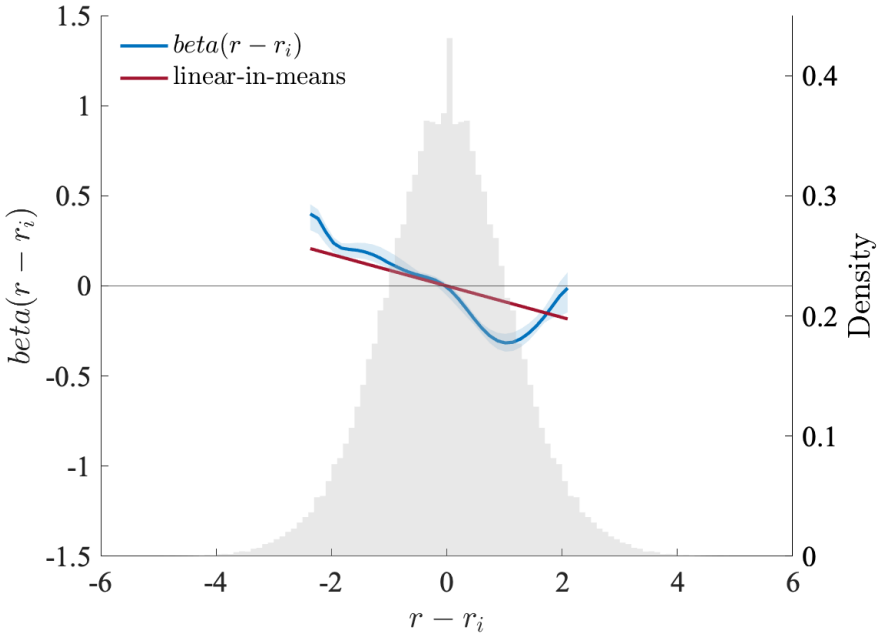
Table 2: Tests for Identification Assumptions

	aRC-CIA Assumption			Validity of IV		
	Attendance (%)	List Length	log(Med. Income)	Attendance (%)	List Length	log(Med. Income)
Joint test: p-value	0.071	0.161	0.780	0.214	0.242	0.638
c_1^{ij}	-0.011	-0.431	0.013			
	(0.006)	(0.282)	(0.043)			
c_2^{ij}	-0.002	0.058	-0.002			
	(0.002)	(0.211)	(0.013)			
c_3^{ij}	-0.004	0.214	0.003			
	(0.004)	(0.324)	(0.027)			
$z_1^{ij} - E[z_1^{ij}\|\theta]$				-0.006	0.270	0.025
				(0.005)	(0.228)	(0.053)
$z_2^{ij} - E[z_2^{ij}\|\theta]$				-0.013*	0.113	0.129
				(0.006)	(0.350)	(0.066)
$z_3^{ij} - E[z_3^{ij}\|\theta]$				0.003	0.638	0.475**
				(0.009)	(0.472)	(0.099)
N	123,757	125,686	125,257	123,757	125,686	125,257
R^2	0.180	0.405	0.623	0.137	0.159	0.216

Notes: Each column reports coefficients from regressions of the specified pre-treatment outcome on the first three FPCA coordinates of realized peer distribution (testing for aRC-CIA) or recentered peer distribution (testing for IV validity). Attendance (%) is the share of present days out of total school days in students' 5th grade, list length is the number of schools listed in submitted middle school applications, and log(Med. Income) is the log of median income of census tracts where students reside in their 5th grade. Columns (1)-(3) control for expected peer distribution coordinates, propensity scores, school enrollment dummies, and baseline observable characteristics included in x_i. p-values correspond to wild-bootstrap tests of the joint significance of coordinates. Standard errors, clustered at the year by school district level, are shown in parentheses. * $p < 0.05$, ** $p < 0.01$.

3.3 Peer Effect

Figure 3: Non-linear Peer Effect under aRC-CIA



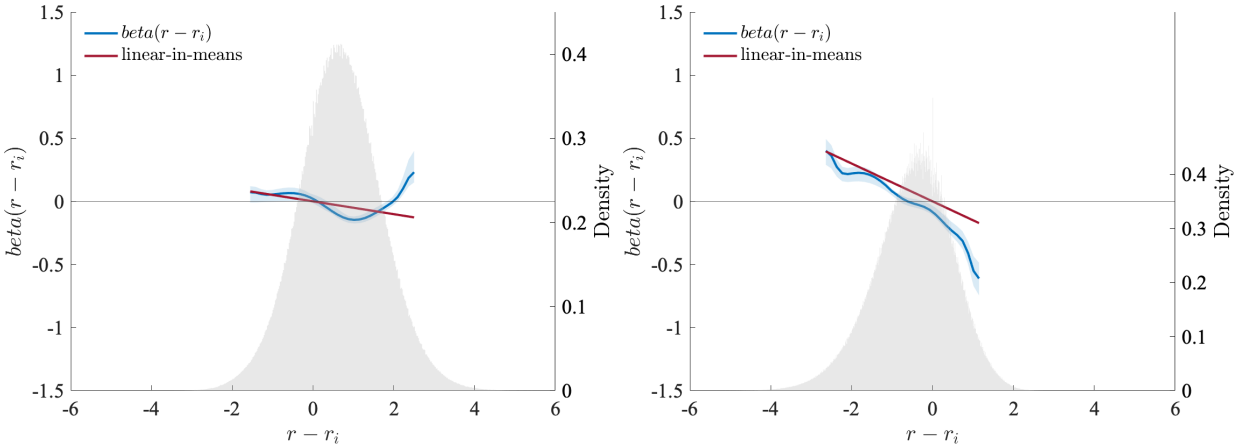
Notes: The figure plots the estimated peer effect function $\hat{\beta}(\tau - \tau_i)$ under the augmented Conditional Independence Assumption (aRC-CIA). The horizontal axis represents the difference between peers' baseline test scores and the student's own baseline score $(\tau - \tau_i)$. The blue curve shows the non-parametric estimate based on functional principal components with 95% confidence intervals, while the red line represents the linear-in-means specification. The shaded gray area shows the empirical density of $(\tau - \tau_i)$ in the data. Standard errors are computed using a wild bootstrap clustered at the school-market level.

Figure 3 presents our main estimate of peer effects under the augmented conditional independence assumption (aRC-CIA). The blue curve plots the nonparametric estimate of $\beta(\tau - \tau_i)$, the marginal effect of peers whose baseline achievement differs from student i's own baseline by $r - r_i$. For comparison, the red line shows the corresponding estimate from a conventional linear-in-means model. The shaded gray area depicts the empirical density of $r - r_i$, i.e., the distribution of relative peer achievement in the data.

The figure highlights two key findings. First, peer effects are distinctly non-linear: students are negatively affected by peers slightly stronger than themselves but benefit from exposure to much higher-achieving peers. This produces the inverted-S shape evident in the blue curve. For example, replacing 10% of one's peers who are only one standard deviation higher with peers two standard deviations lower increases the student's achievement by about 0.1 standard deviations. By contrast, replacing those same peers

with much higher-achieving peers—two standard deviations above—yields a smaller but still positive effect of about 0.05 standard deviations. In other words, students appear to be discouraged when surrounded by slightly stronger peers but gain when exposed to much stronger peers. Thus, our result is consistent with both rank effect found in other settings (Murphy and Weinhardt, 2020; Carneiro, Aguayo, Salvati, and Schady, 2022) *and* positive peer effect, especially from peers who are way better achieving than students themselves ('shining light'). This inverted-S shape stands in stark contrast to the red linear-in-means line, which forces a monotone relationship and thus completely misses this pattern. Importantly, these results are robust to varying the number of functional components used in estimation as well as across student subgroups. Second, the linear-in-means specification misses these patterns, instead imposing a monotone relationship that averages away the heterogeneity across the distribution.

Figure 4: Peer Effect by Own Test Scores



Notes: The figure plots the estimated peer effect function $\hat{\beta}(\tau - \tau_i)$ separately for low-performing (left panel) and high-performing (right panel) students, defined by if a student's baseline test score is above median or not. The horizontal axis represents the difference between peers' baseline test scores and the student's own baseline score $(\tau - \tau_i)$. The blue curve shows the non-parametric estimate based on functional principal components with 95% confidence intervals, while the red line represents the linear-in-means specification. The shaded gray area shows the empirical density of $(\tau - \tau_i)$ within each subgroup. Standard errors are computed using a wild bootstrap clustered at the school-market level.

Figure 4 explores heterogeneity in peer effects by student baseline achievement. The left panel shows estimates for low-performing students, while the right panel focuses on high-performing students. For low-performing students, the peer effect curve (blue) is non-monotonic: exposure to slightly higher-achieving peers initially reduces achieve-

ment, but larger jumps in peer quality generate positive gains. In contrast, high-performing students display a more consistently negative slope—being surrounded by stronger peers tends to lower their own outcomes. The discrepancy between our results and the linear-in-means specification is more pronounced among low-performing students given the non-monotonic and nuanced patterns in peer effects. Together, these results highlight that the consequences of shifting peer composition depend strongly on a student's own starting point: weaker students may benefit from exposure to substantially stronger peers, while stronger students are more likely to be discouraged by upward shifts in peer quality. We show more heterogeneity results by gender, subsidized lunch status, and baseline score quartiles in Figures A.4, A.5, and A.6, respectively.

Figure A.3 reports the estimates of $\beta(\cdot)$ from the 2SLS procedure, as specified in Equations 6 and 7. While the confidence intervals are wider—reflecting the less available exogenous variation relative to OLS—the overall pattern resembles that in Figure 3, exhibiting a generally negative slope, albeit with greater steepness.

3.4 School Effect

Next, this section compares school effects from our unified model to those from a value–added–only specification and documents a systematic compression of school effects when peer effects are ignored.

To do so, we consider the following value-added-only production function (Chetty, Friedman, and Rockoff, 2014).

$$y_i = \tilde{\alpha} + \tilde{\mu}_{j(i)} + \tilde{\varepsilon}_i \tag{10}$$

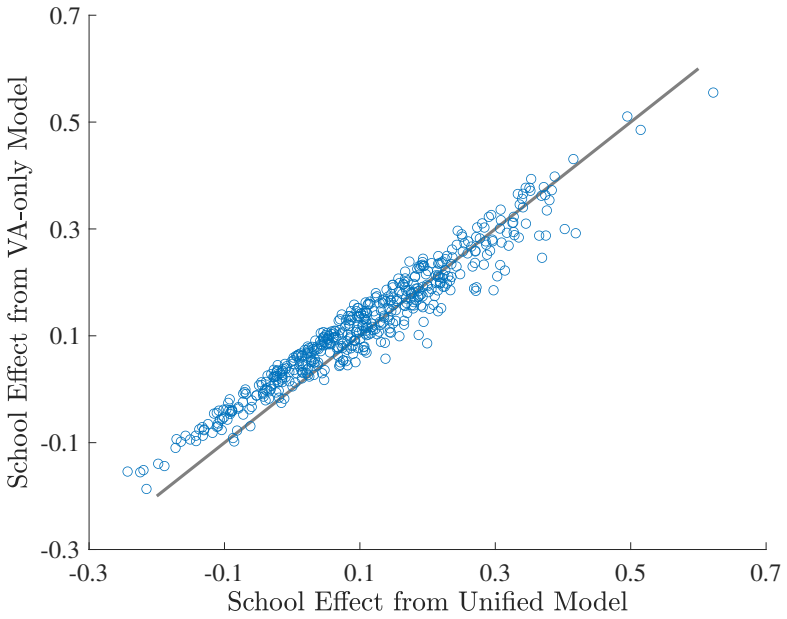
To estimate the causal school effects, we run the regression below, controlling for propensity scores as we described in section 2.4.

$$y_i = \sum_j \tilde{\mu}_j s_{ij} + \sum_j \tilde{\gamma}_j p_{ij} + X_i' \tilde{\alpha} + \tilde{\epsilon}_i \tag{11}$$

Notably, school effects from this equation are what are estimated in Angrist, Hull, Pathak, and Walters (2024).

Then we shrink both school effects from our unified model in equation 5 and from value-added only specification in equation 11 toward the grand mean using an empirical–Bayes (EB) rule.

Figure 5: School Effect from Unified Model vs. Value-added Only Model



Notes: Each dot is a school. The horizontal axis reports school value added from our unified model that jointly estimates school and peer effects. The vertical axis reports school value added from a value-added–only specification that omits peer effects. For both approaches, school effects are Empirical-Bayes shrunken toward the grand mean to reduce noise from small schools. The gray line is 45-degree line.

Figure 5 compares schools' effect on students' test scores from our unified model on the horizontal axis—which jointly identifies school and peer effects—to a school VA-only specification on the vertical axis that omits peer effects. Points lie close to the 45° line, so rankings are broadly similar, but the fitted line is flatter (slope < 1), indicating compression: high-VA schools are pulled down and low-VA schools pushed up when peer effects are ignored. The standard deviation of school effect estimates from our unified model is 0.13, while that from VA-only specification is 0.11.

When peer composition is omitted, the VA-only regression necessarily loads systematic differences in peer effect into the school effect. Two empirical facts in our data make the resulting bias compressing rather than inflating: (i) peer impacts are predominantly negative (section 3.3), and (ii) schools with truly higher value added enroll higher-achieving peers on average. Taken together, the missing peer term pulls high-VA schools *down* (because they also have strong peers with overall negative peer effect) and pushes low-VA schools *up* (because they have weaker peers with overall positive peer effect). The distri-

bution of estimated school effects is therefore squeezed toward the center: high performers are understated, low performers overstated, and the fitted relationship between VA-only and unified estimates has slope below one with a smaller cross-school standard deviation.

4 Varying School Assignments

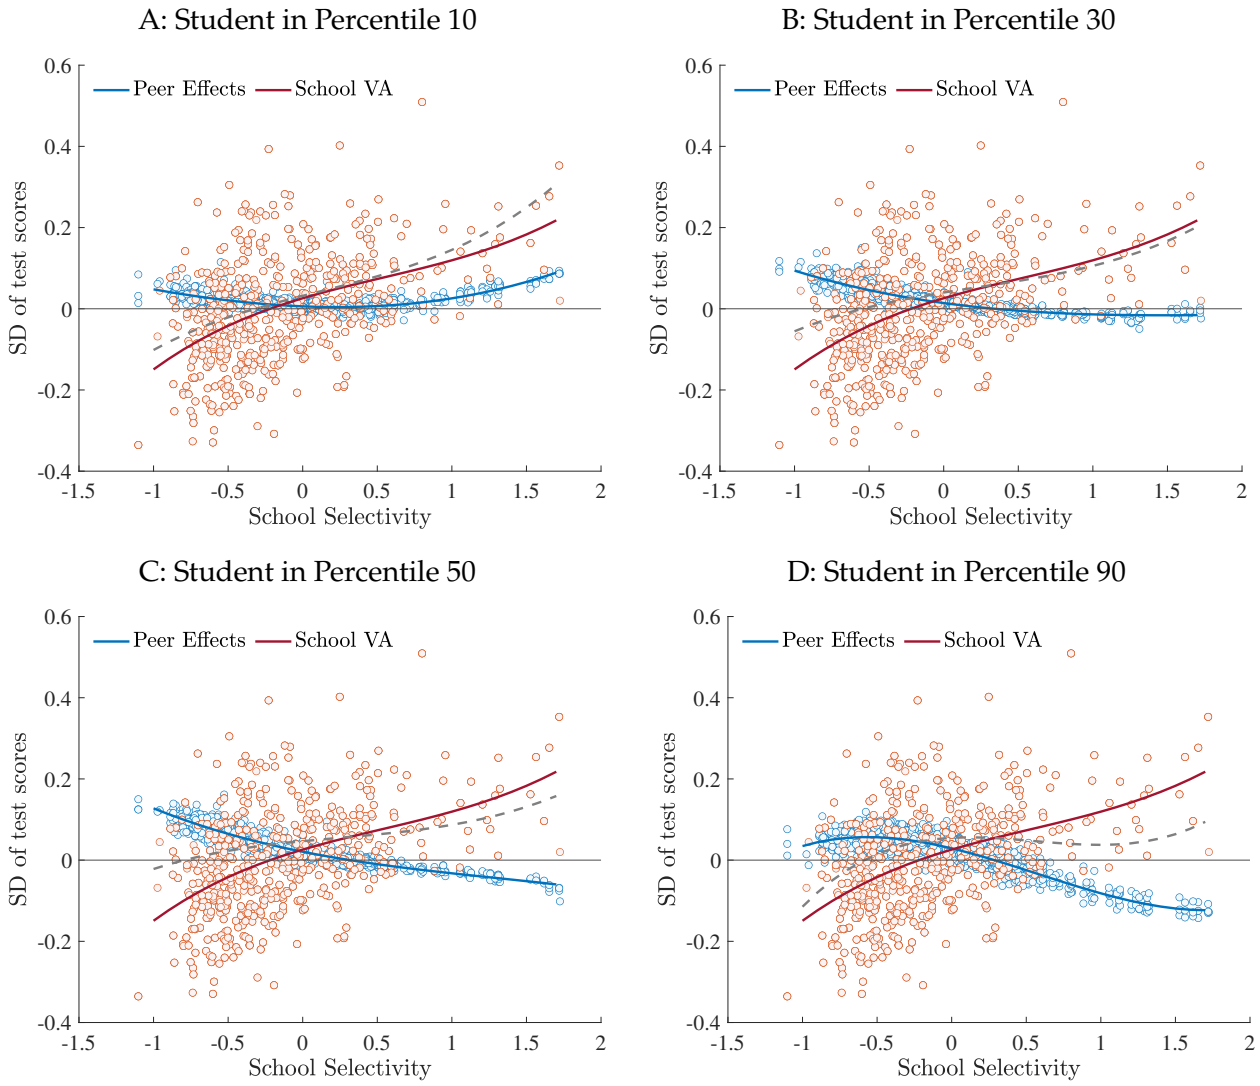
We have argued that by separately estimating peer values and school effects, we can better characterize the value of certain school assignments. In this section, we further explore the contribution of peers and schools to student achievement under alternative assignments of students to schools.

We first consider the allocation of a new student into the system, and simulate the peer and school values for the entrant student at all schools in the city, assuming the observed student assignments are in place. Taking the main estimates of peer and school values from the homogeneous model (Figure 3), we find a non-linear and non-monotone relationship between peer-values and school selectivity –measured by the mean of a school's enrolled student test scores. An incoming student in the bottom 10 percent of the achievement distribution experiences more positive peer effects at the most selective schools and the least selective schools. Since school VA is estimated to increase with selectivity, such a student is expected to experience larger test score gains at the most selective schools (Figure 6, Panel A).

Conversely, for higher-achieving incoming students, the relationship between peer effects and selectivity is decreasing: these students experience larger peer gains in less selective schools. For incoming students in the 90th percentile of achievement, the larger school effects observed in more selective schools are counterbalanced by reduced peer effects when compared to less selective schools. This implies that students at this achievement level could see comparable academic gains whether they attend moderately selective or the most selective schools (Figure 6, Panel D).

Figure 6 also reveals that, holding the observed student assignments, the peer effects distribution tends to be more compressed than the school VA distribution. The standard deviation of school VA is 0.13, and the average standard deviation of peer effects across percentiles of baseline achievement is 0.035. Less than a third of the dispersion in school VA.

Figure 6: Peer Effect and School VA in the Observed Assignment by Achievement Level



Note: Each figure shows scatter plots of school selectivity and school VA (in red) and selectivity and peer effects (in blue) across all schools in our sample, for an incoming student in the 10th, 30th, 50th, or 90th percentiles of the baseline achievement distribution. Solid lines represent fitted cubic relations between the variables in each case. The dashed gray line is the sum of the red and blue solid lines. The VA and the peer value at the school with the average school VA are normalized to zero in each plot. This means, gains in scores are expressed relative to gains at the school with mean VA.

5 Conclusion

This paper proposes a unified framework for the identification and estimation of school and peer effects. Our proposed methodology leverages FPCA to estimate the impacts of the entire distribution of peers on student outcomes, rather than just the mean. By exploit-

ing exogenous variation stemming from lottery-based school assignments, we demonstrate under which conditions both peer and school effects are identified.

Our findings reveal substantial heterogeneity and non-linearity in peer effects. Specifically, low-achieving students exhibit significant gains when exposed to peers in the right tail of the achievement distribution relative to same-scoring peers, but experience losses with higher-achieving, but not exceptional, peers. In contrast, high-achieving students experience monotonically negative peer effects when exposed to higher-achieving peers. These results challenge conventional linear-in-means models and underscore the importance of modeling the full peer distribution to accurately capture peer influences on learning. Moreover, our results highlight the importance of considering peer effects when evaluating policies that alter students' school assignments.

Furthermore, the analysis indicates that accounting for peer effects alters the distribution of estimated school effects, with implications for school accountability and value-added measurement. When netting out peer influences, we observe greater variability in school effects, reflecting underlying student sorting mechanisms and compositional heterogeneity within schools.

Methodologically, our approach delineates the conditions under which lottery-induced exogenous variation can be used to disentangle peer effects from school effects, offering a framework applicable to various educational contexts with centralized assignment systems. This contributes to the ongoing discourse on causal inference in peer effects research and provides practical tools for policy evaluation.

References

Abdulkadiroğlu, A., J. Angrist, and P. Pathak (2014): "The Elite Illusion: Achievement Effects at Boston and New York Exam Schools," *Econometrica*, 82(1), 137–196.

Abdulkadiroğlu, A., J. D. Angrist, Y. Narita, and P. Pathak (2022): "Breaking Ties: Regression Discontinuity Design Meets Market Design," *Econometrica : journal of the Econometric Society*, 90(1), 117–151.

Abdulkadiroğlu, A., J. D. Angrist, Y. Narita, and P. A. Pathak (2017): "Research Design Meets Market Design: Using Centralized Assignment for Impact Evaluation," *Econometrica : journal of the Econometric Society*, 85(5), 1373–1432.

Allende, C. (2019): "Competition Under Social Interactions and the Design of Education Policies," p. 104.

Angrist, J., P. Hull, P. A. Pathak, and C. Walters (2024): "Credible School Value-Added with Undersubscribed School Lotteries," *Review of Economics and Statistics*, 106(1), 1–19.

Boozer, M., and S. E. Cacciola (2001): "Inside the'Black Box'of Project STAR: Estimation of Peer Effects Using Experimental Data," *Available at SSRN 277009*.

Borusyak, K., and P. Hull (2023): "Nonrandom Exposure to Exogenous Shocks," *Econometrica*, 91(6), 2155–2185.

Carneiro, P., Y. C. Aguayo, F. Salvati, and N. Schady (2022): "The Effect of Classroom Rank on Learning throughout Elementary School: Experimental Evidence from Ecuador," *WORLD (Oakland, Calif. : 1993)*, 1.

Carrell, S. E., R. L. Fullerton, and J. E. West (2009): "Does Your Cohort Matter? Measuring Peer Effects in College Achievement," *Journal of Labor Economics*, 27(3), 439–464.

Carrell, S. E., B. I. Sacerdote, and J. E. West (2013): "From Natural Variation to Optimal Policy? The Importance of Endogenous Peer Group Formation," *Econometrica*, 81(3), 855–882.

Chang, Y., S. N. Durlauf, S. Lee, and J. Y. Park (2023): "A Trajectories-Based Approach to Measuring Intergenerational Mobility," Discussion paper, National Bureau of Economic Research.

Chetty, R., J. N. Friedman, and J. E. Rockoff (2014): "Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates," *American Economic Review*, 104(9), 2593–2632.

Duflo, E., P. Dupas, and M. Kremer (2008): "Peer Effects and the Impact of Tracking: Evidence from a Randomized Evaluation in Kenya," .

Gibbons, S., and S. Telhaj (2008): "Peers and Achievement in England's Secondary Schools," http://www.spatialeconomics.ac.uk/.

Hoxby, C., and G. Weingarth Salyer (2006): "School Reassignment and the Structure of Peer Effects, Manuscrito," *Universidad de Harvard*.

Hoxby, C. M. (2006): "Program Report: Economics of Education," Discussion paper.

Lavy, V., O. Silva, and F. Weinhardt (2009): "The Good, the Bad and the Average: Evidence on the Scale and Nature of Ability Peer Effects in Schools," .

Lyle, D. S. (2009): "The Effects of Peer Group Heterogeneity on the Production of Human Capital at West Point," *American Economic Journal: Applied Economics*, 1(4), 69–84.

Murphy, R., and F. Weinhardt (2020): "Top of the Class: The Importance of Ordinal Rank," *The Review of Economic Studies*, 87(6), 2777–2826.

Rosenbaum, P. R., and D. B. Rubin (1983): "The Central Role of the Propensity Score in Observational Studies for Causal Effects," *Biometrika*, 70(1), 41–55.

Sacerdote, B. (2001): "Peer Effects with Random Assignment: Results for Dartmouth Roommates*," *The Quarterly Journal of Economics*, 116(2), 681–704.

Woodcock, S., J. Friesen, and R. Meilman Lomaz Cohn (2025): "Sorting, Peer Effects and School Effectiveness in Private and Public Schools," .

Zimmerman, D. J. (2003): "Peer Effects in Academic Outcomes: Evidence from a Natural Experiment," *The Review of Economics and Statistics*, 85(1), 9–23.

A.1 Interpretation of $\beta(r)$

The value $\Delta f \times \Delta r \times [\beta(r') - \beta(r)]$ encapsulates the approximate effect of relocating the area $\Delta f \times \Delta r$ in the density $f(\cdot)$ from r to r'. We show this using the Riemann sum approximation for f:

$$\int \beta(r) f(r) dr \approx \Delta r \sum_{i=1}^{n} \beta(r_i) f(r_i)$$

Now, with the modified density g:

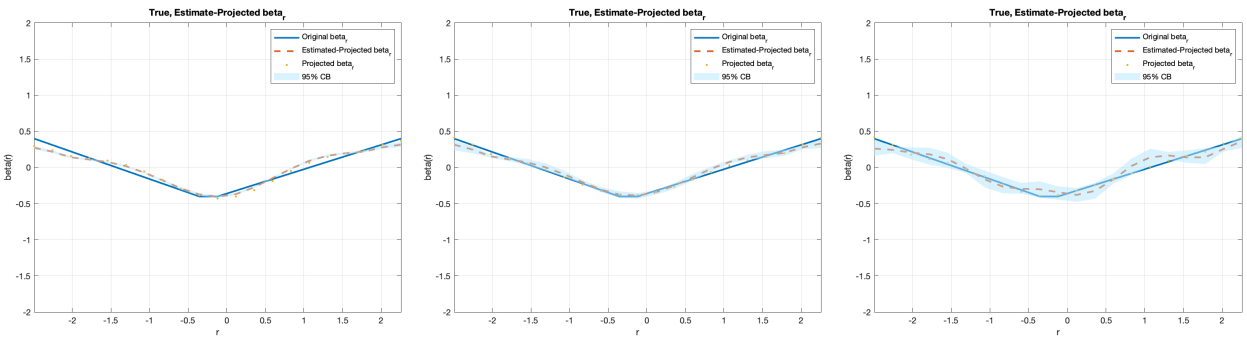
$$\begin{aligned} g(r_k) &= f(r_k) + \Delta f \\ g(r_j) &= f(r_j) - \Delta f \\ g(r_i) &= f(r_i) \qquad \text{for all } i \neq j, k \end{aligned}$$

The difference in the Riemann sums for f and g is:

$$\Delta r \left[\beta(r_k) g(r_k) + \beta(r_j) g(r_j) - \beta(r_k) f(r_k) - \beta(r_j) f(r_j)\right] = \Delta f \times \Delta r \times \left[\beta(r_k) - \beta(r_j)\right].$$

A.2 Appendix Figures and Tables

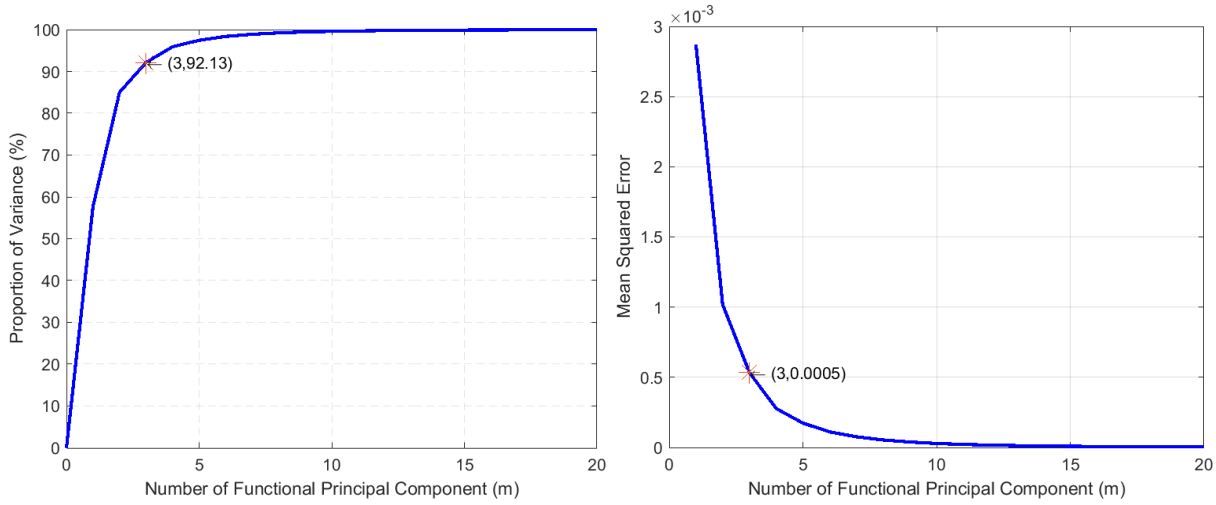
Figure A.1: True, Estimated, Projected $\beta(.)$ with m of Functional Principal Components



(a) $m = 3$ (b) $m = 5$ (c) $m = 7$

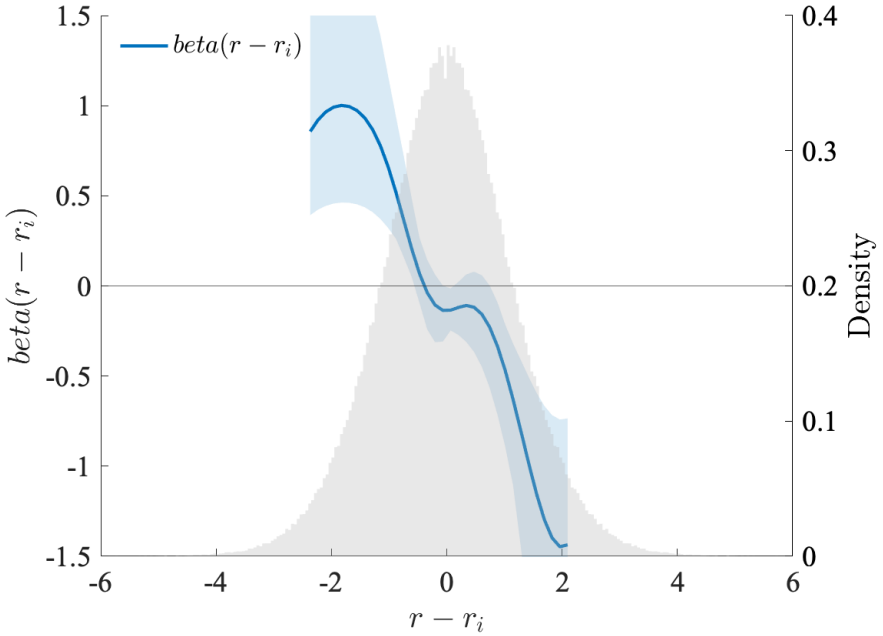
Note: Each figure shows simulated $\beta(\cdot)$, projection of $\beta(\cdot)$ onto the functional principal components of $(f_i)_i$, and the estimate of $\beta(\cdot)$ according to our Monte Carlo analyses. m denotes the number of leading principal components used.

Figure A.2: Cross-Validation Errors and Cumulative Screeplot



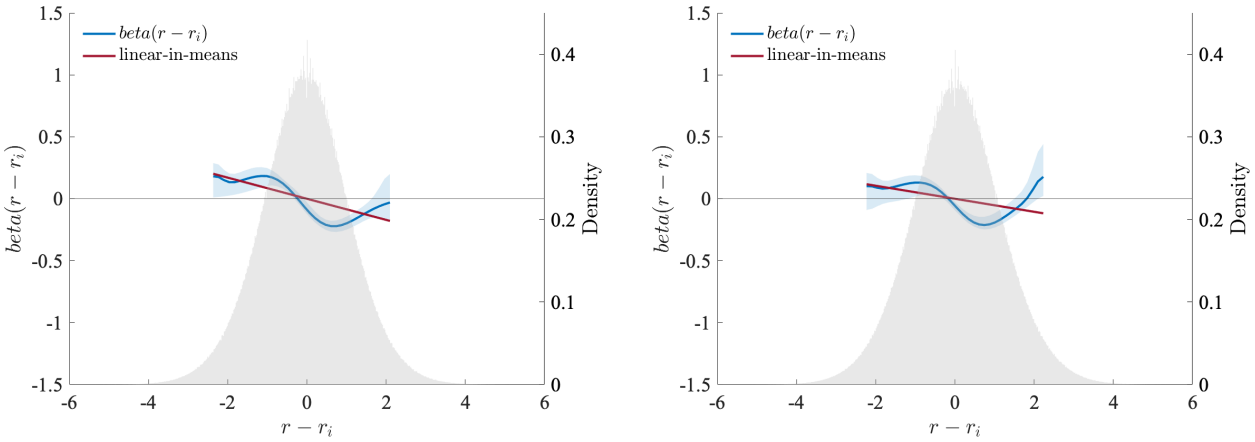
Note: The left panel is the cumulative scree plot, which presents the total percentage of variation explained by the leading m-functional principal components also as a function of m. In the right panel is the leave-one-out cross-validation (LOOCV) errors as a function of the number m of functional principal components used to approximate parental income trajectory.

Figure A.3: Non-linear Peer Effect from 2SLS



Notes: The figure plots the estimated peer effect function $\hat{\beta}(\tau - \tau_i)$ from 2SLS. The horizontal axis represents the difference between peers' baseline test scores and the student's own baseline score $(\tau - \tau_i)$. The blue curve shows the non-parametric estimate based on functional principal components with 95% confidence intervals. The shaded gray area shows the empirical density of $(\tau - \tau_i)$ in the data. Standard errors are computed using a wild bootstrap clustered at the school-market level.

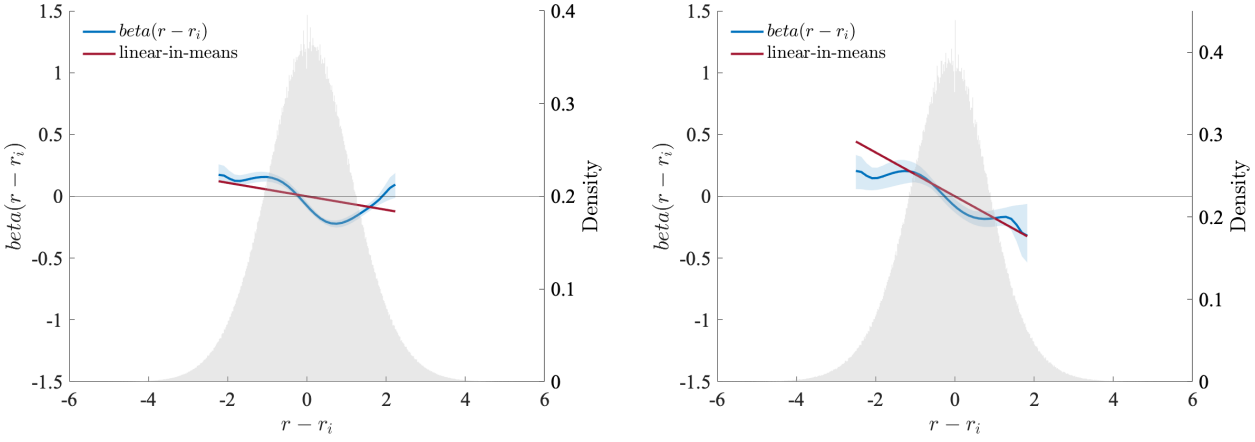
Figure A.4: By Gender



(a) Female (b) Male

Note: Peer effects functions $\beta(r - r_i)$ estimated by gender. The blue curves show the non-parametric estimates based on functional principal components with 95% confidence intervals, while the red lines represent the linear-in-means specifications. Gray areas represent the empirical density of $r - r_i$ within each subgroup.

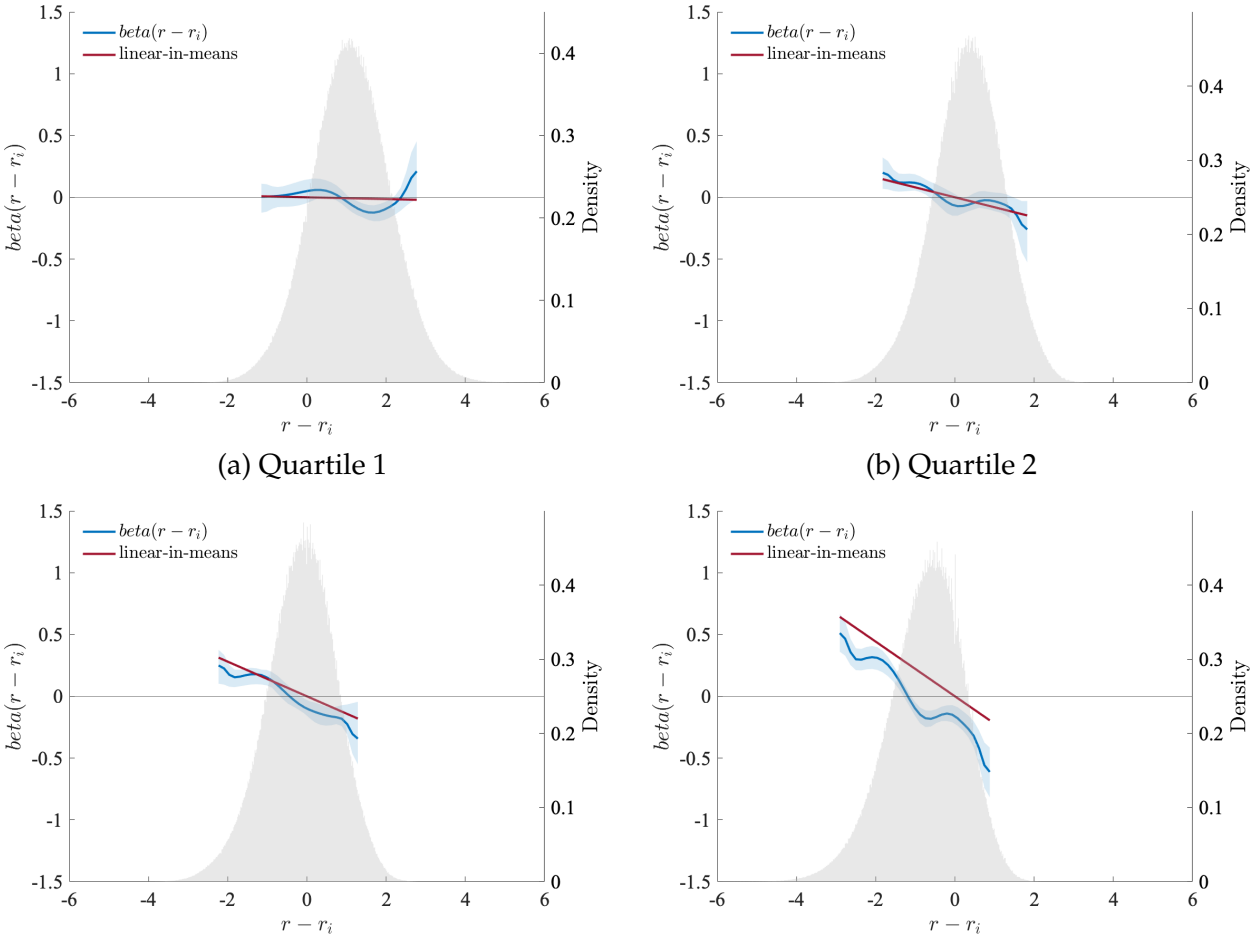
Figure A.5: By Socioeconomic Status



(a) Free/Reduced-Price Lunch (Low SES) (b) Paid Lunch (High SES)

Note: Peer effects functions $\beta(r - r_i)$ estimated by socioeconomic status. The blue curves show the non-parametric estimates based on functional principal components with 95% confidence intervals, while the red lines represent the linear-in-means specifications. Gray areas represent the empirical density of $r - r_i$ within each subgroup.

Figure A.6: By Baseline Score Quartiles



(a) Quartile 1

(b) Quartile 2

(c) Quartile 3

(d) Quartile 4

Note: Peer effects functions $\beta(r - r_i)$ estimated by performance quartiles. The blue curves show the non-parametric estimates based on functional principal components with 95% confidence intervals, while the red lines represent the linear-in-means specifications. Gray areas represent the empirical density of $r - r_i$ within each subgroup.