

This PDF is a selection from a published volume from the
National Bureau of Economic Research

Volume Title: The Economics of School Choice

Volume Author/Editor: Caroline M. Hoxby, editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-35533-0

Volume URL: <http://www.nber.org/books/hox03-1>

Conference Date: February 22-24, 2001

Publication Date: January 2003

Title: Does Public School Competition Affect Teacher Quality?

Author: Eric A. Hanushek, Steven G. Rivkin

URL: <http://www.nber.org/chapters/c10084>

Does Public School Competition Affect Teacher Quality?

Eric A. Hanushek and Steven G. Rivkin

Vouchers, charter schools, and other forms of choice have been promoted as a way to improve public schooling, but the justification for that position is largely based on theoretical ideas. Until quite recently there was little evidence on public school responsiveness to competition from private schools, other public school districts, or charter schools, and empirical research remains quite thin. Under most conceivable scenarios of expanded choice, even with private school vouchers, the public school system will still remain the primary supplier of schooling. Therefore, it is important to know what might happen to quality and outcomes in the remaining public schools. This research is designed to provide insights about that from an analysis of how public schools respond to competition from other public schools.

The empirical analysis has two major components. First, estimates of average school quality differences in metropolitan areas across Texas are compared to the amount of public school competition in each. At least for the largest metropolitan areas, the degree of competition is positively related to performance of the public schools. Second, the narrower impact of metropolitan area competition on teacher quality is investigated. Because teacher quality has been identified as one of the most important determinants of student outcomes, it is logical to believe that the effects of competition on

Eric A. Hanushek is the Paul and Jean Hanna Senior Fellow at the Hoover Institution of Stanford University and a research associate of the National Bureau of Economic Research. Steven G. Rivkin is professor of economics at Amherst College and a research associate of the National Bureau of Economic Research.

This research has been supported by grants from the Smith Richardson Foundation and the Packard Humanities Institute. The authors would like to thank Joe Altonji, Patrick Bayer, Caroline Hoxby, and participants at the NBER Conference on the Economics of School Choice for their many helpful comments.

hiring, retention, monitoring, and other personnel practices would be one of the most important aspects of any force toward improving public school quality. The results, although far from conclusive, suggest that competition raises teacher quality and improves the overall quality of education.

Prior to the analysis of Texas public schools we briefly consider the various margins of competition for public schools. Although many simply assume that expanded availability of alternatives will lead to higher public school quality, the institutional structure of public schools raises some questions about the strength of any response.

1.1 The Margins of Competition

Competition for public schools may emanate from a variety of sources. Neighborhood selection places families in particular public school districts and specific school catchment areas within districts. Families also choose whether to opt out of the public schools and send their children to parochial or other private school alternatives.¹ Although these choices have operated for a long time, recent policy innovations have expanded competition within the public school sector. The ability to attend school in neighboring districts, charter schools, and private schools with public funding enhances choice and potentially imposes additional competitive pressures on public schools.

Most of the attention to private schools has concentrated on student performance in Catholic schools.² The literature on Catholic school performance is summarized in Neal (1998) and Grogger and Neal (2000). The evidence has generally indicated that Catholic schools on average outperform public schools.³ This superiority seems clearest in urban settings, where disadvantaged students face fewer options than others.

Our main interest, however, centers on the reactions of public schools to the private sector. In an important article about the impact of private schools on schools in the public sector, Hoxby (1994) demonstrates that

1. Magnet schools have also existed for a long time. However, their small numbers, targeted curricula, and frequent use of entrance examinations limit the extent to which they provide competition for other public schools. Moreover, because they are often introduced to meet school desegregation objectives, choice is frequently limited by racial quotas (Armor 1995).

2. Currently, almost 90 percent of all students attend public elementary and secondary schools. This percentage has been stable for some time, although the exact character of the alternative private schooling has changed. The percentage of private school students in Catholic schools has declined, whereas other religious based schooling has increased to offset this decline. Nonetheless, adequate data on non-Catholic schools have not been readily available.

3. As has been recognized since some of the earliest work on the topic (Coleman, Hoffer, and Kilgore 1982), it is difficult to separate performance of the private schools from pure selection phenomena. A variety of alternative approaches have dealt with the selection problem, and a rough summary of the results of those efforts is that there remains a small advantage from attending Catholic schools. Grogger and Neal (2000) suggest, however, that there is no advantage to attending private elite schools—a surprising result given the high average tuitions.

public schools in areas that have larger concentrations of Catholic schools perform better than those facing less private competition. This analysis provides the first consistent evidence suggesting that public schools react to outside competition.

The most important element of competition comes from other public schools. Specifically, households can choose the specific jurisdiction and school district, à la Tiebout (1956), by their choice of residential location. Although adjustment is costly, these choices permit individuals to seek high-quality schools if they wish. Residential location decisions are of course complicated, involving job locations, availability of various kinds of housing, school costs and quality, and availability of other governmental services. Nonetheless, given choice opportunities plus voting responses, this model suggests pressure on schools and districts to alter their behavior; competitive alternatives that lead families to choose other schools would yield downward pressure on housing prices and perhaps even an enrollment decline.

The ensuing public pressures might be expected to lead administrators and teachers to respond. For example, job performance may affect a superintendent's ability to move to another district or a principal's autonomy or ability to remain in a school. Better performance by teachers may make the school more attractive to other high-quality teachers, thereby improving working conditions.

Offsetting forces may, nonetheless, mute any competitive pressures. The current structure of many school systems including tenure for teachers and administrators likely lessens the impact of competitive forces. Institutionally, district survival is virtually guaranteed under plausible changes in the competitive environment.

The empirical analysis of Borland and Howsen (1992) and its extension and refinement in Hoxby (2000) investigate public school responses to Tiebout forces using the concentration of students in school districts within metropolitan areas as a measure of competition. Borland and Howsen find that metropolitan areas with less public school competition have lower school quality. Noting, however, that the existing distribution of families across districts reflects endogenous reactions to school quality, Hoxby pursues alternative strategies to identify the causal impact of concentrations. She finds that consideration of endogeneity increases the estimated impact of competition on the performance of schools. Our analysis builds on these specifications of public school competition.

The general consideration of Tiebout competition, however, leaves many questions open. For example, it is not obvious how to define the "competitive market." Although the district is the fundamental operating and decision-making unit in most states, districts themselves can be very large and heterogeneous. This heterogeneity could lead to competition, and responses, that are more local in nature—say, at the school rather than the

district level. For example, Black (1999) and Weimer and Wolkoff (2001) suggest that school quality differences are capitalized into housing prices at the individual school rather than the district level. This ambiguity motivates our use of alternative measures of the level of competition.

Much recent attention has focused on more radical forms of competition such as vouchers or charter schools. Again, whereas most debate focuses on the performance of these alternatives, our interest is the reaction of public schools to these competitive alternatives. With the exception of Hoxby (chap. 8 in this volume), however, little consideration has been given to the actions of public schools.

1.2 The Importance of Teacher Quality

The difficulty of identifying and measuring school quality constitutes a serious obstacle to learning more about the effects of competition. A substantial body of work on the determinants of student achievement has failed to yield any simple descriptions of the key school and teacher factors. Although class size and other variables may significantly affect outcomes for specific populations and grades, financial measures (spending per pupil and teacher salaries) and real resources (teacher experience and degrees, class size, facilities, and administration) do not appear to capture much of the overall variation in school or teacher quality (Hanushek 1986, 1997).⁴

On the other hand, schools and teachers have been shown to be dramatically different in their effects on students. A variety of researchers have looked at variations among teachers in a fixed effect framework and have found large differences in teacher performance (see, e.g., Hanushek 1971, 1992; Murnane 1975; Armor et al. 1976; Murnane and Phillips 1981). The general approach has been to estimate value added achievement models and to assess whether or not performance gains differ systematically across teachers. It is important to note that value added models control for differences in entering achievement and thus remove a number of potential sources of bias, including differences in past performance and school factors, individual ability, and so forth. In every instance of such estimation, large differences have been found. Of particular significance for the work here, these differences have generally been weakly related to the common measures of teachers and classrooms found in the more traditional econometric estimation.

These analyses have not, however, conclusively identified the impacts of different teachers. Because parents frequently set out to choose not just spe-

4. Although parts of this discussion have generated controversy—largely over the policy conclusions that might be drawn—none of the discussion has suggested that any of these resource measures are good indicators of overall school quality. The focus in the discussion has been whether policy changes in any of these measures could be expected to yield positive effects on student performance. See, for example, the paper by Burtless (1996).

cific schools but also specific teachers within schools, the makeup of individual classrooms may not be random. This possibility is compounded by two other influences. First, teachers and principals also enter into a selection process that matches individual teachers with groupings of children.⁵ Second, if the composition of the other children in the classroom is important—that is, if there are important peer group effects on achievement—the gains in an individual classroom will partially reflect the characteristics of the children and not just the teacher assigned to the classroom.⁶ These considerations suggest a possibility that classroom outcome differences reflect more than just variations in teacher quality.

A recent paper by Rivkin, Hanushek, and Kain (2001) uses matched panel data for individual students and schools to estimate differences in teacher quality that are not contaminated by other factors. Because that work forms the basis for the investigation here, it is useful to understand the exact nature of it. The authors use a value added model that compares the pattern of school average gains in achievement for three successive cohorts as they progress through grades five and six. The value added model, by conditioning on prior achievement, eliminates unmeasured family and school factors that affect the level of beginning achievement for a grade and permits concentration on just the flow of educational inputs over the specific grade. The analysis then introduces fixed effects for individual schools and for specific grades in each school, allowing for effects of stable student ability and background differences, of overall quality of schools, and of the effectiveness of continuing curricular and programmatic elements for individual grades. This basic modeling provides what is essentially a prediction of achievement growth for individuals based on each one's past performance and specific schooling circumstances. The central consideration, then, is how much changes in teachers affect the observed patterns of student achievement growth within each school.⁷

This analysis shows that cohort differences in school average gains rise significantly as teacher turnover increases. By controlling for other potentially confounding influences, the methodology generates a lower bound estimate of the variance in teacher quality based on within-school differences in test score gains among the cohorts.

The estimation of “pure” teacher quality differences reveals that the variation in teacher quality within schools (i.e., ignoring all variation across

5. Hanushek, Kain, and Rivkin (2001b) show that teachers both leave schools and select new districts based on the achievement and racial composition of the schools.

6. Hanushek et al. (forthcoming) and Hanushek, Kain, and Rivkin (2002) show that peer characteristics related to achievement and race influence individual student achievement.

7. The analysis does not look at the achievement growth in individual classrooms, because that could confound individual student placement in specific classrooms with differences in teacher skills. Therefore, it aggregates students across classrooms in a specific grade, effectively instrumenting with grade to avoid selection effects. Even if we had wished to pursue individual classrooms, however, we could not because of data limitations.

schools) is large in Texas elementary schools. One standard deviation of teacher quality—for example, moving from the median to the 84th percentile of the teacher quality distribution—increases the annual growth of student achievement by at least 0.11 standard deviations, and probably by substantially more. This magnitude implies, for example, that having such an 84th percentile teacher for five years in a row rather than a 50th percentile teacher would be sufficient to eliminate the average performance gap between poor students (those eligible for free or reduced lunch) and nonpoor students.

Evidence on the importance of teacher quality forms the basis for a major segment of the empirical analysis here. Specifically, if the degree of local competition is important, it should be possible to detect its impact on teacher quality by examining performance variation along with the amount of local competition across the state of Texas. In particular, it would be surprising for competition to exert a substantial effect on students without influencing the quality of teaching, and investigating these effects provides information about the mechanism behind any observed impacts of competition.

1.3 Empirical Analysis

We investigate how varying amounts of public school competition in the classic Tiebout sense affect student performance and the hiring of teachers. It is important to note that our efforts are not general. We leave aside many of the issues discussed previously and in the other papers of this volume about possible details and dimensions of competition and concentrate entirely on issues of academic performance across broadly competitive areas. Nonetheless, the importance of this topic for individual labor market outcomes and for the politics of schools justifies the choice.

The empirical work exploits the rich data set on student performance of the University of Texas at Dallas Texas Schools Project. Because Texas is a large and varied state, a wide range of local circumstances is presented. Indeed, there are twenty-seven separate metropolitan statistical areas (MSAs) in Texas. These areas, described in table 1.1, vary considerably in size and ability to mount effective competition across districts. The basic Tiebout model assumes a wide variety of jurisdictional choices such that people can choose among alternative public service provision while retaining flexibility in housing quality and commuting choices. Clearly, the smaller MSAs of Texas offer limited effective choice in all dimensions, so it will be interesting to contrast results across the various areas of the state.

We employ the Texas Schools Project data first to estimate overall quality differences between MSAs and to compare these results with the degree of public school competition. Following that, we investigate whether or not competition raises the quality of teaching.

As suggested by the previous discussions, this analysis is best thought of as a reduced-form investigation. We do not observe the underlying decision-

Table 1.1 Metropolitan Statistical Areas (MSAs) in Texas

Metropolitan Area	1997 Population (1000s)	Density (persons per square mile)	Population Change 1990–97 (%)	Number of Districts	Number of Elementary Schools
Houston	3,852	651	15.9	45	699
Dallas	3,127	505	16.8	77	590
Fort Worth-Arlington	1,556	533	14.4	37	311
San Antonio	1,511	454	14.1	25	318
Austin-San Marcos	1,071	253	26.6	29	210
El Paso	702	693	18.6	9	137
McAllen-Edinburg- Mission	511	326	33.2	15	121
Corpus Christi	387	253	10.6	20	96
Beaumont-Port Arthur	375	174	3.8	16	72
Brownsville-Harlingen- San Benito	321	354	23.3	10	81
Killeen-Temple	300	142	17.4	14	63
Galveston-Texas City	243	609	11.8	9	64
Odessa-Midland	243	135	7.9	2	28
Lubbock	231	256	3.6	8	57
Brazoria	225	163	17.6	8	48
Amarillo	208	114	11.0	5	59
Longview-Marshall	208	118	7.5	20	44
Waco	203	195	7.3	18	58
Laredo	183	55	37.5	4	53
Tyler	167	180	10.2	8	40
Wichita Falls	137	89	5.2	9	35
Bryan-College Station	133	227	9.1	2	18
Texarkana	123	82	2.7	13	24
Abilene	121	133	1.5	5	39
San Angelo	103	67	4.3	6	32
Sherman-Denison	102	109	6.9	13	30
Victoria	82	93	10.3	4	21

making by school officials; nor do we have detailed and precise measures of the competition facing individual schools and districts. Instead we use aggregate indicators of potential competition from public schools and concentrate on whether or not there are systematic patterns to student outcomes.

1.3.1 The Texas Database

The data used in this paper come from the data development activity of the UTD Texas Schools Project.⁸ Its extensive data on student performance are compiled for all public school students in Texas, allowing us to use the universe of students in the analyses. We use fourth-, fifth-, and sixth-grade

8. The UTD Texas Schools Project has been developed and directed by John Kain. Working with the Texas Education Agency (TEA), this project has combined a number of different data

data for three cohorts of students: fourth-graders in 1993, 1994, and 1995. Each cohort contributes two years of test score gains. Students who switch public schools within the state of Texas can be followed just as those who remain in the same school or district, a characteristic we use in our analysis.

The Texas Assessment of Academic Skills (TAAS), which is administered each spring, is a criterion-referenced test used to evaluate student mastery of grade-specific subject matter. We focus on test results for mathematics, the subject most closely linked with future labor market outcomes. We transform all test results into standardized scores with a mean of zero and variance equal to one. The bottom 1 percent of test scores and the top and bottom 1 percent of test score gains are trimmed from the sample in order to reduce measurement error. Participants in bilingual or special education programs are also excluded from the sample because of the difficulty in measuring school and teacher characteristics for these students.

The empirical analysis considers only students attending public school in one of the twenty-seven MSAs in Texas (identified in table 1.1). A substantial majority of all Texas public school students attend schools in one of these MSAs. Each MSA is defined as a separate education market, and measures of competition are constructed for each. The analysis is restricted to MSAs because of the difficulty of defining school markets for rural communities. Below we discuss potential problems associated with defining education markets in this way.

1.3.2 Competition and School Quality

How will public school competition affect the provision of education? Although Tiebout-type forces would be expected to raise the efficiency of schooling, it is not clear that more competition will necessarily result in higher school quality. If wealth differences or other factors related to school financing lead to more resources in areas with less competition, the efficiency effects of competition could be offset by resource differences. Therefore we consider differences in both school quality and school efficiency across metropolitan areas.

A second important issue is precisely how to define the relevant competition. The importance of district administrators in allocating funds, determining curriculum, hiring teachers, and making a variety of other decisions suggests that much if not most of the effects of competition should operate at the district level. However, anecdotal evidence on school choice provides strong support for the notion that parents actively choose among schools

sources to compile an extensive data set on schools, teachers, and students. Demographic information on students and teachers is taken from the Public Education Information Management System (PEIMS), which is TEA's statewide educational database. Test score results are stored in a separate database maintained by TEA and must be merged with the student data on the basis of unique student identifiers. Further descriptions of the database can be found in Rivkin, Hanushek, and Kain (2001).

within urban and large suburban school districts, consistent with the view that principals and teachers exert substantial influence on the quality of education. This anecdotal information is reinforced by the aforementioned research on housing capitalization.

We treat the basis for competition as an empirical question. In the estimation, we conduct parallel analyses where competition is measured on the basis of the concentration of students both in schools and separately in districts.

Although Hoxby (2000) provides the empirical context within which to place this study of school efficiency, the methodology employed here is much closer to the work by Abowd, Kramarz, and Margolis (1999) on inter-industry wage differences. Just as interindustry wage differences reflect both worker heterogeneity and industry factors, interschool or district differences in student performance reflect both student heterogeneity and school factors. However, a comparison of wage differences for a worker who switches industries or of achievement differences for a student who switches schools effectively eliminates problems introduced by the heterogeneity of workers or students. In this way the availability of matched panel data facilitates the identification of sector effects.

Equation (1) describes a value added model of learning for student i in grade g in MSA m at time t :

$$(1) \quad \Delta \text{Achievement}_{igt} = \text{family}_i + \text{family}_{igt} + \text{MSA}_m + \text{error}_{igt},$$

where the change in achievement in grade g equals test score in grade g minus test score in grade $g - 1$.

The overall strategy concentrates on estimation of metropolitan area fixed effects (MSA_m) for each of the twenty-seven MSAs in Texas. Importantly, this model removes all fixed family, individual, and other influences on learning (family_i) as well as time-varying changes (family_{igt}) in family income, community type (urban or suburban), specific year effects, and the effect of moving prior to the school year (students may or may not move prior to fifth grade).⁹

In this model of student fixed effects in achievement gains, the MSA quality fixed effects are identified by students who switch metropolitan areas. These twenty-seven MSA fixed effects provide an index of average school quality for the set of metropolitan areas. Although most of the variation in school quality likely occurs within an MSA, such variation is ignored because of the focus on competition differences among MSAs.¹⁰ Importantly,

9. The empirical specification with the emphasis on the effects of moving reflects our prior analysis that shows an average decline in learning growth in the year of a move (Hanushek, Kain, and Rivkin 2001a).

10. The choice of individual schools does introduce one complication. The average effects will be weighted by the student choices of schools within metropolitan areas instead of the overall distribution of students across an area. Our econometric estimates assume that any differential selectivity of schools within metropolitan areas (after allowing for individual fixed effects for movers) is uncorrelated with the level of competition.

by removing student fixed effects in achievement gains, this approach effectively eliminates much of the confounding influences of student heterogeneity present in analyses based on cross-sectional data.

Nevertheless, we do not believe that students switch districts at random, and changes in circumstances not captured by the student fixed effects may dictate the characteristics of the destination school as well as affecting student performance. For example, families who experience job loss or divorce may relocate to inferior districts, whereas families who experience economic improvements may tend to relocate to better districts.¹¹ If the limited number of time-varying covariates does not account for such changes in family circumstances, the estimates of metropolitan area school quality will reflect both true quality differences and differences in family circumstances.

However, even if the rankings of metropolitan area average quality are contaminated, regressions of these rankings on the degree of competition may still provide consistent estimates of competition effects as long as the omitted student and family effects are not related to the degree of competition. The fact that mobility across regions is most importantly linked to job relocations and less to seeking specific schools or other amenities certainly mitigates any problems resulting from nonrandom mobility (Hanushek, Kain, and Rivkin 2001a). On the other hand, other factors, including school resources, that may be correlated with the measures of competition may confound the estimated effects of competition on school quality and, more importantly, on school efficiency. We do include average class size as a proxy for school efficiency. Although average class size captures at least a portion of any difference in resources, there is a good chance that influences of confounding factors remain.¹²

Two other important issues more specific to the study of school competition are the measurement of the degree of competition and the identification of separate public school markets. Following general analyses of market structure, we calculate a Herfindahl index based alternatively on the concentration of students by district and by school across the metropolitan areas.¹³ As Hoxby (2000) points out, the Herfindahl index is itself endogenously determined by the location decisions of families. Any movement of families into better districts within a metropolitan area will change the

11. Note that the direction of any bias is ambiguous. Negative or positive shocks that precipitate a move may affect performance prior to the move as much as or even more than performance following a move. The range of responses and the effects over time are discussed in Hanushek, Kain, and Rivkin (2001a).

12. Inter-metropolitan area differences in the price of education quality raise serious doubts about the validity of expenditure variables as measures of real differences in resources. Such differences result from cost-of-living differences, variability in working conditions, and differences in alternative employment opportunities for teachers, as well as other factors.

13. The Herfindahl index is the sum of squared proportions of students (by district or school) for the MSA. A value of one indicates all students in a single location (no competition), whereas values approaching zero show no concentration and thus extensive competition.

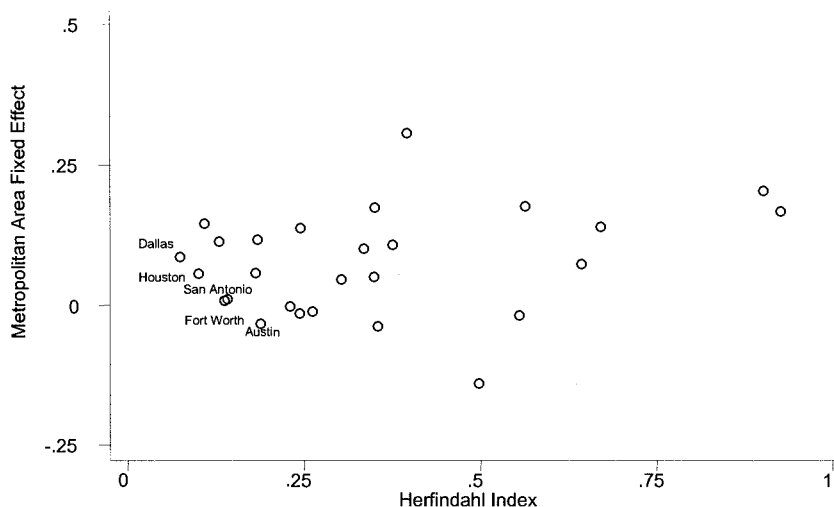


Fig. 1.1 School quality and district concentration

value of the Herfindahl index, raising it if families concentrate in larger districts and lowering it if families move to smaller districts, as would be the case with urban flight. In essence, the Herfindahl index reflects both the initial administrative structure of schools and districts as well as within-metropolitan area variation in school or district quality. Only the former provides a good source of variation, and that is the source of variation Hoxby attempts to isolate with her instrumental variable approach, which deals with the endogeneity of school districts. We do not have available instruments, so the second source of variation may introduce bias of an indeterminate direction.

The identification of the relevant education market (i.e., defining the appropriate set of schools from which parents choose) also presents a difficult task. It is certainly the case that a number of families who work in an MSA choose to live outside the MSA, and thus measuring school competition using the census definitions of MSAs almost certainly introduces some measurement error in the calculation of the Herfindahl index that would tend to bias downward the estimated effects of competition.

Results

Figures 1.1 and 1.2 plot the metropolitan area school quality fixed effects against the Herfindahl index, the measure of competition.¹⁴ The estimates of school quality are obtained from student fixed effect regressions of

14. Average enrollments in fifth and sixth grade for the three years of data are used to construct the Herfindahl index. The district (school) Herfindahl index is the sum of squared proportions of enrollment in each district (school).




Fig. 1.2 School quality and school concentration

achievement gain on twenty-seven metropolitan area dummy variables and controls for free lunch eligibility, community type, and whether the student moved prior to the grade. The five largest metropolitan areas are specifically identified in the figures. Figure 1.1 measures competition by the concentration of students into school districts, whereas figure 1.2 measures competition by the concentration of students into schools, implicitly permitting competition to occur both within and across districts. The overall patterns presented in figures 1.1 and 1.2 do not reveal a strong positive relationship between competition at either the district or school level and school quality. Rather, the scatter of points moves roughly along a horizontal line regardless of whether competition is measured at the school or district level. Not surprisingly, the coefficient on Herfindahl index from a regression of the metropolitan area fixed effect on the Herfindahl index is small and not significantly different from zero regardless of whether competition is measured at the school or district level (table 1.2).¹⁵ Note that competition varies far less when measured at the school level, because any dominance of large districts is ignored.

In contrast to the lack of an overall positive relationship between competition and school quality, the school fixed effects for the five largest metropolitan areas suggest the presence of a positive relationship between school quality and competition: the ordering of Dallas, Houston, San Antonio, Fort Worth, and Austin according to school quality exactly matches

15. All regressions are weighted by the number of students in each metropolitan area in the first-stage regressions.

Table 1.2 Relationship between MSA Average School Quality and Competition, by MSA Size

	Competition between Districts		Competition between Schools	
	All Areas	By Area Size	All Areas	By Area Size
Herfindahl index	0.11 (1.43)		0.82 (1.24)	
Herfindahl index · large MSA ^a		-1.07 (2.76)		-17.3 (-2.53)
Herfindahl index · (1 – large MSA)		0.09 (1.01)		0.16 (0.18)
Large MSA		0.11 (1.82)		0.03 (0.69)
R^2	0.08	0.36	0.06	0.31

Notes: All regressions use the estimated average quality of the MSA schools from equation (1). Observations are weighted by students in the MSAs. *T*-statistics are presented below each coefficient. For competition between districts, Herfindahl index is defined by proportionate shares of students across districts. For competition between schools, Herfindahl index is defined by proportionate shares of students across schools.

^aInteraction of Herfindahl index and dummy variable for the five largest MSAs (Houston, Dallas, Ft. Worth, San Antonio, and Austin).

the ordering by competition regardless of how competition is measured. This is confirmed by the regression results in table 1.2 that allow for separate slope coefficients for the five largest MSAs. Although there is little or no evidence that competition at the school or district level is significantly related to school quality for the smaller MSAs, the competition effect is positive and significant at the 1 percent level for the five largest metropolitan areas. Because some of the smaller MSAs in Texas actually get quite small and offer far fewer choices of districts (see table 1.1), it would not be surprising if the incentive effects of competition were much weaker in comparison to the effects in the large MSAs. Effective competition may require a minimum range of housing and public service quality (Tiebout 1956).

Regardless of MSA size, however, competition should have its sharpest effects on reducing inefficiencies in resource use and education production. In a coarse effort to isolate competition effects on efficiency, the first-stage regressions underlying figures 1.3 and 1.4 include average class size as a control for resource differences. Not surprisingly, given the strong evidence that class size and other resource differences explain little of the total variation in school quality, the inclusion of class size has little impact on either the observed patterns in the figures or the Herfindahl index coefficients (see table 1.3).

All in all, the figures and regression results suggest that competition improves school quality in larger areas with substantial numbers of school and district choices. However, a sample size of twenty-seven with only five very




Fig. 1.3 School efficiency and district concentration

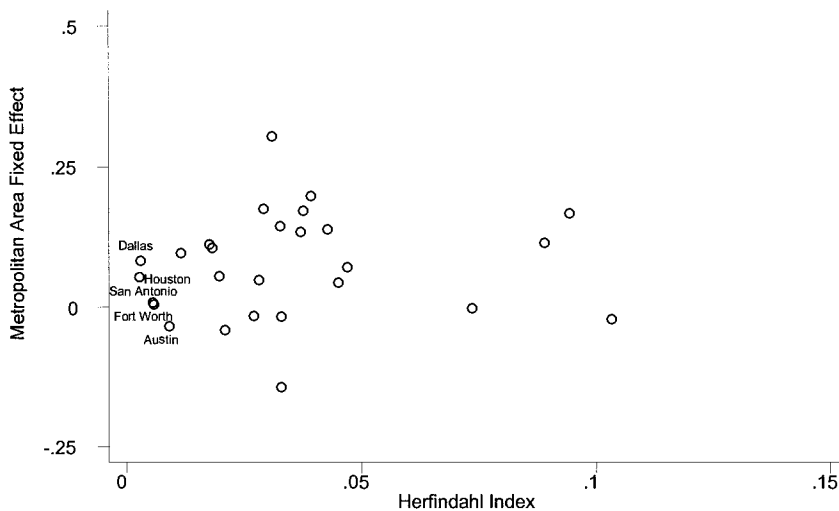


Fig. 1.4 School efficiency and school concentration

large MSAs is quite small, and there may simply not be enough variation to identify more precisely the effect of competition on average school quality and efficiency. Moreover, although the matched panel data remove many of the most obvious sources of bias, the limited number of time-varying characteristics may fail to control for all confounding family and student influences.

Table 1.3 Relationship between MSA Average School Efficiency and Competition, by MSA Size

	Competition between Districts		Competition between Schools	
	All Areas	By Area Size	All Areas	By Area Size
Herfindahl index	0.11 (1.44)		0.83 (1.26)	
Herfindahl index · large MSA ^a		-1.07 (2.27)		-17.2 (-2.51)
Herfindahl index · (1 – large MSA)		0.09 (1.01)		0.18 (0.19)
Large MSA		0.11 (1.80)		0.03 (0.69)
R ²	0.08	0.36	0.06	0.31

Notes: See notes to table 1.2. School quality is adjusted for resources (class size).

^aInteraction of Herfindahl index and dummy variable for the five largest MSAs (Houston, Dallas, Ft. Worth, San Antonio, and Austin).

1.3.3 Competition and Teacher Quality

The suggestive results for the effects of competition on overall quality leave uncertainty about the strength of the Tiebout forces. This portion of the empirical analysis investigates a much narrower question with a methodology that likely does a far better job of controlling for confounding influences on student outcomes. Although the quality of teaching is only one of many possible determinants of school quality, evidence in Rivkin, Hanushek, and Kain (2001) strongly suggests that it is the most important factor. Consequently, it would be highly unlikely that competition would exert a strong effect on school quality without affecting the quality of teachers.

At first glance the problem might appear to be quite simple: More competitive areas should lead schools to hire better teachers as measured by teacher education, experience, test scores, and other observable characteristics. However, two issues complicate any simple analysis: (a) evidence overwhelmingly shows that observable characteristics explain little of the variation in teacher quality in terms of student performance (see Hanushek 1986, 1997); and (b) competition could lead schools to raise teacher quality per dollar spent but not the level of quality, and it is quite difficult to account for cross-sectional differences in the price of teacher quality.

Isolating the contributions of teachers to between-school or between-district differences in student performance is inherently very difficult. Given the added difficult task of accurately capturing cross-sectional differences in the price of teacher quality, we do not believe that an analysis of the effect of competition on teacher quality per dollar in salary is likely to produce

compelling evidence. We pursue a very different empirical approach focusing on the within-school and -district variations in the quality of teaching, using the methodology developed in Rivkin, Hanushek, and Kain (2001).

In essence, our approach here examines the link between competition and the variance in teacher quality, testing the hypothesis that more competition should lead to less variance in the quality of teaching within schools and districts. Lower variance would result if competition pushes schools to hire the most qualified applicants and to be more aggressive in pushing teachers to perform better and in dismissing teachers who do not teach well. Schools not facing much competition would be free to pursue other considerations in hiring and to avoid potentially unpleasant retention decisions and serious monitoring.

Consider a job search framework in which firms differ in the extent to which they maximize profits (student outcomes such as achievement in the case of schools). Firms facing few competitive pressures in which management cannot capture residual profits would probably not place much emphasis on teacher quality in making personnel decisions. Rather, other aspects of teachers might play a more important role. If these other characteristics are not highly correlated with those related to instructional quality, the variance in instructional quality would be quite large. Firms facing substantial competitive pressures, on the other hand, would probably place much more emphasis on characteristics related to instructional quality. The variation in quality at each firm would be much smaller as management attempted to hire and retain the best staff possible given the level of compensation.

Such factors suggest that quality variation within schools or districts in part reflects administrator commitment to instructional quality. Notice that schools with low salaries and poor working conditions will be likely to attract lower-quality teachers than other schools despite the best attempts of administrators. It is not that variance measures quality; rather, it is the case that the variance in instructional quality should decline the stronger the commitment to such quality.

The linkage between competition and the variance in teacher quality reflects two underlying assumptions. First, improving the quality of teachers is primarily achieved by hiring and retention decisions for teachers. Second, because of the incentives, average quality will not decline while a teaching force with a more homogeneous impact on student performance is selected.

If competition leads schools to concentrate on the hiring and retention of better teachers, the prior arguments suggest that the variance in teacher quality will also decline. However, teacher quality may also be raised by increased effort or an improvement in the skills of teachers through in-service training. Such skill or effort improvements have a more ambiguous effect on the within-school or -district variance in teacher quality, because the effect depends on where attention is focused or is most effective. Holding constant

the distribution of teachers in a school, increases in the effort or skill (through training) of teachers initially at the lower end of the skill distribution would unambiguously reduce variance.¹⁶ Moreover, as long as effort and skill are not negatively correlated, policies aimed at increasing effort at the bottom of the effort distribution would generally reduce the variance in teacher quality. On the other hand, if policies to increase effort or to provide effective in-service training were to exert larger impacts at the upper end of the initial skill distribution, the variance could be raised. We believe, based on past research of the inefficacy of teacher development programs and additional education, that there is more support for the teacher selection route than for the effort or development route as a way to increase teacher quality. Nonetheless, we cannot rule out that there is some potential for a lessened effect of competition on the variance in teacher performance to the extent that the latter approaches are employed, are effective, and have a larger impact at the top of the skill distribution.

Although it is possible that schools move to more homogeneous but lower quality with competition, this movement seems very unlikely in a system that also emphasizes accountability, as the Texas system does. Common conceptual models of school behavior, even those based on nonmaximizing approaches, would generally not support a lessened variance with a lower mean from competition.

The within-school variance in teacher quality, measured in terms of the student achievement distribution, is estimated from year-to-year changes in average student test score gains in grades five and six.¹⁷ We hypothesize that teacher turnover should lead to greater variation in student performance differences among cohorts if there is less competition. This would result from the inferior personnel practices of schools facing little competition, which would increase both the variance of the quality of new hires and the variance of the quality of teachers retained following the probationary period. Of course, other factors that contribute to differences across cohorts might be systematically related to teacher turnover and competition, and we take a number of steps to control for such confounding influences.

The methodology and identifying assumptions are described in detail in Rivkin, Hanushek, and Kain (2001) and are only summarized here. Throughout, we look only at within-school variance in teacher quality and ignore any between-school variance. This approach, while giving a lower bound on teacher quality differences, avoids any possible contamination of family selection of schools. In order to sort out teacher quality effects from other things that might be changing within a school, we concentrate on the

16. This impact would be consistent with the focus of many in-service programs on remediation or improving basic teaching skills.

17. We concentrate on grade average achievement largely because the data do not permit us to link students with individual teachers, but also to avoid problems of within-grade sorting of students and teachers.

divergence of patterns of achievement gains across cohorts for each school. The idea behind the estimation is that the pattern of achievement gains across grades and cohorts of students within a school should remain constant (except for random noise) if differences among individual students are taken into account and if none of the characteristics of the school (teachers, principal, curriculum, etc.) change. When teachers change, however, variations in teacher quality will lead to a divergence of achievement patterns over time. We then relate systematic changes in teachers and in other aspects of schools to any changes in the pattern of achievement gains that are observed.

The basic framework regresses the between-cohort variance in school average test score gains on the proportion of teaching positions occupied by new people in successive years. The dependent variable generally analyzed is

$$(2) \quad [(\overline{\Delta A}_{6s}^c - \overline{\Delta A}_{5s}^c) - (\overline{\Delta A}_{6s}^{c'} - \overline{\Delta A}_{5s}^{c'})]^2.$$

Each term in this expression involves the average growth in achievement (ΔA) for a given grade (five or six) and a given cohort (c or c') in a specific school (s). This measure focuses on the pattern of achievement changes and how it differs across cohorts. The term can be interpreted as the degree that achievement patterns diverge over time: If nothing changes in the grade pattern of achievement across cohorts, this term will be zero.

Intuitively, if teacher quality differences are important, high turnover of teachers should lead to more variation in teacher quality over time; this should show up in lack of persistence of student gains across cohorts.¹⁸ To be precise, if no teachers in a school change and if the effectiveness of teachers is constant across adjacent years, teachers would add nothing to the divergence of achievement patterns across cohorts. On the other hand, if all of the teachers change and if teachers are randomly selected into schools, the divergence of achievement across cohorts would reflect the underlying variance in teacher quality. Our estimation strategy formalizes these notions and shows that there is a precise relationship between achievement changes as in equation (2) and the proportion of teachers who are different between years of observation (given the prior assumptions). In general, if the assumptions are violated, it would imply that we have underestimated the variation in achievement. In looking at the effects of competition, the most likely effect is that any effects of competition will be underestimated.¹⁹

To control for other influences on the variation in cohort performance,

18. Rivkin, Hanushek, and Kain (2001) show that the magnitude of the coefficient on teacher turnover has a simple interpretation. It equals four times the within-school variance in teacher quality.

19. The exception would be if the teacher development/effort models were dominant and if the cost-effective response of school districts is to emphasize the top of the existing distribution of teachers—something that we believe is unlikely.

the regressions also standardize for the inverse of the number of teachers in the grade,²⁰ the inverse of student enrollment, and a dummy variable for one of the cohort comparisons. We also restrict the sample to students who remain in the same school for both grades, effectively removing fixed student effects.²¹ Some regressions also remove school or even school-by-grade fixed effects, identifying effects on competition by differences in the rate of teacher turnover between the 1993 and 1994 cohorts and the 1994 and 1995 cohorts.

Finally, we aggregate across the teachers and classrooms within each grade of each school. Aggregation overcomes what is possibly the largest form of selection within schools: that which occurs when parents maneuver their children toward specific, previously identified teachers or when principals pursue purposeful classroom placement policies. Looking at overall grade differences, which is equivalent to an instrumental variable estimator based on grade rather than classroom assignment, circumvents this within-grade teacher selection.

The new contribution of this work is the introduction of competition into the analysis in the form of an interaction between the proportion of teachers who are different and the Herfindahl index for the MSA. If competition works to reduce the within-school or within-district variance in teacher quality, the coefficient on the interaction term should be positive. (Because variation over time in the Herfindahl index is not used, the main effect of the index cannot be identified in the fixed effect specifications, so it is not included.)

The fact that districts exert substantial control over teacher hiring suggests that it is the competition between districts that should have the strongest influence on teacher quality. However, there are a number of reasons to believe that the competition should be measured at the school level. First, principals exert a great deal of control over hiring, retention, and monitoring; second, within-district variation in working conditions in the absence of flexible salaries could lead to substantial variation in quality; and, third, on the practical side, the methodology depends on variation in the proportion of new teachers divided by the number of teachers. In districts with many teachers, high values of the denominator overwhelm any variation in teacher turnover, and it may preclude detecting the effects of variations in quality. Nevertheless, for completeness, the empirical analysis measures competition at both the school and district level.

20. The proportion of teachers who are different must be divided by the number of teachers in a school because of the aggregation to grade averages. The total within-school variance in teacher quality includes not only variation across grades but also variation within grades. The variance of grade averages equals the total variance divided by the number of teachers per grade as long as the hiring process is identical for adjacent cohorts and grades.

21. Some specifications (not reported here) include controls for new principals and new superintendents. Each may directly affect the variation in achievement and be correlated with teacher turnover. However, the results are quite insensitive to the inclusion of these variables.

It is important to note that extensive teacher sorting among schools and districts on the basis of teacher quality could make it difficult to disentangle any behavioral effects of additional competition from a reduction in the average metropolitan area within-school variance in teacher quality that followed structurally from reducing teacher concentration in schools or districts. Compare the cases of three equally sized districts and four equally sized districts where teachers are sorted into a perfect quality hierarchy across districts. If teachers are drawn from the same initial distribution in both areas, the additional district will mechanically lead to smaller within-school variance. However, evidence from Rivkin, Hanushek, and Kain (2001) suggests that sorting of teachers on the basis of quality may in fact be quite limited in many areas. Work by Ballou (1996) and Ballou and Podgursky (1997) documents teacher hiring practices in which applicant skill does not play a primary role. Moreover, we have very little information about the nature of the teacher quality pool, but one would expect competitive pressures to change the hiring practices regarding where teachers are drawn from the overall distribution of skills in the labor market.

Nevertheless, the possibility remains that a lower Herfindahl index may be associated with more extensive sorting without inducing a behavioral change. In an effort to address this issue, we divide metropolitan areas up into separate school markets on the basis of student income, under the assumption that an expansion in the number of wealthy districts, while permitting increased teacher sorting, does not effectively increase the number of choices for poor children and vice versa. A finding that income-specific competition measures are more strongly related to the within-school variance in teacher quality than the overall competition measure would support the belief that competition induces a behavioral response.

Results

Tables 1.4 and 1.5 report the main results on the effects on variations in teacher quality of school and district competition, respectively.²² The focus of attention is the interaction of the proportion of different teachers and the Herfindahl index. The main effect for the within-school variance in teacher quality is the proportion of different teachers divided by the number of teachers, and the interaction term identifies how the variance in teacher quality is affected by different degrees of competition within metropolitan areas. Consistent with expectations, the estimates in table 1.4 using school-level competition are much more precise. No interaction coefficients using district competition are significant even at the 10 percent level.

The school competition results support the hypothesis that competition raises school quality through its effect on teacher personnel practices. All

22. The sample is restricted to schools with at least ten students in a grade, and the results are somewhat sensitive to this assumption.

Table 1.4 Estimated Effect of Math Teacher Turnover on the Squared Difference in School Average Test Score Gains Between Cohorts (absolute value of *t*-statistics in parentheses)

	No Fixed Effects	School Fixed Effects	School-by-Grade Fixed Effects
% math teachers different/# teachers	0.013 (0.85)	0.036 (2.58)	0.039 (1.46)
% different/# teachers · school Herfindahl Index	1.35 (2.60)	1.18 (2.38)	2.05 (2.01)
<i>N</i>	1,140	1,140	1,140

Notes: See text and Rivkin, Hanushek, and Kain (2001) for a description of the specification and estimation approach. All regressions control for the inverse of student enrollment and restrict the sample to students remaining in the same school.

Table 1.5 Estimated Effect of Math Teacher Turnover on the Squared Difference in District Average Test Score Gains between Cohorts (absolute value of *t*-statistics in parentheses)

	No Fixed Effects	School Fixed Effects	School-by-Grade Fixed Effects
% math teachers different/# teachers	-0.040 (1.55)	-0.030 (1.38)	0.100 (1.91)
% different/# teachers · school Herfindahl Index	0.11 (1.25)	0.06 (0.93)	-0.28 (1.56)
<i>N</i>	832	832	832

Notes: See notes to table 1.4.

interaction terms are positive and significant at the 5 percent level, even in the specification that includes school-by-grade fixed effects. In other words, less competition leads to a larger within-school variance in teacher quality. The magnitude of the interaction coefficients in the fixed effects model suggests that a one standard deviation increase in the degree of competition (a 0.02 point decline in the Herfindahl index) would reduce the within-school variance of teacher quality by roughly 0.09 standard deviations in the teacher quality distribution.

Although this effect size might appear small, it is in fact large relative to that of measured inputs such as class size. Rivkin, Hanushek, and Kain (2001) find that a one standard deviation reduction in class size (roughly three students per class) would lead to an increase in achievement of 0.02 of a standard deviation. In other words, effect sizes for class size reduction are between one-fourth and one-fifth as large as the effect size for competition and teacher quality.

Importantly, a metropolitan area-wide variable may provide a noisy measure of competition for most students and be susceptible to the structural

problems described earlier. Although the estimation cannot be easily divided in terms of individual high- and low-income students, it is nevertheless informative to focus on schools serving a large proportion of low-income students and those serving a small proportion. Therefore we divide the sample into schools in which at least 75 percent of students are eligible for a subsidized lunch and those in which fewer than 25 percent are so eligible (the middle category is excluded) and compute two Herfindahl indexes for each metropolitan area.

Table 1.6 reports the results for these two samples of schools. The results suggest that public school competition is much more important for lower-income students, for whom the interaction coefficients are positive and strongly significant. In contrast, the estimates for schools with very few lower-income students are small and statistically insignificant. To the extent that private school alternatives are much more relevant and place much more pressure on schools serving middle- and upper-middle-class students, this result is not altogether surprising, and it is consistent with the belief that the observed effects capture a behavioral response. At the very least, more should be learned about competition effects for lower-income and minority students, because most of the large urban districts in the country serve increasingly lower income populations.

In summary, these results provide support for the notion that competition affects teacher quality. Importantly, the inferences drawn about quality from estimates of effects on within-school variance rest upon the assumption that administrators do not systematically act to ensure the highest quality of teaching possible. Evidence from Ballou and Podgursky (1995) and Ballou (1996) of school hiring decisions not driven primarily by applicant quality supports the view that there is a great deal of slack in the hiring process. Moreover, the small number of teachers released on the ba-

Table 1.6 Estimated Effect of Math Teacher Turnover on the Squared Difference in School Average Test Score Gains between Cohorts, by School Demographics (absolute value of *t*-statistics in parentheses)

	Schools with >75% Eligible for Subsidized Lunch			Schools with <25% Eligible for Subsidized Lunch		
	No Fixed Effects	School Fixed Effects	School-by-Grade Fixed Effects	No Fixed Effects	School Fixed Effects	School-by-Grade Fixed Effects
% math teachers different/ # teachers	0.006 (0.23)	0.044 (1.61)	0.060 (1.15)	0.110 (2.90)	0.053 (1.55)	0.056 (0.94)
% different/# teachers · school Herfindahl Index	1.15 (2.50)	0.97 (3.71)	1.19 (2.11)	-0.18 (1.07)	-0.08 (0.55)	0.06 (0.21)
<i>N</i>	306	306	306	272	272	272

Notes: See notes to table 1.4.

sis of poor performance and anecdotal evidence of weak efforts by many teachers are consistent with lax monitoring procedures.

On the other hand, the positive coefficient on the interaction term could reflect the fact that schools in more competitive metropolitan areas hire more systematically but not any better than others—that is, that they hire more similar teachers but not ones of higher quality. Because it is not possible to identify the part of the variation in which the decline in variation is concentrated, statistical evidence cannot be brought to bear on this question. However, to the extent that alternative job opportunities tend to thin out or even truncate the right side of the distribution, it is quite likely that the reduction in variance occurs predominantly in the left tail. Moreover, there is little *a priori* reason to support the hypothesis that competition leads to more systematic but not necessarily any better personnel practices.

1.4 Conclusions

These results provide the first piece of evidence on the mechanisms through which competition may affect school quality; they suggest that more competition tends to increase teacher quality, particularly for schools serving predominantly lower-income students. Given the evidence that teacher quality is an important if not the primary determinant of school quality, a finding that competition was not related to the quality of teaching would have raised doubts about the strength of the link between competition and overall school quality.

Future work in this area should explore specific aspects of the teacher/management relationship such as hiring, tenure, and monitoring, as well as the effects of competition on the use of inputs and other aspects of school operations. Such information will provide a much better understanding of the processes that generate the observed link between competition and school quality. It is also relevant for efforts to improve the existing public schools.

References

- Abowd, John M., Francis Kramarz, and David N. Margolis. 1999. High wage workers and high wage firms. *Econometrica* 67 (2): 251–333.
- Armor, David J. 1995. *Forced justice: School desegregation and the law*. New York: Oxford University Press.
- Armor, David J., Patricia Conry-Oseguera, Millicent Cox, Niceima King, Lorraine McDonnell, Anthony Pascal, Edward Pauly, and Gail Zellman. 1976. *Analysis of the school preferred reading program in selected Los Angeles minority schools*. Santa Monica, Calif.: RAND Corporation.

- Ballou, Dale. 1996. Do public schools hire the best applicants? *Quarterly Journal of Economics* 111 (1): 97–133.
- Ballou, Dale, and Michael Podgursky. 1995. Recruiting smarter teachers. *Journal of Human Resources* 30 (2): 326–38.
- . 1997. *Teacher pay and teacher quality*. Kalamazoo, Mich.: W. E. Upjohn Institute for Employment Research.
- Black, Sandra E. 1999. Do better schools matter? Parental valuation of elementary education. *Quarterly Journal of Economics* 114 (2): 577–99.
- Borland, Melvin V., and Roy M. Howsen. 1992. Student academic achievement and the degree of market concentration in education. *Economics of Education Review* 11 (1): 31–39.
- Burtless, Gary. 1996. *Does money matter? The effect of school resources on student achievement and adult success*. Washington, D.C.: Brookings Institution.
- Coleman, James S., Thomas Hoffer, and Sally Kilgore. 1982. *High school achievement: Public, Catholic, and private schools compared*. N.Y.: Basic Books.
- Grogger, Jeffrey, and Derek Neal. 2000. Further evidence on the effects of Catholic secondary schooling. In *Brookings-Wharton papers on urban affairs, 2000*, ed. William G. Gale and Janet Rothenberg Pack, 151–93. Washington, D.C.: Brookings Institution.
- Hanushek, Eric A. 1971. Teacher characteristics and gains in student achievement: Estimation using micro data. *American Economic Review* 60 (2): 280–88.
- . 1986. The economics of schooling: Production and efficiency in public schools. *Journal of Economic Literature* 24 (3): 1141–77.
- . 1992. The trade-off between child quantity and quality. *Journal of Political Economy* 100 (1): 84–117.
- . 1997. Assessing the effects of school resources on student performance: An update. *Educational Evaluation and Policy Analysis* 19 (2): 141–64.
- Hanushek, Eric A., John F. Kain, Jacob M. Markman, and Steven G. Rivkin. Forthcoming. Does peer ability affect student achievement? *Journal of Applied Econometrics*.
- Hanushek, Eric A., John F. Kain, and Steve G. Rivkin. 2001a. Disruption versus Tiebout improvement: The costs and benefits of switching schools. NBER Working Paper no. 8479. Cambridge, Mass.: National Bureau of Economic Research, September.
- . 2001b. Why public schools lose teachers. NBER Working Paper no. 8599. Cambridge, Mass.: National Bureau of Economic Research, November.
- . 2002. New evidence about *Brown v. Board of Education*: The complex effects of school racial composition on achievement. NBER Working Paper no. 8741. Cambridge, Mass.: National Bureau of Economic Research, January.
- Hoxby, Caroline Minter. 1994. Do private schools provide competition for public schools? NBER Working Paper no. 4978. Cambridge, Mass.: National Bureau of Economic Research, December.
- . 2000. Does competition among public schools benefit students and taxpayers? *American Economic Review* 90 (5): 1209–38.
- Murnane, Richard J. 1975. *Impact of school resources on the learning of inner city children*. Cambridge, Mass.: Ballinger.
- Murnane, Richard J., and Barbara Phillips. 1981. What do effective teachers of inner-city children have in common? *Social Science Research* 10 (1): 83–100.
- Neal, Derek. 1998. What have we learned about the benefits of private schooling? Federal Reserve Bank of New York *Economic Policy Review* 4 (March): 79–86.
- Rivkin, Steven G., Eric A. Hanushek, and John F. Kain. 2001. Teachers, schools,

- and academic achievement. NBER Working Paper no. 6691 (revised). Cambridge, Mass.: National Bureau of Economic Research.
- Tiebout, Charles M. 1956. A pure theory of local expenditures. *Journal of Political Economy* 64 (October): 416–24.
- Weimer, David L., and Michael J. Wolkoff. 2001. School performance and housing values: Using noncontiguous district and incorporation boundaries to identify school effects. *National Tax Journal* 54 (2): 231–53.

