Priced Out: Aggregate Income Shocks and School Pricing in the Chilean Voucher Market

Benjamin Feigenberg *

September 8, 2014

Abstract

Private school market shares are rising steadily in many developing countries, but we have a limited understanding of how private schools set prices, how parents respond, and how this affects enrollment and performance in equilibrium. To shed light on demand behavior and supply response, I present a model of school pricing that incorporates an unusual feature of schooling compared to other goods – a potential preference by parents for small classes, and hence low school enrollment – that interacts with schools having market power. I show that, for a relatively broad range of parameter values, these two features can lead to the surprising result that an increase in aggregate household income, and hence an increase in willingness to pay for private schooling, can actually cause equilibrium private school enrollment to *decrease*. To investigate how private school enrollment responds to rising household income in practice, I exploit aggregate community-level income shocks in Chile, which has had a nationwide school voucher system since 1981. These shocks are caused by different responses to the price of copper in different municipalities. I show that private school prices rise by 0.9% in response to a shock that causes a 1% rise in income while private school enrollment falls by 2.0%. I find that falling private school enrollment is primarily caused by the middle-income students at the top schools. Those middle-income students induced to downgrade by rising prices do not experience the test score gains from the income shock experienced by students in the rest of the income distribution. I structurally estimate an extended version of the model and find that both market power and parental preferences for reduced class size are contributing to the observed declines in enrollment.

JEL: I24, I25, L1, O15

*Department of Economics, University of Illinois at Chicago (email: bfeigenb@uic.edu). I am grateful to my advisors David Autor, Esther Duflo, and Ben Olken for their guidance and to Abhijit Banerjee and Paulo Somaini for their extensive feedback. I also thank Dave Donaldson, Glenn Ellison, Eliza Forsythe, Francisco Gallego, Anil Jain, Conrad Miller, Rohini Pande, Juan Passadore, Miikka Rokkanen, Nancy Rose, Adam Sacarny, Annalisa Scognamiglio, Brad Shapiro, Tavneet Suri, Henry Swift and participants at the MIT applied micro, development, labor, and industrial organization field lunches for their valuable comments. I thank Amanda Dawes and the Chilean Ministry of Education for help in accessing the data.

1 Introduction

In recent years, parents in developing countries, dissatisfied with the quality and availability of public schooling, have increasingly turned towards the private school sector. In surveys of urban Indian slums, for example, the majority of students report attending private school (Tooley et al., 2007). In Colombia, one-third of students nationwide attend private school and that rate is even higher in urban areas, such as Bogotá, where over 70% of secondary schools are private (Bettinger et al., 2010; King et al., 1997). As the private school sector expands, the market for schooling will play an increasingly critical role in determining the quality of education to which children have access. Consequently, to be able to analyze how private school expansion will affect educational opportunities, we need to better understand how private schools set prices and enrollment levels, how parents respond, and how this affects performance and enrollment in equilibrium.

There are two reasons to think that understanding school pricing may be more complicated in this setting than in standard competitive market models. First, private schools may have substantial market power. Parents' idiosyncratic preferences for particular school attributes allow for extensive horizontal differentiation. Moreover, while students' potential choice set is typically large, they are often unwilling to travel large distances to school and so, in practice, choose from only a small number of schools. The fact that reputation plays an important role in schooling decisions (and takes time for schools to build) further suggests that school supply may be constrained in the short run in many environments. Second, school quality is decreasing in class size and, because the number of classrooms is typically fixed in the short run, school quality is decreasing in enrollment. This unusual feature of the educational market increases the likelihood that schools face downward-sloping demand curves in the price-enrollment space.

In this paper, I propose a model of private school supply and demand which incorporates schools' market power and the fact that school quality (e.g. class size) is a normal good, but adjustments to the number of classrooms in a school are infrequent.[1] Based on

[1] The model builds on industrial organization demand estimation research, including Berry et al. (1995) and Nevo (2000), that analyzes consumer choice within differentiated products markets. The

this model, I simulate schools' profit-maximizing pricing behavior in a simplified market in which consumers (i.e., parents) experience a positive aggregate income shock. I show that, for a wide range of parameter values, private schools' profit-maximizing strategy is to increase prices to the point that the equilibrium private school enrollment share declines as students switch to public schools. The finding that rising aggregate income could lead private school enrollment to fall in equilibrium with non-trivial probability makes the market for private schooling unusual even among the class of imperfectly competitive markets.

To investigate the response of private school prices and enrollment to aggregate changes in income in practice, I use Chilean labor force survey data to construct aggregate income shocks based on variation in global copper prices and in local elasticities of income with respect to copper prices. I then use administrative school price data provided by the Chilean Ministry of Education to examine how school prices respond to a positive aggregate income shock. I estimate equilibrium impacts on students' enrollment decisions and examine the extent to which enrollment impacts vary based on the education level of students' parents. I find that private school prices rise and enrollment levels fall in this setting. These effects are driven by those schools that were already the most expensive at baseline. I proceed to structurally estimate an extended version of the model in order to characterize the relative importance of the market power and class size channels in explaining estimated enrollment declines. I show that both features of the market are needed to explain the observed equilibrium declines in private school enrollment. Finally, I present estimates of test score impacts and compare them to observed changes in student enrollment patterns.

Chile provides an ideal setting for studying how the pricing behavior of profit-maximizing schools mediates the impact of municipality-level aggregate shocks that would unambiguously improve outcomes for students in a non-market educational environment. The advantages of the Chilean setting are two-fold. First, its private school market is well-

model is most closely related to Urquiola and Verhoogen (2009), which models education supply and demand in Chile in order to investigate how schools' enrollment and pricing decisions affect estimates of the impact of class size on student outcomes.

established and covers the whole country. Chile introduced its nationwide school voucher system in 1981 and allows private schools to accept vouchers while charging additional fees. At present, only 7% of Chilean students attend primary or secondary schools which do not participate in the voucher market. Second, the Chilean Ministry of Education collects data on private school enrollment and prices that is not available in other settings, and the Ministry links this data to individual student records. This linkage allows me to investigate how school pricing decisions impact students' enrollment responses and their academic performance.

Rigorous evidence on demand behavior and supply response within large-scale private school markets more broadly has been limited by data constraints as well as a lack of plausibly exogenous variation in demand for or supply of private schooling that can be isolated and studied.[2] While there is an extensive school voucher literature (see, for instance, Rouse, 1998; Angrist et al., 2002; Krueger and Zhu, 2004; Howell and Peterson, 2006) in which the authors use voucher lotteries to identify the causal impact of gaining access to private school on educational outcomes, voucher experiments have typically taken place in settings in which the group of voucher recipients was small relative to overall private school enrollment. As a result, researchers have been unable to use these experiments to study school price responses and the implications for students' school choices and academic achievement. There is also a large body of research that estimates the causal impact of market competition (i.e., the penetration of private voucher schools) on educational outcomes in Chile using cross-sectional data.[3] By identifying a source of ag-

[2] One exception is Muralidharan and Sundararaman (2013), in which the authors conduct a two-stage market-level and student-level school voucher randomization in Andhra Pradesh, India. In this paper, the authors find no evidence of spillover effects on students who did not receive vouchers.

[3] Here, authors are limited by the fact that the voucher system was introduced nationwide in 1981 as part of a larger educational reform. A number of early studies of the Chilean voucher system, such as Mizala and Romaguera (2000), employ OLS regressions of test scores on school type (private versus public) and include student demographic characteristics in an effort to control for selection. An alternative approach, employed in Sapelli and Vial (2002), uses a Roy-style selection model to estimate test score gains associated with public versus private schooling. More recently, researchers have sought out plausibly exogenous variation in the degree of market competition across Chilean municipalities. In Hsieh and Urquiola (2006), the authors instrument for municipality-level exposure to the voucher system using baseline municipality population, urbanization, and degree of inequality and find that increased school choice did not affect test scores or educational attainment but did lead to increased sorting based on student background. In contrast, Gallego (2013) uses the historical

gregate income shocks at the municipality by year level and by exploiting rich educational panel data, my work sheds light on how schools behave strategically *within* private school markets and how school pricing impacts students.[4] In doing so, this project builds on computational studies of education markets, such as Epple and Romano (1998), Nechyba (2000), and Ferreyra (2007). In that research, the authors simulate aggregate responses to tuition voucher policies based on varied assumptions about the structure of the market and the determinants of parental demand.

To conduct the reduced-form analysis, I construct plausibly exogenous municipality-level aggregate income shocks using variation in global copper prices and in how income in a given municipality responds to global copper prices. Specifically, I use historical global copper prices and pre-period labor force survey data on household income to construct municipality-specific measures of the elasticity of income with respect to global copper prices from the period before educational microdata is available. For the purposes of exposition, I normalize these municipality-specific elasticities. The interaction of the normalized elasticities with contemporaneous copper prices serves as the aggregate income shock. In the analysis, I focus on the effect of "positive copper shocks," which correspond to copper price increases in municipalities with positive normalized elasticities. Based on this identification strategy, I find that a positive copper shock has a relatively uniform impact across the log income distribution. Importantly, I confirm that the copper shocks constructed based on this methodology do not have any statistically significant impact on private school prices from previous years. I then proceed to estimate the effect of these positive copper shocks on school prices, student enrollment patterns, and student test scores.

distribution of Catholic priests to instrument for the concentration of voucher schools and finds that an increase in the ratio of voucher to public schools led to increased test scores in both public and private schools.

4 Andrabi et al. (2013) examines the strategic behavior of schools in low-information environments. In this study, the authors provide parents in rural Pakistan with report cards on school and student test scores, and they identify significant school price, quality, and enrollment changes in response. However, the Chilean market is distinct in that parents appear to be better-informed about school quality. Mizala and Urquiola (2013) presents evidence that a government program designed to identify effective schools had little impact on enrollment or tuition levels.

School price and enrollment responses to rising aggregate income make it clear that the education market does not behave like a standard competitive market. I estimate that a 1% increase in aggregate household income has a positive (0.9%) impact on private school prices. However, this 1% aggregate income increase causes a 2.0% decline in private school enrollment as students move to public schools. This negative enrollment response is inconsistent with a standard perfect competition model of the market in which private schools expand (or enter the market) to meet increased demand and so both private school prices and enrollment shares rise with aggregate income. Changes in school prices and enrollment are not uniform within the private school sector. While the most elite private voucher schools increase prices by 1.4% in response to the 1% aggregate income shock, those private schools charging the lowest prices at baseline do not raise prices at all. Correspondingly, top schools become significantly more exclusive (reducing enrollment by 3.2%), while low-end private schools increase enrollment levels.

Given rising prices at elite private schools, heterogeneity in baseline enrollment patterns and in parents' willingness to pay for private schooling are the key determinants of whether declines in the quality of school attended are universal or are driven by a particular subpopulation of students. To examine changes in student enrollment decisions, I exploit the availability of unique student identifiers that allow students to be tracked across years and across schools. Grouping students into bins based on their parental education level, I find that declines in the average baseline price of school attended are driven by students in the two middle quartiles of the parental education distribution who attended elite private schools at baseline.[5] Middle parental education students correspondingly exhibit the smallest test score gains (although all subgroups improve their average test score performance in response to the income shock). I provide suggestive evidence that smaller gains for moderate parental education students are indeed driven by their relatively higher rates of school downgrading, and this implies that those students

[5] Low parental education students do not significantly reduce the baseline price of schools attended because they are unlikely to be attending elite private schools at baseline. At the other end of the distribution, high parental education students experience smaller declines in the baseline price of schools attended because they are more likely to be able and willing to pay to stay in elite private schools after these schools increase prices.

who downgrade actually experience large test score *losses* in response to rising aggregate income.

In order to characterize the mechanism driving observed declines in private school enrollment, I proceed to estimate the model based on the demand curve estimation approach introduced in Berry et al. (1995).[6] I calculate how the elasticity of residual demand faced by private schools changes in response to a positive aggregate income shock and then simulate price and enrollment impacts by finding a Nash equilibrium in the market (i.e., a set of prices at which all private schools are best-responding). Estimates reveal that schools' market power and parents' class size preferences cause reductions in demand elasticity which make it profitable for schools to raise prices substantially. At the same time, parents' preferences for reduced class size are essential in ensuring that elite private schools prefer to raise prices and reduce enrollment rather than leave prices relatively fixed and instead profit from substantial enrollment increases. Structural simulations produce an increase in public school enrollment and private school prices that qualitatively matches the reduced-form findings. In addition, the simulations predict larger price increases and enrollment declines at elite private schools (as observed in the data).

The remainder of the paper is structured as follows. Section 2 presents a model of the private school education market that predicts how private school prices and enrollment levels respond to rising aggregate income. Section 3 documents institutional details of the Chilean educational market. Section 4 discusses the data used in the empirical analysis, and Section 5 outlines the identification strategy. Section 6 estimates the impact of positive copper shocks on school prices and enrollment. Section 7 investigates the impact of positive copper shocks on student school switching patterns and test scores. Section 8 structurally estimates an extended version of the model, and Section 9 concludes.

[6] The particular two-step algorithm used is outlined in Berry et al. (2004) and detailed in Gallego and Hernando (2009).

2 Model

In standard competitive market models, increased aggregate demand causes increases in equilibrium prices and quantities. The market for private school education, however, has two key features that make it distinct from standard markets and suggest that schools' strategic pricing behavior could potentially cause private school enrollment to fall in response to an increase in household income and aggregate demand. First, schools have market power that results from short-run supply constraints, idiosyncratic parental preferences, and the fact that most students choose between a relatively small number of schools located in close proximity to their homes. Second, school quality is a normal good, and school quality is decreasing in average classroom size. In this section, I propose a model that incorporates these two features of the private school market and outlines the conditions under which they will cause equilibrium private school enrollment to fall in response to increased aggregate demand.

I assume that the utility function of individual i is defined as in the standard vertical differentiation model but with the inclusion of an additional error term:

$$U_{ij} = \alpha v_i s_j - p_j + \epsilon_{ij} \tag{1}$$

In this framework, α represents an aggregate willingness to pay shifter, each individual has willingness to pay for quality v_i, and each school j has an associated quality s and price p. As is common in the literature, this expression assumes that schools choose an overall price level but do not have the capacity to price discriminate across students.[7] The error term captures the household's idiosyncratic utility gain from having a child attend school j. The error term has a mean zero Type I extreme value distribution, as in the standard aggregate logit model.[8] The scale parameter of the error term, defined as σ, can be interpreted as a measure of the extent of school differentiation within the market. I

[7] In practice, my own analysis of Chilean educational survey data suggests that there is limited but non-zero variation across students in tuition paid for a given school and grade.

[8] See, for instance, Nevo (2000).

assume that willingness to pay has probability density function $f(v)$ and is distributed on the interval $\underline{v}, \bar{v}$. Given these assumptions, the expression for the share of individuals with willingness to pay v who attend school j is defined as:

$$\Gamma_j(\underline{s}|\alpha, v, p, s) = \frac{exp^{\frac{\alpha v s(\underline{s}) - p(\underline{s})}{\sigma}}}{\int_{\underline{\hat{s}}} exp^{\frac{\alpha v s(\underline{\hat{s}}) - p(\underline{\hat{s}})}{\sigma}} g(\underline{\hat{s}}) d\underline{\hat{s}}} \tag{2}$$

where $\underline{s}$ represents the baseline quality of a school and $\underline{\hat{s}}$ spans the support of the baseline school quality distribution. If the total population is N, the number of students choosing to attend school j with price p and quality s is given by:

$$D_j(\underline{s}|\alpha, v, p, s) = N \int_{\underline{v}}^{\bar{v}} \Gamma_j(\underline{s}|\alpha, v, p, s) f(v) dv \tag{3}$$

Before advancing to the dynamics of school price and enrollment decisions, I introduce assumptions about school operating costs and the school quality production function. I assume that there are two types of costs faced by schools: (1) F_n, a cost of operating each classroom, and (2) c, a per-student variable cost. Then, the profit function for school j can be expressed as follows:

$$\Pi_j = (p_j - c)q_j - n_j F_n \tag{4}$$

where q_j is total enrollment and n_j is the number of classrooms in school j.

In the market for private school education, school reputation is typically an important determinant of perceived quality and is relatively fixed in the short run.[9] Consequently, I model school quality s as a function of $\underline{s}$ and q. Enrollment is a key determinant of perceived quality due to parents' preference for reduced classroom size (and possibly for school exclusivity).[10] One important modeling decision I make is in assuming that the number of classrooms in a school is fixed in the short run. This assumption improves the

[9] For more discussion on this topic in the Chilean context, see Gallego (2013).

[10] For a discussion of the relationship between classroom size and school quality in the Chilean context, see Urquiola and Verhoogen (2009). McEwan (2013) presents a meta-analysis of randomized experiments in developing-country settings and finds that a group of treatments that includes classroom size reductions has a significant positive impact on student test scores.

tractability of the model by reducing the number of choice variables from three to two and appears to be satisfied in the Chilean setting (in which I test the predictions of the model).[11] I assume that the functional form of the relationship between school quality and enrollment is characterized as follows:

$$s_j = \underline{s}_j - \tau(\frac{q_j}{n_j}) \tag{5}$$

In this equation, τ reflects the relative importance of classroom size in the school quality production function. The expression implies that school quality is decreasing linearly in classroom size.[12]

Given the above expressions for school profit, school quality, and household demand for private schooling, I derive comparative statics for school price and enrollment responses to shifting demand.[13] I focus on the case of a positive shock to aggregate demand, represented by an increase in α. To start, I define the school's profit maximization problem in terms

[11] Specifically, this assumption is justified by the fact that observed changes to the number of classrooms never serve to offset the impact of observed enrollment changes on average classroom size in Chile. In practice, public schools and low-cost private schools slightly increase the number of classrooms as they increase enrollment, but average classroom size still rises significantly at these times. High-cost private schools slightly reduce the number of classrooms as they decrease total enrollment, but average classroom size still falls significantly as a result of these changes. Thus, changes in enrollment, rather than changes in the number of classrooms, appear to determine aggregate changes in average classroom size.

[12] This parameterization of the school quality production function differs from that in Urquiola and Verhoogen (2009). There, the authors model the impact of classroom size reductions on school quality to be increasing in baseline school quality and proportional to percentage changes in classroom size. My assumption that the impact of classroom size reductions on school quality is independent of baseline quality is motivated by the finding (discussed in Section 8) that heterogeneity in parental preferences for reduced classroom size by baseline characteristics is relatively limited. Since wealthier students are more likely to attend schools with higher baseline quality, these students would exhibit stronger preferences for reduced classroom size if associated quality gains were indeed larger in higher baseline quality schools. I assume that quality decreases linearly in classroom size to reduce the incentive for schools to dramatically reduce class size (to near zero) in order to raise prices as this behavior is not consistent with the data. Nonetheless, model simulations look similar if I alternatively assume that school quality changes are proportional to percentage changes in classroom size.

[13] I assume that the number of schools in the market is fixed in the short run. This assumption is supported by the fact that there is no significant change in Chilean school entry/exit in response to aggregate income shocks in the empirical application (see the discussion in Section 6.1).

of prices:

$$\max_p \Pi(p, q(p, \underline{s})) \tag{6}$$

Here, equilibrium enrollment $q^*(p, \underline{s})$ is determined by parental demand for a school with baseline quality $\underline{s}$ that charges price p. In this setting, parents are fully informed about the distribution of willingness to pay and about school costs. As a result, the market clearing set of prices and enrollment levels is a fixed point at which q^* parents are willing to pay p^* for their children to attend a school with quality $s(\underline{s}, q^*)$.[14]

Based on the first order condition from the school's profit maximization and the expressions for demand and expected willingness to pay of parents, I arrive at a set of three equations characterizing equilibrium prices, enrollment, and expected willingness to pay of parents whose children attend school j:

$$p^* - [\sigma + c + \alpha\tau V^* \frac{q^*}{n}] = 0 \tag{7}$$

$$q^* - \int_{\underline{v}}^{\bar{v}} \Gamma_j(\underline{s}|\alpha, v, p^*, s^*) N f(v) dv = 0 \tag{8}$$

$$V^* - \int_{\underline{v}}^{\bar{v}} v \frac{\Gamma_j(\underline{s}|\alpha, v, p^*, s^*)}{\int_{\underline{v}}^{\bar{v}} \Gamma_j(\underline{s}|\alpha, v, p^*, s^*) f(v) dv} f(v) dv = 0 \tag{9}$$

Here, $p^*, q^*,$ and V^* represent equilibrium prices, enrollment, and expected willingness to pay of parents whose children attend school j, respectively.

Finally, I apply the implicit function theorem to find expressions for $\frac{dp^*}{d\alpha}$ and $\frac{dq^*}{d\alpha}$. These expressions are derived and presented in the Mathematical Appendix, but they are algebraically complex and provide little intuition regarding the sign or magnitude of price and enrollment impacts. This intractability is a consequence of the fact that all schools adjust prices in response to a change in α. As a result, a school's own price setting decision must take into account the full set of price changes undertaken by other schools in the market. To simplify the problem, I simulate equilibrium price and enrollment changes in a market with one public school, one private school, and 45 students (the maximum

[14] Regarding the dissemination of class size information, Mizala and Urquiola (2013) notes that Chilean schools may make commitments to prospective parents regarding class size.

classroom size permitted by law). In this market, the public school has a fixed price equal to marginal cost, while the private school adjusts price and enrollment to maximize profits. The private school is assumed to have a baseline quality that is equal to two times the baseline quality of the public school. I present simulation results averaged over 50 replications.

In Figure 1, the blue-shaded region characterizes the set of parameter values at which rising aggregate demand leads to a price increase and a decline in enrollment at the private school. The simulations reveal that the model can produce a private school price increase and enrollment decline for a range of values of the classroom size preference parameter (τ) and the horizontal differentiation parameter (σ). Notably, the simulations suggest that both classroom size preferences and horizontal differentiation are necessary to generate a private school enrollment decline. The region in which the simulation produces an enrollment decline is primarily characterized by values of the τ/σ ratio close to 0.1. There are two instances in which simulations do not produce private school enrollment declines. First, when the ratio τ/σ is sufficiently large (denoted by the yellow region of Figure 1), the model is unstable as small changes in initial conditions (i.e., the assumed willingness-to-pay distribution) produce very different equilibrium prices and enrollment levels. This is a result of the fact that when classroom size preferences are strong and the market is close to competitive, small changes in class size cause large shifts in enrollment patterns. Second, in the red-shaded region, private school enrollment increases in equilibrium as private school price rises. In the upper portion of the graph, strong differentiation and weak class size preferences prevent students from switching to public school when prices and class size rise. In the bottom portion of the graph, strong class size preferences and relatively weak differentiation constrain the magnitude of price and class size increases such that the net effect on private school enrollment is positive.

3 Education in Chile

The model highlights the possibility that the structure of education markets could lead changes in market conditions to have unanticipated impacts on private school enrollment levels. However, whether rising aggregate income causes enrollment declines in practice is an empirical question. Chile provides an ideal environment for studying private schools' price-setting behavior and the implications for students' enrollment decisions and academic performance because Chile's school voucher system is both expansive (covering the whole country) and well-established. The voucher system in Chile was introduced in 1981 as part of a nationwide educational reform which (1) significantly reduced government funding for education, (2) decentralized educational decision-making to the local (municipal) level, (3) relaxed curriculum standards, (4) revoked teacher union contracts, and (5) established a system of school voucher funding whereby a given voucher value is paid to the school that a student attends regardless of whether it is public or private (Bravo et al., 2010, Hsieh and Urquiola, 2006). At the time that the reform was enacted, private schools had only a 22% market share in Chile (Gallego, 2013).

At present, 50% of students in grades one through twelve attend private (voucher) schools, 43% attend public schools, and 7% attend unsubsidized private schools.[15] Under the voucher system, all public primary schools are free, and all public secondary schools can charge at most $7 USD per month.[16] In contrast, primary and secondary voucher schools are permitted to charge a "top-up" that is up to three times the annual voucher amount. Voucher values are determined by grade level and length of school instruction (full- or half-day), with the voucher value in secondary school depending additionally on

[15] Students attending unsubsidized private schools are excluded from the analysis because the government does not maintain data on prices charged by these schools, and because these schools are not subjected to the same regulations as public and private voucher schools. In practice, the key distinguishing feature of unsubsidized private schools is that they cater to a much wealthier population than either public or private voucher schools. Other research on the Chilean voucher system, including Urquiola and Verhoogen (2009) and Gallego (2013), also excludes unsubsidized private schools from the analysis. Using enrollment data, I have confirmed that unsubsidized private school enrollment does not increase in response to observed declines in enrollment at elite private voucher schools.

[16] Mizala and Urquiola (2013) notes that few public secondary schools charge any fee in practice.

whether a student is studying in the vocational or college-preparatory track.[17] In 2011, the average voucher value was \$110 USD, and the average "top-up" charged was \$34 USD. Officially, private schools can selectively admit students while public schools that are not "at capacity" are required to admit all applicants. However, as noted in Gallego and Hernando (2009), school-side screening appears limited based on the following evidence: 93% of parents report that their children attend the parents' preferred school, the average number of schools to which a student applies is 1.1, and only 4% of parents say their child was rejected from at least one school.

Most private voucher schools in Chile are profit-maximizing. Indeed, Elacqua (2009) finds that over 75% of Chilean private voucher schools are for-profit, and Urquiola and Verhoogen (2009) notes that even those schools that are officially not-for-profit can distribute dividends to principals and/or school board members. There is also reason to believe that individual schools may have substantial local market power. The median number of primary schools within 2.5km, 5km, and 10km of a municipality center is 3, 4, and 9, respectively. Additionally, the median primary school student travels only 2.5 km to school (based on the population centroid of his or her home municipality). Consequently, it appears that most students choose from a small number of schools. Moreover, while there is limited governmental regulation of school openings, private school supply may be constrained by reputational factors (Gallego, 2013), and the expansion of existing schools is constrained, at least in the short run, by the capacity of the school's physical plant.

In Chile, school switching, which represents an important outcome in the analysis, is relatively common: 17% of students switch schools each year, including 12% of students not entering ninth grade. While students are free to attend school in any municipality, high cross-municipality commuting times imply that students' choice set is limited, in practice, to those schools within their home municipality: 83% of all students, and 88%

[17] I adjust school prices to reflect differences in voucher values arising from grade level and type of secondary school track. However, I do not account for differences in voucher values determined by whether a school provides full- or half-day instruction as I cannot distinguish instruction schedules from the data.

of primary school students, attend school in their home municipality. Previous research, such as Gallego (2013) and Hsieh and Urquiola (2006), has consequently defined education markets by municipality borders.

4 Data

4.1 Copper Shocks

My identification strategy requires that I estimate municipality-specific elasticities of income with respect to copper prices. To do so, I collect annual global copper prices (denominated in 1998 USD) from the United States Geological Survey within the Department of the Interior (U.S. Geological Survey, 2013). Survey data on historical municipality-level incomes comes from the Chilean National Socio-Economic Survey (CASEN) for which data is available for the following years: 1990, 1992, 1994, 1996, 1998, 2000, 2003, 2006, 2009, and 2011. The survey provides a repeated cross-section that lists respondents' municipality of residence. As one of its objectives, the CASEN survey is designed "[t]o characterize the population according to household income, quality of housing, education, participation in the labor market, composition of family income, and other relevant variables" (Ministry of Social Development, 2013). The CASEN survey is representative at the national, regional, and, in some cases, municipality level. Previous work, such as Auguste and Valenzuela (2006), has also used CASEN survey data to construct municipality-level socioeconomic variables.

4.2 Educational Administrative and Survey Data

Administrative data from the Chilean Ministry of Education provides a roster of all students enrolled in Chilean schools in each year from 2002 to 2012. Each student is tracked with a unique identifier, which allows researchers to follow students across years and to merge administrative data with educational survey data. The administrative data file provides the school attended and grade level of each student in each year along with a measure of academic performance, municipality of residence, and basic demographic information such as birth date, age, and gender.

In addition to administrative records on Chilean students, I obtained a data file from the Ministry of Education containing mean school prices for fee-charging private voucher schools for the years 2004 to 2011.[18] The data file also contains prices for a small number of public schools, which are allowed to charge nominal fees for secondary school students. For all other public secondary schools, I have assumed that no price is charged in excess of the school voucher. Given the low maximum price that may be charged by public secondary schools, measurement error along this margin is unlikely to significantly affect the results.

The Chilean Ministry of Education also releases annual test score data. Prior to 2006, the Ministry administered the national Educational Quality Measurement System Exam (SIMCE) to one grade level across the country each year, rotating among grades four, eight and ten. After 2006, they tested fourth graders annually and eight and tenth graders biennially, rotating between the two. According to the Ministry, the test is designed to improve educational outcomes by providing an external measure of students' mastery of the curriculum (Agency of Education Quality, 2013). In the analysis, I average student scores on the language and math components of the exam and normalize scores at the grade-year level.[19] For each cohort that takes the SIMCE, the Ministry of Education also collects detailed survey data from teachers and parents. The data collected from teachers includes proxies for school quality, such as teacher training, teacher experience, and classroom technology use. The survey of parents provides information on household demographic characteristics, such as household size, approximate household income, and parental education. This measure of parental education plays a central role in the student-level heterogeneity analysis detailed in Section 7.

[18] This data file contains prices charged by 65% of private voucher schools which serve 80% of voucher school students. Based on the literature (see, for instance, Mizala and Urquiola, 2013), it appears that virtually all of the missing schools do not charge fees. School-level estimates are robust to setting the fee equal to zero for these schools and including them in the analysis.

[19] Analysis is limited to the language and math components because these are the only subject scores that are available for each year and grade.

5 Identification Strategy

This research examines how private school price and enrollment responses to aggregate income shocks can shed light on demand behavior and supply response in the Chilean voucher market. To generate aggregate income shocks, I rely on cross-sectional variation in municipality-specific elasticities of income with respect to copper prices in combination with time-series variation in global copper prices. The preferred income measure is referred to in the CASEN survey as "autonomous income" and includes employment income, non-employment income, and government transfers associated with individual work history, but excludes other government transfers. This measure has been used in previous research aimed at characterizing the distribution of income in Chile (see, for instance, Solimano and Torche, 2008) and has the advantage of being available in all survey years.[20] The sample included for the calculation of municipality-specific income statistics is all respondents aged 18-65. To construct municipality-specific elasticities, I estimate the following equation using municipality-level data from 1990-2000 (the "pre-period"):

$$Y_{mt} = \alpha_m + \beta_m * P_t + \epsilon_{mt} \tag{10}$$

In this equation, Y_{mt} represents log mean earnings in municipality m in year t, and P_t is the log of the world copper price in year t, denominated in 1998 USD. β_m coefficients can be interpreted as the municipality-specific elasticities of income with respect to global copper prices. In practice, each municipality will have at most six observations over this time period, and not all municipalities are included in the CASEN labor force survey in each year. Consequently, the sample is restricted to the 270 of 346 Chilean municipalities that appear in at least three rounds of the CASEN survey between the years 1990 and 2000.[21] Excluded municipalities are, in general, either smaller than covered municipalities

[20] Results look very similar if a measure of employment income alone is used instead.

[21] β_m is just identified with two observations. However, the 20 municipalities with only two observations are excluded because estimated elasticities for these municipalities are very noisy compared to elasticities estimated based on 3+ observations. Appendix Table 1 confirms that results are robust to estimating elasticities at the regional level, in which case data is available for all regions in all years.

or are too new to have been covered in a sufficient number of survey years.

Figure 2 overlays the estimated municipality-level coefficients onto a map of Chile that includes municipality borders. There does not appear to be any clear pattern in the geographical distribution of coefficient values, although the map does make clear that municipalities without estimated copper coefficients are most likely to be located in the sparsely populated north-south extremes of the country. While the constructed municipality-specific coefficients are only important in relative terms since year fixed effects and municipality fixed effects are included in all subsequent specifications, it may initially seem counterintuitive that the mean municipality-level coefficient is negative-valued. The negative relationship between log incomes and log copper prices during this period appears to be explained by the canonical "Dutch Disease" phenomenon. In the case of Chile, increasing global copper prices are negatively correlated with the Chilean Peso (CLP) to USD exchange rate during the relevant years.[22] Consequently, as global copper prices rise, export-oriented industries suffer. The fact that Chilean exports represent 34% of gross domestic product (compared to, for example, 20% in Argentina) makes the relationship between copper prices and export industry competitiveness particularly salient (World Bank, 2013). Indeed, in the past two years, as copper prices have risen dramatically, exporters concerned about currency appreciation have advocated for capital controls and the Chilean government has invested $12 billion in a program aimed at weakening the currency (Pica and Wisnefski, 2012).

To provide additional evidence that currency appreciation drives the negative income responses to copper price increases calculated for most municipalities, I estimate industry-specific elasticities of income with respect to copper prices using CASEN labor force survey data. The four industries with the negative elasticities largest in magnitude are: timber extraction (-4.5), restaurants (-3.4), construction (-3.3), and "large-scale businesses" (-2.2). The industries with the most positive elasticities are: communications (1.0), public instruction (1.0), garments (2.5), and non-ferrous metals (3.7). It is unsurprising that the non-ferrous metals elasticity is largest in magnitude as this category includes copper

[22] See Figure 3 for the time-series plots of copper prices and the exchange rate.

production. More generally, the set of industries with negative elasticities is primarily export-oriented or tourism-related (wood products are Chile's second largest export after copper).[23] In contrast, industries with positive elasticities (other than copper) appear to be more domestically-focused. When I estimate the cross-sectional correlation between regional elasticities and regional industry shares, shares of the following industries are positively correlated with local elasticities: agriculture, public administration, and electricity/gas/water.[24] Finance, manufacturing, housing, business, and hotel/restaurant shares are negatively correlated with regional elasticities, while mining is positively correlated with regional elasticities in copper-producing regions and negatively correlated in regions that do not produce copper. These correlations are again consistent with the hypothesis that negative income responses to copper price increases are caused by reduced export competitiveness. Indeed, the value of regional exports relative to total regional economic activity negatively predicts regional elasticities in non-copper producing regions, while the relationship is positive in those regions that produce copper.

Given that the municipality-specific elasticities are only important in relative terms, I normalize the coefficients to have a mean of zero and a standard deviation equal to one. Then, I define the normalized copper shock C_{mt} assigned to municipality m in year t as being equal to this normalized municipality-specific coefficient $\bar{\beta}_m$ multiplied by P_t, the log of the global copper price in year t (denominated in 1998 USD). For the purposes of exposition, an increase in C_{mt} is referred to in the remainder of the paper as a "positive copper shock."

To test whether copper shocks can predict income levels in the post-2000 period for which school price and enrollment data are available, I estimate the following equation using income data from the CASEN survey waves of 2003, 2006, 2009, and 2011:

$$Y_{mt} = \alpha + \beta * C_{mt} + \gamma_m + \lambda_t + \epsilon_{mt} \tag{11}$$

[23] For a ranking of Chilean exports, see Chilean Customs (2013).

[24] As noted, Appendix Table 1 shows that reduced-form estimates generated using regional copper shocks are consistent with estimates from the benchmark specification that identifies impacts using municipality-level variation.

In this equation, γ_m and λ_t represent municipality and year fixed effects, respectively. In this and in all subsequent specifications, standard errors are clustered at the municipality level and calculated by bootstrapping the two-step procedure in which municipality-specific elasticities are first estimated (based on Equation 10) and then included in the set of regressors in a second specification (such as Equation 11).

Table 1 provides evidence that the copper shocks constructed using historical data on incomes and copper prices have a significant impact on household incomes during the period for which educational data is available. Column (1) implies that a one standard deviation positive copper shock is associated with a 2.9% increase in the mean income of municipality residents, and the estimate is statistically significant at the 1% confidence level.[25] Columns (2)-(6) estimate the impact of positive copper shocks on income deciles by replacing the log mean income dependent variable with dependent variables measuring the logs of the 50th through 90th percentiles of the earnings distribution.[26] While estimates differ somewhat across columns, there is no evidence of a monotonic relationship between income decile and estimated copper shock impact that might explain the evidence of heterogeneity in school choice impacts based on family background that is detailed in Section 7.

6 School Prices and Enrollment

6.1 Panel Estimates

When aggregate private school demand rises with aggregate income, equilibrium private school prices and enrollment must adjust in order for the market to clear. As outlined in

25 In terms of interpretation, a one standard deviation positive copper shock implies that copper prices double in a municipality that is at the 85th percentile in the distribution of elasticities. This makes clear that a positive copper shock of even one standard deviation is quite uncommon in practice given that the average annual change in the log copper price during the period being studied is only 0.23.

26 Below the 40th percentile of the earnings distribution, a significant share of municipalities have income measures that are always equal to zero and so sample selection issues emerge. Using a Poisson QMLE specification to account for zero-valued observations, I confirm that estimated impacts at the 40th percentile look comparable to 50th percentile estimates.

Section 2, the enrollment response to a change in aggregate income is ambiguous. This section estimates both price and enrollment relationships. The benchmark specification used to estimate average school-level price impacts is the following:

$$P_{smt} = \alpha + \beta * C_{mt} + \gamma_s + \lambda_t + \epsilon_{smt} \tag{12}$$

P_{smt} is the log mean total price charged by school s in municipality m in year t. This measure is calculated as the log of the sum of the mean "top-up" charged by school s in year t plus the mean voucher value received by school s in year t.[27] The remaining variables are as defined in Equation 11, except that municipality fixed effects are subsumed by school fixed effects, γ_s.[28]

Since public school prices are determined by voucher values, they should not respond to changes in aggregate income. The following specification incorporates this feature of the market by allowing aggregate income shocks to differentially affect public and private school prices:

$$P_{smt} = \alpha + \beta_1 * C_{mt} + \beta_2 * C_{mt} * V_s + \gamma_s + \lambda_t + \epsilon_{smt} \tag{13}$$

Here, V_s is an indicator variable equal to one for private voucher schools. While this specification allows price changes to depend on whether a school is public or private, it does not allow for price changes to vary heterogeneously *within* the private school sector. Since low-end private schools more closely resemble public schools than higher-priced private voucher schools with regards to the students they attract and the prices they charge, it seems likely that average private school price impacts mask substantial heterogeneity. Any such heterogeneity will, in turn, significantly alter the impact of rising aggregate income on students' enrollment decisions. To test for differential school price changes within the private school sector, I define an indicator variable P_{sq} which is equal

27 Survey evidence suggests that within-cohort price discrimination based on student characteristics is limited in Chilean private schools. Consequently, in analyzing school price responses to shifting willingness to pay, school-level specifications are appropriate.

28 Although a small number of schools switch municipalities during the relevant years, all schools are assigned to the municipality in which they are located at baseline to eliminate concerns about endogenous relocation.

to one if the tuition charged at baseline (in 2004, the first year for which price data is available) by school s falls into quintile q within municipality m.[29] Then, I estimate the following specification:

$$P_{smt} = \alpha + \beta_1 * C_{mt} + \sum_{q=2}^{5} \beta_q * C_{mt} * P_{sq} + \gamma_s + \lambda_t + \epsilon_{smt} \tag{14}$$

In each municipality, quintile one roughly corresponds to public schools and quintiles two through five to the four quartiles of the private school distribution. There are an average of 24 schools per municipality in the sample and quintiles are defined so that cutoffs do not separate schools charging the same price within a single municipality.[30] Research on the Chilean education sector suggests that private voucher school quality is, on average, higher than public school quality.[31] Within the private school sector, price is highly positively correlated with measures of perceived school quality, such as average test scores and students' average household income. As a result, estimates of impact heterogeneity based on baseline school price can be readily interpreted as reflecting heterogeneity in impacts based on baseline school quality.[32]

[29] In practice, 92.9% of schools in the sample appeared in 2004 and results are robust to including only these schools or assigning the remaining schools to a baseline price quintile based on the price charged in the first year in which they were in operation.

[30] Appendix Figure 1 plots the distribution of prices by baseline school price quintile for a sample year (2006).

[31] In a meta-analysis, Drago and Paredes (2011) concludes that private school test scores are approximately one-tenth of a standard deviation higher than public school scores conditional on student characteristics.

[32] In Appendix Table 2, I estimate the following specification to provide evidence that this is the case:

$$P_{sgmt} = \alpha + \beta * X_{sgmt} + \gamma_{sg} + \lambda_t + \epsilon_{sgmt} \tag{15}$$

where P_{sgmt} is the log price for grade level g in school s in municipality m in year t (defined as the log of the sum of the average school price and the grade-specific voucher value) and γ_{sg} represent school-grade fixed effects. This equation is estimated at the grade four level, as SIMCE data is most frequently available for grade four students. X_{sgmt} represent fourth grade-specific quality measures: mean test scores, mean parental education of classmates, mean log household income of classmates, whether classmates' parents expect them to graduate from college, average years of teacher experience, fraction of certified teachers, fraction of teachers with graduate degrees, and the fraction of teachers that frequently use computers for work. In Appendix Table 2, I also estimate the cross-sectional correlation between these characteristics and school prices. Coefficients on student test scores, mean parental education of classmates, mean log household income of classmates, and whether classmates' parents expect them to graduate from college are positive across specifications

Column (1) of Table 2 estimates Equation 12 and indicates that a one standard deviation positive copper shock causes an average price increase of 0.42% (significant at the 1% confidence level). Column (2) estimates Equation 13 and reveals that this small aggregate price change masks a significantly larger 2.74% increase in private school prices. The level term coefficient estimated in Column (2) is not statistically different from zero, reflecting the fact that public school prices do not change in response to copper shocks. Column (3) estimates Equation 14 and interaction term coefficients provide evidence that there is substantial heterogeneity in school price impacts based on baseline school price/quality. The coefficients on the interactions between the copper shock measure and price quintiles two and three are small in magnitude. In contrast, the coefficients on the interaction terms for quintiles four and five imply that a one standard deviation positive copper shock leads to a 2.24% increase in private school prices in the fourth quintile of schools and a 4.21% increase in prices in the fifth quintile. These findings indicate that the average private school price impacts identified in Column (2) are driven almost entirely by private schools in the upper half of the baseline price distribution.

Columns (4)-(6) of Table 2 re-estimate Equations 12-14, but replace the dependent variable P_{smt} with E_{smt}, a measure of enrollment in school s in municipality m in year t. Enrollment data is available for the full universe of public and private schools and for a number of years for which school price data is not available. Nonetheless, for the sake of comparability, the school-level sample is restricted to include only those school-year observations for which price data is available. Column (4) estimates mean school-level enrollment impacts and reveals that there is a small but statistically significant increase in average enrollment in response to a positive copper shock (the coefficient of 6.4 is equivalent to a 1.8% enrollment increase). Column (5) reveals that small average enrollment increases mask large enrollment declines in private schools and large increases in public school enrollment. Estimates suggest that average private school enrollment falls by 34.1 students (5.8%) per school and public school enrollment rises by 14.4 students

and are significant at the 1% level in most cases.

(5.6%) per school in response to a one standard deviation positive copper shock.[33] Column (6) indicates that evidence of price impact heterogeneity is matched by heterogeneous changes in average enrollment. The level term in Column (6) implies that enrollment in quintile one schools rises by 12.9 students (5.2%) per school in response to a one standard deviation positive copper shock. Interaction terms for school price quintiles two and three are positive but not statistically significant, implying that the enrollment impacts in these schools are statistically indistinguishable from those estimated for public schools. In contrast, coefficients on the interaction terms for quintiles four and five are negative and statistically significant. These coefficients indicate that enrollment declines by 26.4 students (4.2%) per school in quintile four schools and by 55.8 students (9.4%) per school in quintile five schools.

Table 2 provides evidence that positive copper shocks cause private school prices to rise and enrollments to fall, with the magnitude of impacts increasing in baseline school price. These estimates reflect (scaled) measures of the causal impacts of changes in aggregate income to the extent that copper shocks only affect local educational markets through their impact on aggregate income conditional on year and municipality fixed effects.[34] Even if this exclusion restriction were to be satisfied, however, there are multiple mechanisms that may drive the link between aggregate income and school prices and enrollment levels. In order to gauge the extent to which rising aggregate income impacts school prices by shifting aggregate demand, I investigate the relative importance of alternative channels other than the demand for school quality that I have emphasized here. Seemingly, the most relevant alternative explanation for the link between aggregate income and school prices and enrollment levels is that rising aggregate income affects funding for public schools and so affects quality of and demand for public schools. Previous research indicates, however, that municipal tax revenues do not significantly affect

[33] The average private school enrollment is 587 students while the average public school enrollment is only 258 students.

[34] Technically, if this exclusion restriction is satisfied, then the causal impact of rising aggregate income can be identified based on split-sample instrumental variables estimates in which the reduced form is obtained using administrative educational data and the first stage is estimated with CASEN labor force survey data.

local educational expenditures in Chile (Auguste and Valenzuela, 2006). Consequently, changes in public school funding in response to copper shocks would likely come from national government sources. While vouchers are the primary mechanism through which central government funds are distributed, voucher values are determined nationally. The most important remaining source of local education funding to consider is the National Fund for Regional Development (FNDR), which distributes funds to select municipalities in order to increase their public school revenues. Appendix Table 3 provides evidence that FNDR funding is uncorrelated with municipality-specific copper shocks.

Rising aggregate income may also affect school prices and enrollment levels by changing the marginal cost curves that schools face. The largest component of a school's marginal cost is teacher incomes. In Appendix Table 4, I test for teacher income changes in response to positive copper shocks when the data is aggregated up to the regional level. I find that teacher incomes do not change significantly in response to positive copper shocks (coefficients are negative and I can rule out economically significant increases in teacher incomes).[35] While teacher income changes cannot be estimated at the school level based on available data, Appendix Table 5 reveals that average teacher experience and the school-level share of certified teachers do not change in response to the aggregate income shock. These characteristics are correlated with teacher incomes, and so these findings provide additional evidence that schools' cost curves are not changing significantly in response to aggregate income changes. Appendix Table 6 confirms that rental incomes also do not change dramatically in response to positive copper shocks (estimates are marginally significant and smaller in magnitude than estimated average changes in total income). This finding implies that schools' marginal costs are also not significantly affected by changes in the rental prices they face.

A remaining concern is that price and enrollment impacts estimated by including only those schools in operation in a given year may suffer from selection issues if rising

[35] This analysis is conducted at the regional level given that there are relatively few teachers observed in the CASEN survey data and so municipality-level estimates are quite noisy. As shown in Appendix Table 1, price and enrollment results are robust to estimating income elasticities at the regional level.

aggregate income induces changes in the number of schools in operation in particular municipalities. To assuage this concern, I show in Appendix Table 7 that positive copper shocks do not have a statistically significant impact on the number of schools in operation in a given municipality.[36]

Table 2 price and enrollment results imply that elite private voucher schools experience revenue losses in response to the aggregate income shock. For these schools' price and enrollment adjustments to be profitable, it must then be the case that reductions in school costs are larger. Case study evidence suggests that personnel costs constitute 87% of total non-capital costs in Chilean schools, and so changes in contracted teacher hours are the key determinant of whether schools can reduce costs in the short run (Ugarte and Williamson, 2012). In Appendix Table 8, Table 2 specifications are re-estimated for dependent variables that measure the total number of teacher contract hours per week at a given school and the total number of teachers employed at that school. The results imply that private schools reduce contracted hours by 57.1, on average. These reductions are driven by quintile four and five schools, which decrease contracted hours by 66.5 (8.0%) and 67.6 (8.0%), respectively. On average, private schools decrease the number of teachers employed by 1.6, while private schools in baseline price quintiles four and five reduce the number of employed teachers by 1.9 and 1.7, respectively. In contrast, the average public school increases the number of teachers employed by 0.4. Given that the average teacher contract is for 32 hours, these results indicate that private school cost reductions are explained primarily by reductions in the number of teachers in a school rather than in the number of contracted hours per teacher.[37] Under reasonable assumptions, the magnitude of estimated cost savings in elite private voucher schools implies that school profits are indeed increasing in response to the aggregate income shock.[38]

[36] I present results in both levels and logs and for the full sample as well as public school and private school subsamples. All estimates are statistically insignificant. In all specifications, lagged values of the dependent variable are included as controls to deal with spurious correlation driven by non-stationarity.

[37] Survey evidence suggests that teaching staff reductions are largest in non-core disciplines, such as art, music, and physical education.

[38] Private schools in baseline price quintiles four and five reduce revenues by 2.1% and 5.6%, respectively. However, personnel costs are estimated to fall by 8.0% for both sets of schools. If

6.2 Additional Robustness Checks

Tables 3 and 4 present results from a series of specifications that test the robustness of price and enrollment estimates to alternative definitions of the aggregate income shock, the sample of schools, and the unit of analysis. Columns (1)-(6) of Table 3 re-estimate the specifications from Table 2, but replace the copper shock measure C_{mt} with an average shock that is calculated as follows: First, Equation 10 is re-estimated separately using lagged copper values and twice-lagged copper values. Then, the average shock measure is constructed as the average of the current year and two lagged coefficients multiplied by the corresponding current and lagged copper prices. In practice, estimates in Columns (1)-(6) of Table 3 closely resemble Table 2 estimates (Table 3 estimates are slightly larger in magnitude). This suggests that overall price and enrollment impacts are relatively stable across specifications that vary in the time frame over which copper price changes are permitted to induce changes in local school prices and enrollment levels.

Columns (1) and (2) of Table 4 re-estimate Equations 12 and 13, but restrict the sample to exclude the Santiago metropolitan region. These specifications are presented in order to confirm that results are not driven by the country's capital (and largest city), where population density is high and so competition may be strongest. Coefficients are somewhat smaller in magnitude but remain highly significant here and in Columns (3) and (4), which replace price with enrollment as the dependent variable. Column (5) returns to the full sample and confirms that school-level enrollment impacts are robust to the inclusion of a control for the total number of enrolled students in municipality m in year t. Columns (6)-(8) of Table 4 estimate enrollment impacts at the municipality level to confirm that school-level estimates scale up to reflect significant changes in the municipality-level share of students attending public school. The specification employed is the following:

$$E_{mt} = \alpha + \beta * C_{mt} + \gamma_m + \lambda_t + \epsilon_{mt} \tag{16}$$

non-personnel costs fall by a similar percentage, then profits in quintile four and five schools rise if markups are lower than 281% and 43%, respectively.

Here, E_{mt} is an enrollment measure reflecting either the total number of enrolled students or the number of students enrolled in public schools in municipality m in year t. The remaining variables are as defined in Equation 11. In Column (6), the dependent variable is the total number of enrolled students in municipality m in year t. The reported estimate suggests that the impact of positive copper shocks on total enrollment is not statistically significant (although the magnitude of the coefficient is consistent with the average school-level enrollment increase estimated in Column (4) of Table 2). In Columns (7) and (8), the dependent variable is the number of students enrolled in public schools in municipality m in year t. The estimated coefficients reflect an increase in public school enrollment of 5.1 to 6.3%.

Finally, as a falsification exercise, Appendix Table 9 re-estimates Equation 12 with the full sample of private schools, but adds a series of lead terms ($C_{m,t+1}$, $C_{m,t+2}$, $C_{m,t+3}$, and $C_{m,t+4}$) to test whether future copper shocks are correlated with current school prices. None of the lead term coefficients presented in the table are statistically significant or large in magnitude.[39]

7 School Switching and Test Scores

7.1 School Switching Patterns

The evidence on average enrollment impacts reveals that aggregate income shocks induce movement from high-cost private schools to low-cost private schools and public schools. However, these results only tell half of the story. While rates of elite private school attendance at baseline are rising in students' socioeconomic status, the marginal students

[39] When both public and private schools are included in the sample, a subset of the $C_{m,t+1}$ coefficients are statistically significant. This is explained by the fact that the year fixed effects estimated in the public school only subsample are increasing in magnitude over time faster than the year fixed effects estimated in the private school only subsample. Consequently, when lead term coefficients are estimated in a pooled sample with a single set of year fixed effects, the year fixed effects take on intermediate values, and this generates a spurious correlation between school prices and future copper prices (which are increasing over time in expectation). Estimating specifications separately for public and private schools, or including separate year fixed effects for public and private schools in a pooled specification, returns statistically insignificant coefficients.

who move from elite private schools to public schools and low-quality private schools need not be representative of their classmates. To the extent that policymakers are concerned about inequality of educational opportunities, identifying what the characteristics of these marginal students tell us about changes in social stratification is of particular policy relevance.

To characterize those students most affected by positive copper shocks, I exploit the availability of unique student identifiers. These identifiers allow students to be tracked across years and schools. The key measure of student background that I use in the analysis is parental education, which has the advantage of being both time-invariant and highly correlated with other measures of student socioeconomic status such as household income.[40] Given that the distribution of educational attainment varies across regions of Chile and across time, I define parental education quartiles within municipality-year cells, with higher quartiles reflecting higher parental education.[41] To identify heterogeneous impacts based on parental education, I estimate the following specification:

$$
\begin{aligned}
Y_{ismt} = \alpha + \beta_1 * C_{mt} + \sum_{k=2}^{4} \beta_k * C_{mt} * X_{ik} \\
+ \sum_{k=2}^{4} \delta_k * X_{ik} + \gamma_m + \lambda_t + \epsilon_{ismt}
\end{aligned}
\tag{17}
$$

Here, Y_{ismt} reflects the outcome of interest for student i who attends school s and resides in municipality m in year t.[42] X_{ik} is an indicator for whether student i is in within-municipality parental education quartile k.

40 To construct the parental education measure, I take the mean of father's education and mother's education. I exclude respondents for whom one or both of these data points is missing. Results are robust to using only mother's education rather than average parental education.

41 Since parental education is only recorded in a subset of grade levels, I calculate the parental education quartile for the first year in which parental education is observed for a given student. Then, I assign this value to all observations for that student.

42 One important distinction between the school-level and student-level specifications is that the municipality corresponding to each observation in the student-level specifications reflects the municipality in which the students lives in the first year in which he appears in the administrative records. Consequently, the structure of fixed effects is complicated by the fact that school fixed effects no longer subsume municipality fixed effects. This distinction will be discussed in more depth in the subsequent analysis of student test score impacts.

Column (1) of Table 5 uses the specification outlined in Equation 17 to estimate the impact of positive copper shocks on the baseline school price quintile of the school that a student attends. Here, estimates indicate that students in parental education quartiles two and three experience the largest relative decreases in school quality, while students in quartile four experience smaller but still statistically significant relative declines.[43] Since middle and high parental education students attend more elite schools at baseline, these estimates may be explained by the fact that these students are simply more likely to end up in worse schools because they are more likely to be at the margin of attending schools that raise their prices significantly. To estimate average changes in school quality in year t conditional on being exposed to the *same* price increase, Column (2) adds lagged school by municipality fixed effects to the specification estimated in Column (1). These fixed effects are defined such that, for an observation in year t for individual i, a fixed effect is included for the school that individual i attended in year $t-1$ in combination with the municipality in which individual i lived in year $t-1$. In Column (2), the interaction terms for parental education quartiles two and three remain negative and statistically significant but are only 10% as large as Column (1) estimates, while the interaction term for quartile four is no longer statistically significant and is even smaller in magnitude (only 2% of the size of the corresponding Column (1) coefficient). These estimates suggest that the larger average price increases that middle parental education students face (relative to low parental education students) explain nearly all of their differential reduction in school quality.[44] In sum, middle parental education students are most affected by copper shocks

[43] I cannot reject that the sum of the coefficient on the positive copper shock level term plus any one of these interaction coefficients is equal to zero since the coefficient on the level term is positive but not statistically significant. Notably, the coefficients on the quartile indicator variables are all statistically significant at the 1% confidence level and are large in magnitude relative to the interaction terms. These coefficients imply that the average baseline price quintile of schools attended increases by 0.49, 0.53, and 0.70 with each additional one-point increase in parental education quartile. Consequently, even though Column (1) estimates suggest some degree of compression in the quality of schools attended by students from higher parental education quartiles relative to peers in the lowest parental education quartile, the magnitude of this compression is small. For instance, the implied decline in the school quality gap between quartile one students and quartile two students represents only 2.4% of the gap at baseline, while the decline between quartile one students and quartile four students represents a mere 0.6% of the baseline gap.

[44] Nonetheless, interaction coefficients remain negative in Column (2). This likely reflects the fact that, even conditional on attending the same school in year $t-1$, higher parental education quartile

because high parental education students are more willing to pay to attend elite schools even when tuitions rise, while low parental education students did not consider attending elite private schools in the first place.

Observed declines in school quality can be driven by either an increase in school "downgrading" or a decrease in school "upgrading," where downgrading is defined by whether a student attends a school in year t that has a lower baseline price quintile measure than the school she attended in year $t-1$ and upgrading is defined by whether a student attends a school in year t that has a higher baseline price quintile measure than the school she attended in year $t-1$. Columns (3) and (4) of Table 5 examine the impacts of positive copper shocks on indicator variables for whether a student upgrades and downgrades, respectively. Estimates imply that reductions in school quality are driven by a combination of reduced upgrading and increased downgrading (although estimated reductions in upgrading are somewhat larger in magnitude).

The fact that increased enrollment in public and low-cost private schools is driven by middle parental education students suggests that these same students will likely experience the largest relative declines in peer quality. To measure changes in peer quality and social stratification, I focus on changes in school-level mean parental education. Column (5) of Table 5 estimates Equation 17 and reveals that middle parental education students do indeed experience significant relative declines in school-level mean parental education. However, estimates also reveal that all subgroups experience overall increases in this measure. Column (6) re-estimates Equation 14 using mean parental education as the dependent variable, and the results reveal that all students experience an average increase in peer quality because all segments of the market experience non-decreasing average parental education (with parental education rising significantly in quintiles four and five).[45] Estimated impacts on school-level mean parental education suggest that changes

students are ex-ante more likely to attend higher baseline price schools in year t and so there is a greater margin for reducing school quality.

45 The possibility of a positive overall effect on mean parental education may initially appear puzzling given that the overall distribution of parental educational attainment is not shifting in response to positive copper shocks experienced by municipalities. However, an increase in mean parental educational attainment can be rationalized by strong ex-ante sorting across schools based on parental

in social stratification are small relative to average increases in peer quality in response to rising aggregate income. However, the fact that school quality is rising in peer quality but falling in classroom size suggests that school quality may nonetheless be falling in a subset of schools.

7.2 Test Scores

Understanding the relationship between changing enrollment patterns and academic performance in this environment has important implications for how we think about the welfare gains associated with aggregate income shocks in the presence of large-scale private school markets. I focus on student SIMCE test scores as a measure of educational achievement. SIMCE test scores have been used in many previous studies of the Chilean voucher system, including Hsieh and Urquiola (2006) and Gallego (2013). In the analysis, test scores are normalized so that the mean score is zero and the standard deviation is one within a particular grade level in a given year.

Column (1) of Table 6 estimates the overall impact of a positive copper shock on test scores using the following student-level specification:

$$T_{ismt} = \alpha + \beta * C_{mt} + \gamma_m + \lambda_t \tag{18}$$

Here, T_{ismt} is the normalized SIMCE score of student i who lives in municipality m and attends school s in year t. Column (1) estimates reveal that a positive copper shock increases average test scores in a municipality by 0.020 standard deviations. Column (2) examines whether test score impacts vary across parental education quartiles by re-estimating Equation 17 with normalized test scores as the dependent variable. Column (2) estimates reveal that test score gains are largest for students from the lowest and

education in combination with higher rates of school downgrading by those who would have been at the bottom of the parental education distribution within their counterfactual schools. As a simplified example, suppose that school A is composed half of students from parental education quartile one and half from quartile two, while school B is composed half of students from quartile three and half from quartile four. Then, if all school downgrading that occurs is driven by quartile three students moving from school B to school A, mean parental educational attainment in both schools will rise.

highest parental education quartiles. While impacts remain positive for all subgroups, the magnitude of test score increases is approximately 20% lower for students in parental education quartiles two and three relative to quartile one students. Students in parental education quartile four also experience (smaller) marginally significant test score declines relative to quartile one students.

To the extent that non-tuition spending rises in response to positive copper shocks, the results presented in Columns (1)-(2) reveal the combined impact of an increase in willingness to pay for schooling and an increase in household resources available for non-tuition expenditures, such as tutoring or even food, which may positively impact test scores. Consequently, the initial finding of larger test score gains for low and high parental education students could be entirely explained by differential changes in non-tuition expenditures for these subgroups, by changes in the schools that students attend, or by some combination of these two channels.[46]

To disentangle these two possible explanations for differential test score gains, I return to the specification from Column (2) of Table 5 that adds lagged school by municipality fixed effects. Estimates, presented in Column (3) of Table 6, provide evidence that the impact of positive copper shocks on school choice drives the differential test score gains that are estimated in Column (2) of that table. The coefficient on the positive copper shock level term of 0.025 is statistically significant at the 1% confidence level and similar in magnitude to the corresponding coefficient in Column (2). However, the coefficient on the parental education quartile two interaction term is no longer statistically significant, while the interaction terms for quartiles three and four are now positive and statistically significant at the 10% and 5% levels, respectively. Column (4) replaces lagged school by municipality fixed effects with school by municipality fixed effects for the school attended in year t and estimated coefficients are similar.

The school that an individual attends is presumably correlated with unobservable

[46] In addition to increases in non-tuition spending, there may be differential behavioral changes (for instance, changes in parental attentiveness) in response to the income shock that affect academic performance. If non-tuition expenditures or behavioral changes were the primary explanation, however, there would have to be some underlying non-monotonicity in order to explain the larger impacts at the top and bottom of the parental education distribution.

characteristics, such as academic motivation, and this implies that estimates in Columns (3) and (4) cannot be interpreted as revealing the purely causal impact of copper shocks on test scores that is independent of school choice. However, it is difficult to reconcile these findings with the hypothesis that larger test score gains for low and high parental education students are driven primarily by differential returns to income available for expenditures other than school costs.[47] At the extreme, if smaller gains for middle parental education students are driven entirely by those who attend worse schools as a consequence of rising aggregate income, then the implied change in test scores for these downgraded students is equal to *negative* 0.22 standard deviations.[48] This decline in academic performance dwarfs in magnitude the average gains experienced by those students who do not downgrade. Although I cannot isolate the subpopulation of students induced to downgrade by the copper shock, I calculate that downgrading is associated with an average test score decline of 0.15 standard deviations within the full population. This figure is similar in magnitude to the decline imputed for students induced to downgrade by the copper shock and provides additional support for the assertion that relative test score declines for middle parental education students are driven primarily by those students who downgrade.

[47] Appendix Table 10 confirms that evaluations of teacher performance do not change in response to aggregate income shocks. This suggests that changes in teacher effort also do not appear to play a central role in explaining observed test score changes. To provide additional evidence of differential effects on academic performance, Appendix Table 11 shows impacts on students' grade point averages. Changes in students' grade point averages mirror estimated test score impacts.

[48] This figure is based on the following calculation: relative to low parental education students, those in parental education quartiles two and three attend a school that is 0.014 baseline price quintile points lower than the school they would have attended absent the shock (based on Column (1) of Table 5). Relative to low parental education students, these quartile two and three students experience a corresponding test score increase in response to a positive copper shock that is 0.0034 standard deviations smaller. Dividing the second figure by the first, I estimate that a one-point downgrade in the baseline school price quintile distribution is associated with a 0.24 standard deviation decrease in test scores. Since this estimated decrease is relative to an average increase of 0.02 standard deviations for low parental education students, the implied aggregate test score reduction for downgraded students is 0.22 standard deviations.

8 Structural Estimation

The model presented in Section 2 suggested that, under certain conditions, the residual demand curves faced by private school administrators are sufficiently inelastic that their profit-maximizing response to increased demand for private schooling is to raise prices so substantially that equilibrium private school enrollment falls. The empirical results presented in Sections 6 and 7 provide evidence that these market conditions are satisfied in the Chilean setting and that students' academic performance and educational opportunities are significantly affected by private schools' pricing behavior. In this section, I estimate an extended version of the model in order to characterize average willingness to pay for particular school attributes and heterogeneity in willingness to pay for school attributes based on individual characteristics. Simulated school price and enrollment responses to rising aggregate income based on these parameter estimates then allow me to quantify the relative importance of the market power and classroom size preferences channels outlined in Section 2.

The estimation approach uses both micro and aggregate data in order to produce parameter values that match schools' predicted market shares to observed ones. This approach was introduced in Berry et al. (2004). The algorithm most closely follows Gallego and Hernando (2009), which first applied the Berry et al. (2004) approach to the Chilean school choice decision. The equation for household utility that is used to generate parameter estimates is the following:

$$u_{ijt} = \frac{q_{jt}}{n_{jt}}\beta_{i1} - p_{jt}\beta_{i2} + \sum_{k=3}^{4} X_{jkt}\beta_{ik} + \lambda_{ijt}\sigma_i + \gamma_j + \xi_{jt} + \varepsilon_{ijt} \tag{19}$$

Here, u_{ijt} is the utility of student i who chooses school j at time t. Average classroom size and price are represented by $\frac{q_{jt}}{n_{jt}}$ and p_{jt}, respectively. X_{jkt} are additional observable characteristics k of school j at time t, and λ_{ijt} is the distance from the center of individual i's home municipality to school j. γ_j is a time-invariant school fixed effect and

ξ_{jt} represents unobservable time-varying school-specific characteristics.[49] ε_{ijt} is a random preference shock (error term) that is assumed to have a Type I extreme value distribution. This expression nests Equation 1 where school quality is expressed as a scalar s and individuals differ in their willingness to pay for school quality. The utility function here, however, exploits the availability of panel data in which case time-invariant school fixed effects are identified and baseline school quality is subsumed into the school fixed effect.

To analyze the conditions under which Equation 19 is identified, it is informative to characterize heterogeneity in preferences for school attributes based on observable individual characteristics. In doing so, I follow the notation from Berry et al. (2004) and rewrite β_{ik} as $\bar{\beta}_k + \sum_r d_{ir}\beta^o_{kr}$.[50] Here, r indexes individual characteristics and β^o_{kr} captures interactions between observed individual and school characteristics. I can then decompose Equation 19 into (1) a component characterizing the mean utility associated with a given school in a particular year (δ_{jt}), and (2) a component that captures preference heterogeneity:

$$u_{ijt} = \delta_{jt} + \sum_{kr} X_{jkt} d_{ir}\beta^o_{kr} + \lambda_{ijt}\bar{\sigma} + \sum_{r} \lambda_{ijt} d_{ir}\sigma^o_r + \varepsilon_{ijt} \tag{20}$$

where

$$\delta_{jt} = \textstyle\sum_k X_{jkt}\bar{\beta}_k + \gamma_j + \xi_{jt} \tag{21}$$

Equation 21 highlights the fact that $\bar{\beta}_k$ and ξ_{jt} cannot be separately identified without further parametric assumptions. Consequently, I follow the two-step procedure employed in Berry et al. (2004). The first step estimates the δ, β^o, $\bar{\sigma}$, and σ^o parameters without imposing any additional structure on the ξ_{jt} coefficients. This ensures that δ, β^o, $\bar{\sigma}$, and σ^o are consistently estimated. The second step of the algorithm requires a set of instruments Z_{jt} that satisfy the assumption $E[\xi_{jt}|Z_{jt}] = 0$ and so allow the $\bar{\beta}_k$ and ξ_{jt} coefficients to be consistently estimated.

In estimating Equation 19, I include fourth grade students in the Santiago metropoli-

[49] For a detailed discussion of the interpretation of parameters in a model with product-specific dummy variables, see Nevo (2000).

[50] I can similarly rewrite σ_i as $\bar{\sigma} + \sum_r d_{ir}\sigma^o_r$.

tan region in the years 2006-2009. I choose fourth graders because geographic coordinates are more widely available for primary schools than for secondary schools. The sample is restricted to the Santiago metropolitan region in the years 2006-2009 given the computational intensity of the estimation algorithm.[51] In addition to average classroom size and price, the included school characteristics (i.e., the set of X_{jkt}) are indicators for whether school j is public or a low-priced private school (in baseline price quintile two or three). The included individual characteristics (i.e., the set of d_{ir}) are: parental education, parental income, imputed income change in response to a positive copper shock, and whether parents expect their child to graduate from college. The imputed income response to copper shock measure is constructed as 2.9% of income (the average change in income estimated in Table 1) multiplied by the copper shock corresponding to a student's home municipality in a given year. Heterogeneity in willingness to pay based on this imputed income shock is a key parameter of interest given the fact that cross-sectional differences in household income are not exogenous. As a result, changes in preferences based on cross-sectional income differences may be quite different from the changes induced when an individual experiences an income shock (and it is this within-individual income variation that is of interest for the simulations). In the analysis, observations are dropped for students with missing values for any of the individual variables and for schools with missing SIMCE survey responses (which are used to determine school attributes).

The maximum likelihood approach I employ in step one of the algorithm searches for the δ, β^o, $\bar{\sigma}$, and σ^o parameter values that maximize the sum of the predicted log probability that each student attends the school they attend in practice. Initially, a guess is chosen for the β^o, $\bar{\sigma}$, and σ^o parameter values. Then, the δ that maximizes the

[51] This constitutes a subset of the full sample analyzed in Section 6-7. The Santiago metropolitan region serves 41% of Chilean students and includes 24% of schools and 20% of municipalities. Appendix Table 12 re-estimates Table 2 specifications for the Santiago metropolitan region subsample. Price and enrollment impacts are broadly similar to full sample estimates: private school prices increase, private school enrollments decline, and both price and enrollment impacts are largest at elite private schools. Initial estimates identify price decreases in public schools in response to copper shocks, but this is driven entirely by the pooled construction of year fixed effects. Public school price impacts are insignificant (with a coefficient of 0.001) when I estimate price changes separately for public and private schools or include separate year fixed effects for public and private schools in a pooled specification.

log likelihood function is calculated conditional on the chosen β^o, $\bar{\sigma}$, and σ^o values.[52] Calculating the maximum achievable log likelihood for each choice of the β^o, $\bar{\sigma}$, and σ^o parameters, the algorithm searches for the β^o, $\bar{\sigma}$, and σ^o values that globally maximize the log likelihood function. The advantage of this procedure over one that jointly searches for the likelihood-maximizing values of δ, β^o, $\bar{\sigma}$, and σ^o is that it eases the computational burden by reducing the dimensionality of the problem (Nevo, 2000). In estimating parameters, I use a gradient-based algorithm in combination with the Nelder-Mead Simplex method to improve the speed with which likelihood-maximizing parameter values are found. This gradient-based approach requires that the gradient of the log likelihood function be calculated explicitly, and I make use of the expressions for the gradient derived in the Computational Appendix of Gallego and Hernando (2009).

Table 7 presents the likelihood-maximizing β^o, $\bar{\sigma}$, and σ^o parameter values along with estimates of β_1 and β_2.[53] Coefficients in Table 7 reflect the change in utility associated with changes in school and individual characteristics. Estimates related to preferences for reduced classroom size suggest that wealthier and more educated parents actually value classroom size slightly less than their peers, while those experiencing a positive income shock increase their valuation of small class size. Higher income parents, better educated parents, those with high expectations for their children, and those experiencing positive income shocks all appear to care less about school prices than their peers. This reduced price sensitivity and increased demand for small class size by households experiencing a positive income shock will play an important role in determining school price responses to rising aggregate income in the simulation presented below.

In the second step of the structural estimation algorithm, I use an instrumental variables approach to deal with the concern that time-varying observable school characteristics (price and average classroom size) may be correlated with time-varying unobservable

52 Berry et al. (1995) proves that one can solve for δ recursively based on the equation $\delta^{n+1} = \delta^n + \ln \frac{\Gamma_{jt}}{\hat{\Gamma}_{jt}(\beta^o, \delta, \bar{\sigma}, \sigma^o)}$, where Γ_{jt} are schools' observed market shares, $\hat{\Gamma}_{jt}(\beta^o, \delta, \bar{\sigma}, \sigma^o)$ are predicted market shares, and predicted probabilities are re-calculated based on each updated value of δ.

53 Note that main effects cannot be estimated for time-invariant school characteristics, such as whether a school is public, given that these characteristics are subsumed by the included school fixed effects.

school characteristics, ξ_{jt}. The instrument set includes the 20th, 40th, 60th, and 80th percentiles of the income distribution in municipality j and in year t. In addition, the set of instruments for p_{jt} and $\frac{q_{jt}}{n_{jt}}$ includes the lagged number of classrooms in grade three and in grade four in school j. The set of income distribution instruments is valid as long as each student's school choice conditional on own income is only affected by changes to other households' income through the impact of these changes on school prices and class size. The strategy of using market-level demand characteristics to instrument for endogenous product attributes is discussed in Berry and Haile (2010) and is also applied in Gentzkow and Shapiro (2010). Lagged number of third and fourth grade classrooms significantly predicts average classroom size because changes to the number of school classrooms are infrequent (and costly) and the number of available classrooms is a key determinant of realized classroom size. Lagged numbers of classrooms are valid instruments if they affect school preferences solely through their impact on classroom size and prices. This condition is satisfied if ξ_{jt} are drawn independently in each year.[54]

Based on the resulting estimates of the $\bar{\beta}_k$ parameters, and the β^o, $\bar{\sigma}$, and σ^o values estimated in step one, I calculate the average elasticity of demand with respect to price faced by private schools. To do so, I must incorporate the fact that a school price increase leads to reduced enrollment which mechanically reduces classroom size. Since average willingness to pay is falling in classroom size, this implies that a price elasticity estimated without accounting for endogenously changing classroom size will represent an overestimate (in magnitude) of the true residual demand elasticity that schools face. To incorporate class size changes into the derivation of residual demand elasticity, I employ the following formula:

$$\hat{\nu}_{jp} = \frac{\sum_{i=1}^{N}[\hat{P}_{ij}(1-\hat{P}_{ij})(\bar{\beta}_2 + \sum_r d_{ir}\beta^o_{r2})]}{1 - \frac{1}{n_j}\sum_{i=1}^{N}[\hat{P}_{ij}(1-\hat{P}_{ij})(\bar{\beta}_1 + \sum_r d_{ir}\beta^o_{r1})]} * \frac{p_j}{q_j} \tag{22}$$

[54] If ξ_{jt} are correlated over time, then most plausible sources of omitted variables bias would lead to elasticities that are too small in magnitude. However, the estimated parameters reflecting heterogeneity in preferences for price and classroom size based on individual characteristics would not be affected.

Here, $\hat{\nu}_{jp}$ is the true elasticity of demand with respect to price that school j faces, N is the total number of students in the market, $\hat{P}_{ij}$ is the estimated probability that student i attends school j, and n_j is the number of classrooms in school j. Based on this expression, I estimate that the average residual demand elasticity faced by schools in the market is -2.66.

Given estimates of the elasticity of residual demand faced by schools, I apply the standard monopolist markup formula to back out schools' marginal costs at baseline enrollment levels.[55] With a full set of estimated demand parameters and school marginal costs, I am prepared to simulate school price and enrollment responses to an aggregate income shock. To do so, I impose a uniform (2.9%) copper shock-induced income change. This imputed income change is set equal to the change in mean income estimated in Table 1 in response to a copper shock of one standard deviation. Consequently, if I have correctly modeled the structure of the market, I should simulate price and enrollment changes that mirror the reduced-form estimates presented in Table 2.

To estimate equilibrium changes in private school prices and enrollment, I follow an iterative procedure. In each iteration, I construct for each school a grid of five prices (centered around the price chosen in the previous iteration) with corresponding changes in enrollment and marginal cost. Enrollment changes are calculated by linearizing the logit function. To compute marginal cost, I estimate cross-sectionally the relationship between marginal cost and class size within each baseline school price quintile bin. Then, I predict marginal cost at a given enrollment level as a function of a school's baseline marginal cost and this slope. Given expressions for school price, enrollment, and average variable cost at each point on the grid, each school is assigned the grid point corresponding to the profit-maximizing price.[56] After each iteration in which all schools have chosen profit-maximizing prices, students are re-sorted across schools based on these new prices and

[55] The markup formula is $MC = p(1 + \frac{1}{\epsilon})$, where ϵ is the elasticity of residual demand.

[56] Schools are not permitted to continue lowering prices after enrollment has risen by 25% or to continue increasing prices after enrollment has fallen by 25% to reflect the fact that such dramatic enrollment changes are not observed in practice and would likely impose additional costs on a school that the structural model does not incorporate.

endogenously determined average classroom sizes.[57] This process then repeats until I arrive at a Nash equilibrium in which all schools are choosing the local best response to all other schools' prices.

Table 8 presents the results from this simulation algorithm for the 783 schools in the Santiago metropolitan region in 2006. The price and enrollment results provide a reasonable approximation of reduced-form estimates. I simulate a 6.09% increase in public school enrollment relative to a 5.25% reduced-form estimate. As observed in the data, I find that price increases and enrollment declines are driven by quintile four and five schools. The simulation results underpredict price changes at the most elite schools and overpredict corresponding enrollment declines. Both schools' market power and parents' class size preferences are necessary to generate these simulation results. Changes in residual demand elasticity are driven by both reduced price sensitivity and strengthened class size preferences. Moreover, class size preferences prevent private schools from pursuing an alternative strategy of leaving prices relatively intact and significantly expanding enrollment. In additional simulations, I confirm that elite private schools do just this in equilibrium when class size preferences are turned off.

9 Conclusions

The proposed model highlights two key mechanisms that could theoretically cause private school enrollment to decline in response to a positive aggregate income shock. First, schools appear to have market power and so may raise prices more than they would in a competitive market. Second, school quality is a normal good and classroom size is an important component of school quality. Simulations based on the model reveal that a broad range of parameter values can cause private school enrollment declines.

[57] I bound average classroom size impacts so that there are no additional quality gains below 10 students and no additional quality losses above 35 students. This assumption is justified by the fact that preferences for smaller class size are estimated based on changes across intermediate class size values. Moreover, imposing an upper bound on quality gains from class size reductions ensures that schools are not incentivized to reduce enrollment to close to zero in order to charge increasingly higher prices (as this is not observed in the data).

To investigate whether the conditions under which aggregate income shocks would cause private school enrollment declines are satisfied in practice, I construct an aggregate income shock and study private school price and enrollment responses in Chile. I find that private school enrollment shares fall when aggregate income rises in this setting. The analysis reveals that private school price increases and enrollment declines are driven by those schools that were most expensive at baseline. Public schools and low-cost private schools do not adjust prices significantly; they instead expand enrollment to absorb those additional students who would have attended high-cost private voucher schools absent the rise in aggregate income. To determine whether market power or class size preferences drive estimated enrollment impacts, I structurally estimate an extended version of the model using Chilean educational data. Simulations based on structural estimates imply that both market power and parents' preferences for smaller class size play critical roles in making it profitable for schools to raise prices to such an extent that enrollment declines.

In the analysis, I find that middle-income students benefit least from increases in aggregate income. These students are most likely to attend worse schools as a result of the rise in aggregate income, and this result is driven by those middle-income students who would have attended elite private schools absent the shock to aggregate income. Correspondingly, these same students experience the smallest test score gains. I present suggestive evidence that this is causally related to higher rates of school downgrading for this subpopulation.

It is encouraging that estimates show relatively large test score gains for disadvantaged students when aggregate income rises. However, the observed changes in enrollment patterns suggest that rising incomes may widen the gap between the highest socioeconomic status students and everyone else. In any case, a better understanding of school responses to aggregate changes in willingness to pay can inform those interested in designing policies to improve educational opportunities for disadvantaged students, among other goals.

References

Agency of Education Quality (June 2013). What is the SIMCE? http://www.agenciaeducacion.cl/simce/que-es-el-simce/.

Andrabi, T., J. Das, and A. Khwaja (2013). Report Cards: The Impact of Providing School and Child Test Scores on Educational Markets. Working Paper.

Angrist, J., E. Bettinger, E. Bloom, E. King, and M. Kremer (2002). Vouchers for Private Schooling in Colombia: Evidence from a Randomized Natural Experiment. *American Economic Review 92*(5), 1535–1558.

Auguste, S. and J. P. Valenzuela (2006). Is It Just Cream Skimming? School Vouchers in Chile. Mimeo.

Berry, S. and P. Haile (2010). Nonparametric Identification of Multinomial Choice Demand Models with Heterogeneous Consumers. Cowles Foundation Discussion Paper No. 1718.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile Prices in Market Equilibrium. *Econometrica 63*, 841–890.

Berry, S., J. Levinsohn, and A. Pakes (2004). Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market. *Journal of Political Economy 112*(1), 68–105.

Bettinger, E., M. Kremer, and J. E. Saavedra (2010). Are Educational Vouchers Only Redistributive? *Economic Journal 120*, F204–F228.

Bravo, D., S. Mukhopadhyay, and P. E. Todd (2010). Effects of School Reform on Education and Labor Market Performance: Evidence from Chile's Universal Voucher System. *Quantitative Economics 1*(1), 47–95.

Chilean Customs (2013). Estadísticas Regionales. http://www.aduana.cl/estadisticas-regionales/aduana/2007-04-16/170042.html.

Drago, J. L. and R. D. Paredes (August 2011). The Quality Gap in Chile's Education System. *CEPAL Review 104*, 162–174.

Elacqua, G. (July 2009). For-Profit Schooling and the Politics of Education Reform in Chile: When Ideology Trumps Evidence. CPCE Working Paper 5.

Epple, D. and R. E. Romano (1998). Competition Between Private and Public Schools, Vouchers, and Peer-Group Effects. *American Economic Review 88*(1), 33–62.

Ferreyra, M. M. (2007). Estimating the Effects of Private School Vouchers in Multidistrict Economies. *American Economic Review 97*(3), 789–817.

Gallego, F. (2013). When Does Inter-School Competition Matter? Evidence from the Chilean "Voucher" System. *B.E. Journal of Economic Analysis & Policy 13*(2), 525–562.

Gallego, F. and A. Hernando (2009). School Choice in Chile: Looking at the Demand Side. Working Paper 356, Catholic University of Chile.

Gentzkow, M. and J. Shapiro (2010). What Drives Media Slant? Evidence from U.S. Daily Newspapers. *Econometrica 78*(1), 35–71.

Howell, W. G. and P. E. Peterson (2006). *The Education Gap: Vouchers and Urban Schools*. Washington, DC: The Brookings Institution.

Hsieh, C.-T. and M. Urquiola (2006). The Effects of Generalized School Choice on Achievement and Stratification: Evidence from Chile's Voucher Program. *Journal of Public Economics 90*, 1477–1503.

King, E., L. Rawlings, M. Gutierrez, C. Pardo, and C. Torres (1997). Colombia's Targeted Education Voucher Program: Features, Coverage, and Participation. Working Paper 3, Development Economics Research Group, World Bank.

Krueger, A. and P. Zhu (2004). Another Look at the New York City School Voucher Experiment. *American Behavioral Scientist*, 658–698.

McEwan, P. (August 2013). Improving Learning in Primary Schools of Developing Countries: A Meta-Analysis of Randomized Experiments. Working Paper.

Ministry of Social Development (June 2013). CASEN Survey: Description and Objectives. http://www.ministeriodesarrollosocial.gob.cl/casen/en/descripcion_obj.html.

Mizala, A. and P. Romaguera (2000). School Performance and Choice: The Chilean Experience. *Journal of Human Resources 35*(2), 392–417.

Mizala, A. and M. Urquiola (2013). School Markets: The Impact of Information Approximating Schools' Effectiveness. *Journal of Development Economics 103*, 313–335.

Muralidharan, K. and V. Sundararaman (2013). The Aggregate Effect of School Choice: Evidence from a Two-Stage Experiment in India. Working Paper.

Nechyba, T. J. (2000). Mobility, Targeting, and Private-School Vouchers. *American Economic Review 90*(1), 130–146.

Nevo, A. (2000). A Practicioner's Guide to Estimation of Random-Coefficients Logit Models of Demand. *Journal of Economics and Management Strategy 9*(4), 513–548.

Pica, C. and S. Wisnefski (2012, October 19). Chile Peso Seen At Outer End of Equilibrium Range. http://online.wsj.com/article/SB10000872396390444734804578064821736405316.html. The Wall Street Journal.

Rouse, C. E. (1998). Private School Vouchers and Student Achievement: An Evaluation of the Milwaukee Parental Choice Program. *Quarterly Journal of Economics 13*, 553–602.

Sapelli, C. and B. Vial (2002). The Performance of Private and Public Schools in the Chilean Voucher System. *Cuadernos de Economia-Latin American Journal of Economics 39*(118), 423–454.

Solimano, A. and A. Torche (August 2008). La Distribución del Ingreso en Chile 1987-2006: Análisis y Consideraciones de Política. Central Bank of Chile Working Paper.

Tooley, J., P. Dixon, and S. Gomathi (2007). Private Schools and the Millennium Development Goal of Universal Primary Education: A Census and Comparative Survey in Hyderabad, India. *Oxford Review of Education 33*(5), 539–560.

Ugarte, G. and B. Williamson (2012). Hacia La Medición del Costo de una Educación de Calidad. Centro de Estudios Mineduc (Chilean Ministry of Education).

Urquiola, M. and E. Verhoogen (2009). Class-Size Caps, Sorting, and the Regression-Discontinuity Design. *American Economic Review 99*(1), 179–215.

U.S. Geological Survey (July 11, 2013). Copper: Statistics and Information. http://minerals.usgs.gov/minerals/pubs/commodity/copper/.

World Bank (November 2013). Exports of Goods and Services (% of GDP). http://data.worldbank.org/indicator/NE.EXP.GNFS.ZS.

Table 1. Copper Shocks and Income

	Log Mean Income	Log 50th Percentile Income	Log 60th Percentile Income	Log 70th Percentile Income	Log 80th Percentile Income	Log 90th Percentile Income
	(1)	(2)	(3)	(4)	(5)	(6)
Normalized Copper Shock	0.0292***	0.0581***	0.0259***	0.0182***	0.0180**	0.0245***
	(0.0083)	(0.0222)	(0.0075)	(0.0055)	(0.0071)	(0.0074)
Year Fixed Effects	X	X	X	X	X	X
Municipality Fixed Effects	X	X	X	X	X	X
Mean of Dependent Variable	12.099	10.375	11.849	12.210	12.507	12.947
	[0.459]	[3.383]	[1.182]	[0.461]	[0.476]	[0.494]
Observations	1078	1078	1078	1078	1078	1078
Sample	All (Municipality-level)		All (Municipality-level)		All (Municipality-level)	

Notes

Normalized Copper Shock is defined as the product of the normalized municipality-specific elasticity of income with respect to copper prices and the log copper price (denominated in 1998 USD). Regressions are clustered at the municipality level (there are 270 municipalities). * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table 2. School Price and Enrollment Responses to Copper Shocks

	Log School Price	Log School Price	Log School Price	Number of Students	Number of Students	Number of Students
	(1)	(2)	(3)	(4)	(5)	(6)
Normalized Copper Shock	0.0042***	-0.0003	-0.0004	6.4***	14.4***	12.9***
	(0.0008)	(0.0002)	(0.0005)	(1.9)	(3.3)	(3.0)
Normalized Copper Shock* Private School		0.0274***			-48.5***	
		(0.0064)			(12.2)	
Normalized Copper Shock* Baseline School Price in Quintile 2			-0.0040			66.0
			(0.0048)			(72.8)
Normalized Copper Shock* Baseline School Price in Quintile 3			0.0084*			8.8
			(0.0048)			(11.4)
Normalized Copper Shock* Baseline School Price in Quintile 4			0.0224***			-39.3***
			(0.0074)			(11.3)
Normalized Copper Shock* Baseline School Price in Quintile 5			0.0421***			-68.7***
			(0.0076)			(11.4)
Year Fixed Effects	X	X	X	X	X	X
School Fixed Effects	X	X	X	X	X	X
Mean of Dependent Variable	10.643			342.1		
	[0.279]			[407.5]		
Observations	52438	52438	52033	52438	52438	52033
Sample	All (School-level)			All (School-level)		

Notes

Normalized Copper Shock is defined as the product of the normalized municipality-specific elasticity of income with respect to copper prices and the log copper price (denominated in 1998 USD). Regressions are clustered at the municipality level (there are 270 municipalities). * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table 3. School Price and Enrollment Robustness Specifications I

	Log School Price	Log School Price	Log School Price	Number of Students	Number of Students	Number of Students
	(1)	(2)	(3)	(4)	(5)	(6)
Average Copper Shock	0.0057***	-0.0002	-0.0003	8.0***	18.9***	17.2***
	(0.0010)	(0.0003)	(0.0003)	(1.5)	(2.5)	(2.4)
Average Copper Shock* Private School		0.0315***			-58.6***	
		(0.0053)			(10.0)	
Average Copper Shock* Baseline School Price in Quintile 2			-0.0021			67.5
			(0.0053)			(68.4)
Average Copper Shock* Baseline School Price in Quintile 3			0.0062*			-1.8
			(0.0036)			(9.3)
Average Copper Shock* Baseline School Price in Quintile 4			0.0266***			-45.8***
			(0.0054)			(9.3)
Average Copper Shock* Baseline School Price in Quintile 5			0.0495***			-83.3***
			(0.0091)			(10.2)
Year Fixed Effects	X	X	X	X	X	X
School Fixed Effects	X	X	X	X	X	X
Mean of Dependent Variable	10.643			342.1		
	[0.279]			[407.5]		
Observations	52438	52438	52033	52438	52438	52033
Sample		All (School-level)			All (School-level)	

Notes

Average Copper Shock is defined as the mean of current, lagged, and twice lagged copper shocks (constructed as described in Table 1). Regressions are clustered at the municipality level (there are 270 municipalities). * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table 4. School Price and Enrollment Robustness Specifications II

	Log School Price		Number of Students		Number of Students		Number of Public School Students	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Normalized Copper Shock	0.0026***	0.0001	5.4***	11.5***	11.8***	91.6	262.4***	212.6***
	(0.0007)	(0.0002)	(1.4)	(3.0)	(2.9)	(72.6)	(56.8)	(44.5)
Normalized Copper Shock* Private School		0.0205***		-51.5***	-60.7***			
		(0.0041)		(11.1)	(14.6)			
Year Fixed Effects	X	X	X	X	X	X	X	X
School Fixed Effects	X	X	X	X	X			
Control for Total Students					X			X
Municipality Fixed Effects						X	X	X
Mean of Dependent Variable	10.611		256.1			10675.0	4161.7	
	[0.258]		[338.0]			[15083.9]	[5546.1]	
Observations	39907	39907	39907	39907	52438	2160	2160	2160
Sample	Excluding metropolitan region		Excluding metropolitan region		All	Municipality-level		

Notes

Regressions are clustered at the municipality level (there are 270 municipalities). * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table 5. Copper Shocks and Student School Sorting Patterns

	Baseline School Price Quintile		Upgrade	Downgrade	School-Level Mean Parental Education	
	(1)	(2)	(3)	(4)	(5)	(6)
Normalized Copper Shock	0.0123	-0.0004	-0.0014	0.0007	0.0460***	0.0021
	(0.0121)	(0.0035)	(0.0008)	(0.0005)	(0.0102)	(0.0158)
Normalized Copper Shock* Within-municipality Parental Education Quartile 2	-0.0115***	-0.0012***	-0.00035***	0.00010***	-0.0093***	
	(0.0017)	(0.0002)	(0.00007)	(0.00002)	(0.0026)	
Normalized Copper Shock* Within-municipality Parental Education Quartile 3	-0.0164***	-0.0013***	-0.00054***	0.00014***	-0.0130***	
	(0.0033)	(0.0003)	(0.00010)	(0.00004)	(0.0046)	
Normalized Copper Shock* Within-municipality Parental Education Quartile 4	-0.0101**	-0.0002	-0.00041***	0.00019**	-0.0065	
	(0.0044)	(0.0005)	(0.00012)	(0.00007)	(0.0058)	
Within-municipality Parental Education Quartile 2	0.485***	0.051***	0.0145***	-0.0019***	0.702***	
	(0.017)	(0.002)	(0.0006)	(0.0002)	(0.022)	
Within-municipality Parental Education Quartile 3	1.022***	0.103***	0.0287***	-0.0039***	1.374***	
	(0.028)	(0.003)	(0.0009)	(0.0003)	(0.035)	
Within-municipality Parental Education Quartile 4	1.722***	0.162***	0.0437***	-0.0077***	2.275***	
	(0.037)	(0.004)	(0.0009)	(0.0004)	(0.052)	
Normalized Copper Shock* Baseline School Price in Quintile 2						0.0781
						(0.0855)
Normalized Copper Shock* Baseline School Price in Quintile 3						0.0058
						(0.0272)
Normalized Copper Shock* Baseline School Price in Quintile 4						0.1180***
						(0.0449)
Normalized Copper Shock* Baseline School Price in Quintile 5						0.1510***
						(0.0378)
Year Fixed Effects	X	X	X	X	X	X
Municipality Fixed Effects	X				X	
Lagged School-Municipality Fixed Effects		X	X	X		
School Fixed Effects						X
Observations	11,908,109	11,168,582	10,933,347	10,933,347	11,908,109	51,633
Sample	All (Student-level)		All (Student-level)		All (Student-level)	All (School-level)

Notes

Upgrading is defined by whether a student attends a school in year `t' that has a higher baseline price quintile measure than the school she attended in year `t-1' and downgrading is defined by whether a student attends a school in year `t' that has a lower baseline price quintile measure than the school she attended in year `t-1'. Lagged school-municipality fixed effects control for the school that an individual attended in the previous year in combination with the municipality in which the individual lived in the previous year. Regressions are clustered at the municipality level (there are 270 municipalities). * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table 6. Copper Shocks and Student Test Scores

	Normalized SIMCE Score		Normalized SIMCE Score	
	(1)	(2)	(3)	(4)
Normalized Copper Shock	0.0198***	0.0183***	0.0254***	0.0265***
	(0.0046)	(0.0051)	(0.0057)	(0.0059)
Normalized Copper Shock* Within-municipality Parental Education Quartile 2		-0.0031***	-0.0002	-0.00005
		(0.0007)	(0.0003)	(0.0003)
Normalized Copper Shock* Within-municipality Parental Education Quartile 3		-0.0037***	0.0009*	0.0012**
		(0.0010)	(0.0005)	(0.0005)
Normalized Copper Shock* Within-municipality Parental Education Quartile 4		-0.0024*	0.0023**	0.0025***
		(0.0013)	(0.0009)	(0.0009)
Within-municipality Parental Education Quartile 2		0.283***	0.157***	0.154***
		(0.006)	(0.002)	(0.002)
Within-municipality Parental Education Quartile 3		0.517***	0.273***	0.267***
		(0.007)	(0.004)	(0.004)
Within-municipality Parental Education Quartile 4		0.820***	0.420***	0.408***
		(0.009)	(0.008)	(0.008)
Year Fixed Effects	X	X	X	X
Municipality Fixed Effects	X	X		
Lagged School-Municipality Fixed Effects			X	
School-Municipality Fixed Effects				X
Observations	2,581,999	2,238,892	2,235,769	2,238,892
Sample	All (Student-level)		All (Student-level)	

Notes

Regressions are clustered at the municipality level (there are 270 municipalities). * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table 7. Preferences for School Attributes

School Characteristic	Classroom Size (unit: 10 students)	Price (unit: 10,000 CLP=20 USD)	Public School (0/1 Indicator)	Low-Price Private School (0/1 Indicator)	Distance (unit: 10km)
	(1)	(2)	(3)	(4)	(5)
Main Effect	-1.7020	-4.4590	-	-	-10.5561
	(0.2730)	(0.9740)	-	-	(0.3113)
Heterogeneity by:					
Baseline Parental Income (unit: 100,000 CLP=200 USD)	0.0059	0.0691	-0.0784	-0.0462	0.0507
	(0.0017)	(0.0011)	(0.0043)	(0.0040)	(0.0039)
Copper Shock Income (unit: 100,000 CLP=200 USD)	-0.0458	0.0231	-0.0632	-0.1095	-0.5402
	(0.0072)	(0.0044)	(0.0170)	(0.0156)	(0.0158)
Parental Education (unit: 1 year)	0.0550	0.1528	-0.0277	-0.0544	1.5368
	(0.0007)	(0.0009)	(0.0014)	(0.0017)	(0.0349)
High Expectations (0/1 Indicator)	0.097	0.3475	-0.0202	-0.1174	0.2458
	(0.0097)	(0.0118)	(0.0176)	(0.0208)	(0.0227)
Sample	Santiago metropolitan region	Santiago metropolitan region	Santiago metropolitan region	Santiago metropolitan region	Santiago metropolitan region

Notes

Main Effect coefficients reflect the change in mean utility associated with a one-unit change in the school characteristic. Heterogeneity coefficients reflect the change in mean utility associated with a school characteristic that results from a one-unit change in a given individual characteristic. High Expectations is an indicator for whether parents expect their child to graduate from college. Standard deviations are presented in parentheses. Note that main effect coefficients are not identified for time-invariant school characteristics.

Table 8. Simulated Enrollment Changes

	Simulation		Reduced-Form Estimates	
	Percentage Change in School Price	Percentage Change in Enrollment	Percentage Change in School Price	Percentage Change in Enrollment
	(1)	(2)	(3)	(4)
Baseline School Price in Quintile 1	0	6.09%	0.04%	5.25%
	-	-	(0.05)	(1.22)
Baseline School Price in Quintile 2	-0.60%	11.93%	-0.40%	13.11%
	-	-	(0.48)	(8.68)
Baseline School Price in Quintile 3	1.03%	1.67%	0.84%	3.58%
	-	-	(0.48)	(3.09)
Baseline School Price in Quintile 4	2.03%	-8.39%	2.24%	-4.22%
	-	-	(0.74)	(1.89)
Baseline School Price in Quintile 5	2.93%	-12.51%	4.21%	-9.40%
	-	-	(0.76)	(2.55)

Notes

Standard errors are presented in parentheses for reduced-form estimates. Percentage enrollment changes are constructed based on average school size in a given baseline school price quintile bin.

Figure 1: Model Simulations

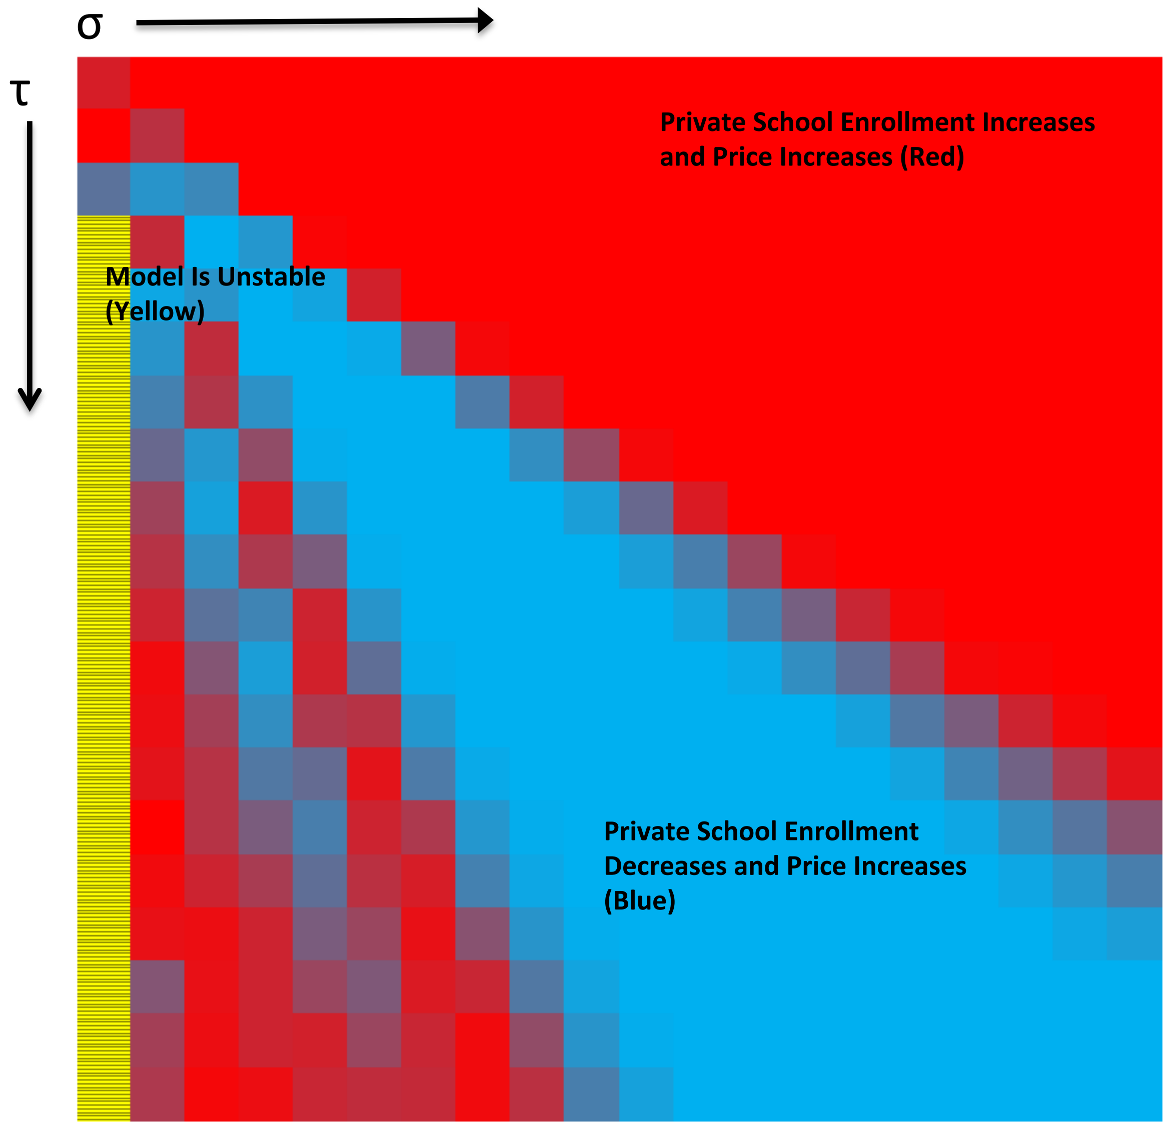

Notes: The figure displays results averaged over 50 model simulations. Y-axis values range from τ=0.25 to τ=5.0 and X-axis values range from σ=2.5 to σ=50.0. Baseline quality ratio between schools is set equal to 2. When printed in grayscale, darker cells are those in which private school enrollment increases, lighter cells are those in which private school enrollment declines, and cells with a horizontal-line pattern are those in which the model is unstable.

Figure 2: Municipality-Specific Elasticities

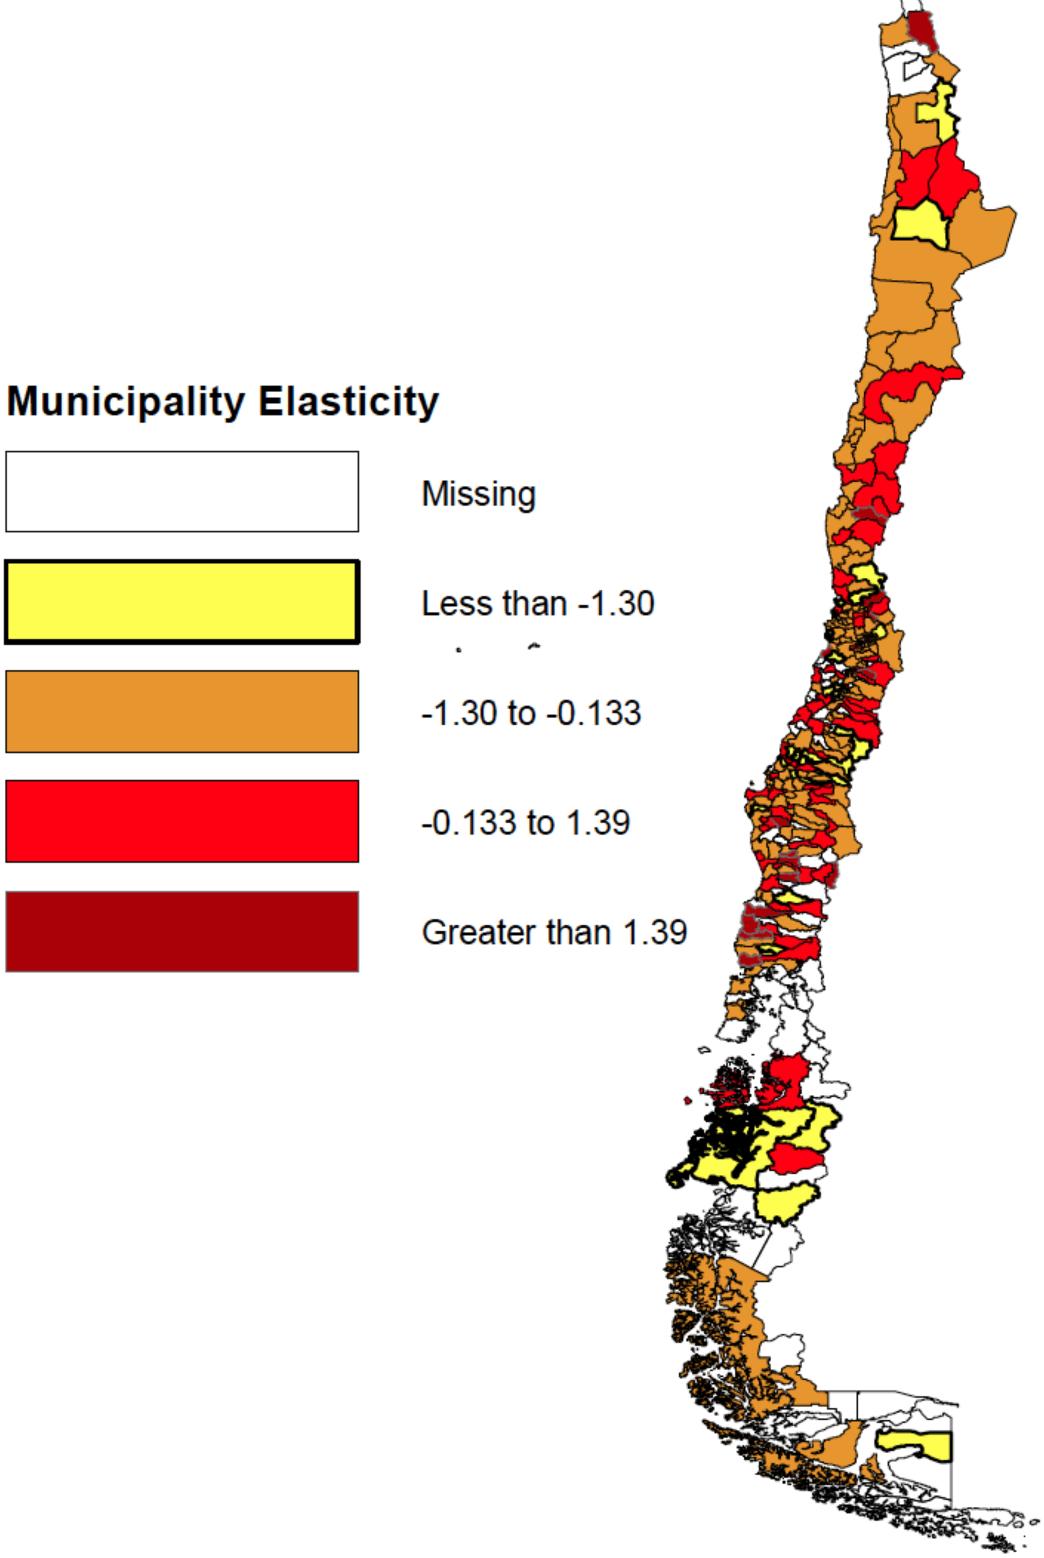

Figure 3: Time Series of Copper Prices and Chilean Exchange Rate

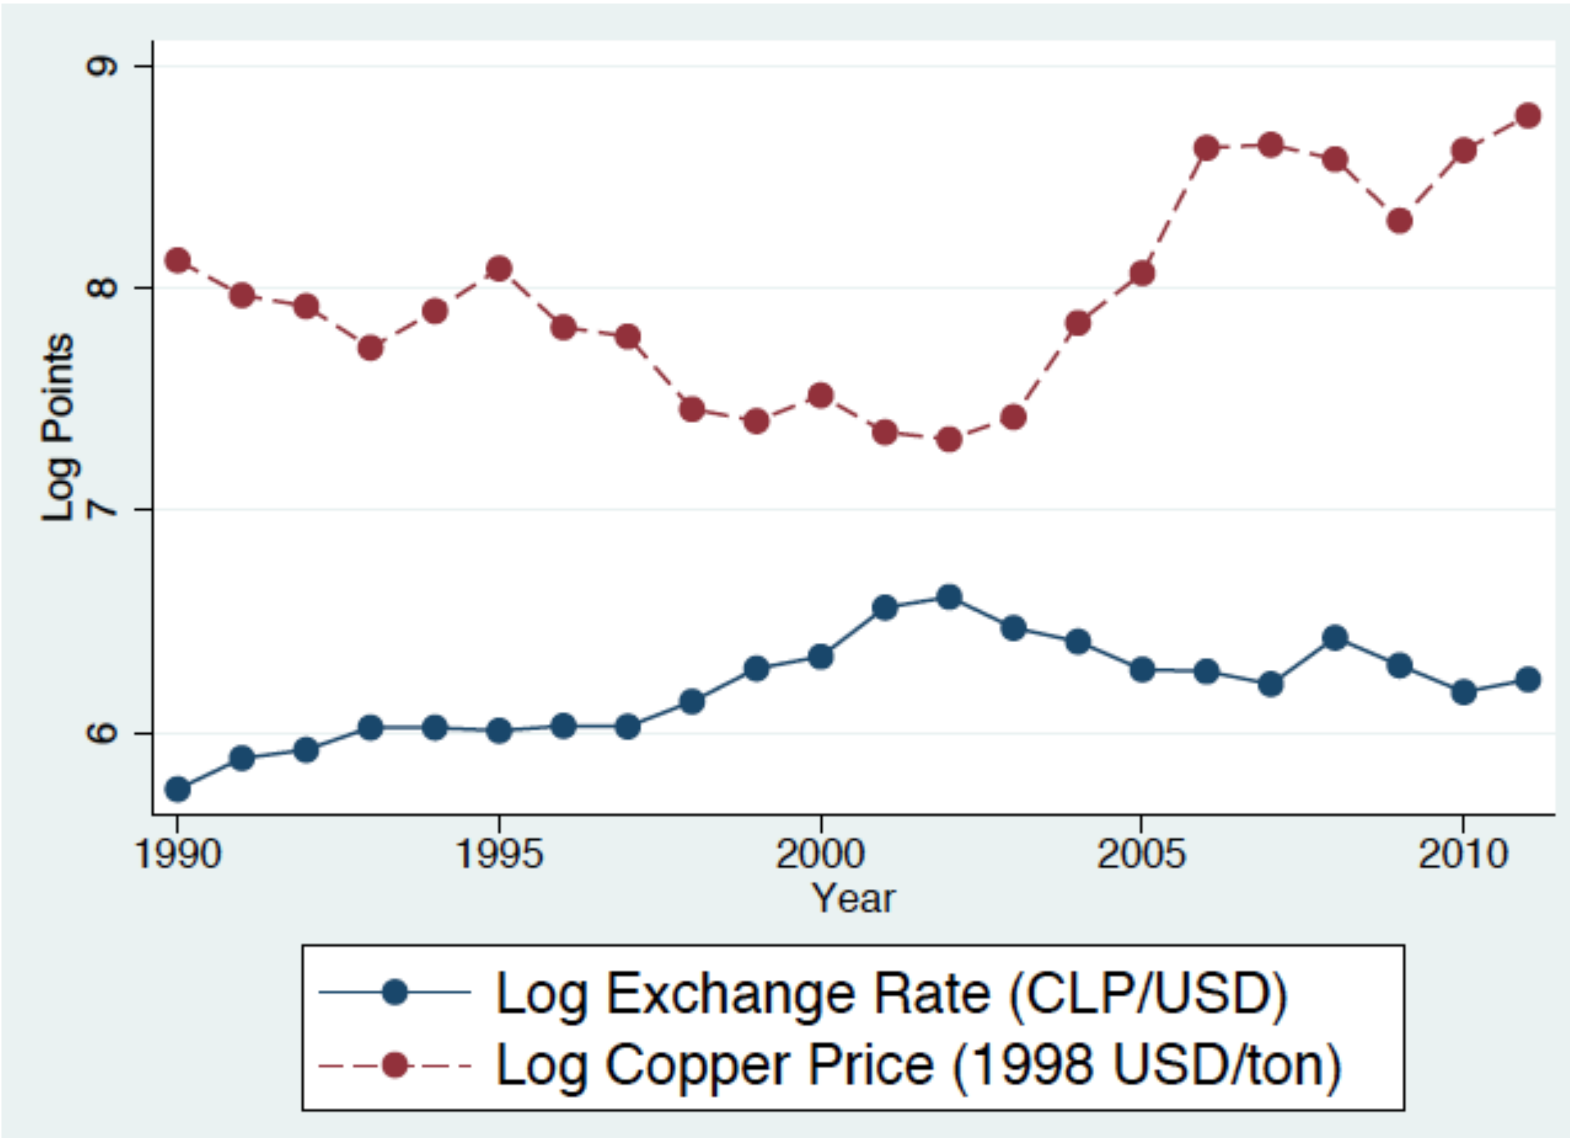

Appendix Table 1. Region-level Analysis

	Log Mean Income	Log School Price	Number of Public School Students
	(1)	(2)	(3)
Normalized Copper Shock (Region-level)	0.063**	0.0060**	5440.3***
	(0.030)	(0.0029)	(2050.3)
Year Fixed Effects	X	X	X
Region Fixed Effects	X		X
School Fixed Effects		X	
Mean of Dependent Variable	12.228	10.643	95,684
	[0.369]	[0.279]	[97,210]
Observations	60	59049	120
Specification	WLS	OLS	WLS
Sample	All (Region-level)	All (School-level)	All (Region-level)

Notes

1 Normalized Copper Shock is defined as the product of the normalized region-specific elasticity of income with respect to copper prices and the log copper price (denominated in 1998 USD). Regressions are clustered at the region level (there are 15 regions). * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Columns (1) and (3) employ a weighted least squares (WLS) specification, and weight observations by number of responses within region*year cell.

Appendix Table 2. School Price Determinants

	Log Price	Log Price	Log Price	Baseline School Price Quintile
	(1)	(2)	(3)	(4)
Mean Normalized Test Score	0.003***	0.013***	0.022***	0.156***
	(0.001)	(0.003)	(0.005)	(0.032)
Mean Parental Education (Years)	0.001***	0.002	0.006***	0.135***
	(0.0002)	(0.002)	(0.002)	(0.014)
Mean of Log Household Income	0.008***	0.037***	0.150***	0.582***
	(0.001)	(0.008)	(0.013)	(0.049)
Mean of High Expectations	0.012***	0.022**	0.111***	1.040***
	(0.002)	(0.009)	(0.013)	(0.096)
Average Teacher Experience (Years)	-0.0005***	-0.001	-0.004***	-0.050***
	(0.0001)	(0.001)	(0.0004)	(0.003)
Fraction of Teachers Certified	-0.012**	0.002	0.072***	0.209
	(0.005)	(0.015)	(0.021)	(0.136)
Fraction of Teachers with Graduate Degrees	0.003***	0.004	-0.007	0.008
	(0.001)	(0.003)	(0.006)	(0.054)
Fraction of Teachers that Use Computer Frequently for Work	0.005***	0.003	0.002	-0.043
	(0.001)	(0.003)	(0.008)	(0.062)
Year Fixed Effects	X	X		
School-Grade Fixed Effects	X	X		
Mean of Dependent Variable	10.572	10.825	10.404	2.260
	[0.263]	[0.271]	[0.161]	[1.631]
Observations	31131	9688	5194	5149
Sample	Fourth Graders (School*Grade-level)	Excluding public schools (School*Grade-level)	Year 2005 (School*Grade-level)	Year 2005 (School*Grade-level)

Notes

All regressions are estimated at the grade four level, as SIMCE data is most frequently available for grade four students. High Expectations is an indicator for whether parents expect their child to graduate from college. Columns (3)-(4) include only observations from the year 2005. Regressions are clustered at the municipality level (there are 270 municipalities). * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Appendix Table 3. Copper Shocks and Public School Funding

	FNDR ($)	FNDR ($)	FNDR ($)	FNDR ($)
	(1)	(2)	(3)	(4)
Normalized Copper Shock	-34848		-0.172	
	(24112)		(0.140)	
Aggregate Copper Shock		-26235		-0.057
		(22932)		(0.128)
Year Fixed Effects	X	X	X	X
Municipality Fixed Effects	X	X	X	X
Mean of Dependent Variable	434,293		10.616	
	[3,992,299]		[0.266]	
Observations	1885	1885	1703	1703
Specification	OLS	OLS	Poisson QMLE	Poisson QMLE
Sample	Municipality-level		Municipality-level	

Notes

Regressions are clustered at the municipality level (there are 270 municipalities). FNDR refers to the National Fund for Regional Development.

Appendix Table 4. Teacher Income

	Log Mean Income	Log Mean Income
	(1)	(2)
Normalized Copper Shock (Region-level)	-0.016	-0.013
	(0.016)	(0.010)
Normalized Copper Shock* Non-Teacher	0.058***	0.079***
	(0.018)	(0.025)
Year*Teacher Fixed Effects	X	X
Region*Teacher Fixed Effects	X	X
Mean of Dependent Variable	12.686	
	[0.583]	
Observations	120	120
Specification	OLS	WLS
Sample	All (Region-level)	

Notes

Normalized Copper Shock is defined as the product of the normalized region-specific elasticity of income with respect to copper prices and the log copper price (denominated in 1998 USD). Regressions are clustered at the region level (there are 15 regions). * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Column (2) employs a weighted least squares (WLS) specification, and weights observations by number of responses within region*year cell.

Appendix Table 5. Copper Shocks and Teacher Characteristics

	Average Experience (Years)	Average Experience (Years)	Average Experience (Years)	Share Teachers Certified	Share Teachers Certified	Share Teachers Certified
	(1)	(2)	(3)	(4)	(5)	(6)
Normalized Copper Shock	-0.061	-0.022	-0.034	0.005	0.005	0.005
	(0.178)	(0.216)	(0.222)	(0.004)	(0.005)	(0.005)
Normalized Copper Shock* Private School		-0.212			-0.002	
		(0.289)			(0.006)	
Normalized Copper Shock* Baseline School Price in Quintile 2			-0.526			0.013
			(0.432)			(0.017)
Normalized Copper Shock* Baseline School Price in Quintile 3			-0.416			0.005
			(0.460)			(0.012)
Normalized Copper Shock* Baseline School Price in Quintile 4			0.068			-0.0001
			(0.324)			(0.008)
Normalized Copper Shock* Baseline School Price in Quintile 5			-0.293			-0.0001
			(0.285)			(0.006)
Year Fixed Effects	X	X	X	X	X	X
School Fixed Effects	X	X	X	X	X	X
Mean of Dependent Variable	15.808			0.923		
	[7.652]			[0.133]		
Observations	47147	47147	46775	47147	47147	46775
Sample	All (School-level)			All (School-level)		

Notes

Normalized Copper Shock is defined as the product of the normalized municipality-specific elasticity of income with respect to copper prices and the log copper price (denominated in 1998 USD). Regressions are clustered at the municipality level (there are 270 municipalities). * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Appendix Table 6. Copper Shocks and Rental Income

	Log Imputed Rental Income	Log 50th Percentile Rental Income	Log 60th Percentile Rental Income	Log 70th Percentile Rental Income	Log 80th Percentile Rental Income	Log 90th Percentile Rental Income
	(1)	(2)	(3)	(4)	(5)	(6)
Normalized Copper Shock	0.019	0.027**	0.010	0.016	0.010	0.030*
	(0.012)	(0.013)	(0.012)	(0.011)	(0.012)	(0.016)
Year Fixed Effects	X	X	X	X	X	X
Municipality Fixed Effects	X	X	X	X	X	X
Observations	1078	1066	1070	1074	1074	1074
Sample	All (Municipality-level)		All (Municipality-level)		All (Municipality-level)	

Notes

Normalized Copper Shock is defined as the product of the normalized municipality-specific elasticity of income with respect to copper prices and the log copper price (denominated in 1998 USD). Regressions are clustered at the municipality level (there are 270 municipalities). * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Appendix Table 7. Copper Shocks and School Entry/Exit

	Number of Schools	Number of Public Schools	Number of Private Schools	Log Number of Schools	Log Number of Public Schools	Log Number of Private Schools
	(1)	(2)	(3)	(4)	(5)	(6)
Copper Shock	-0.050	0.111	-0.014	0.002	0.005	0.002
	(0.170)	(0.108)	(0.119)	(0.006)	(0.005)	(0.008)
Year Fixed Effects	X	X	X	X	X	X
Municipality Fixed Effects	X	X	X	X	X	X
Control for Lagged Dependent Var	X	X	X	X	X	X
Mean of Dependent Variable	31.2	17.5	13.7			
	(26.6)	(12.1)	(18.9)			
Specification		OLS			Poisson QMLE	
Observations	1890	1890	1890	1890	1890	1612
Sample	Municipality-level		Municipality-level		Municipality-level	

Notes

All specifications include a control for the lagged value of the dependent variable. Regressions are clustered at the municipality level (there are 270 municipalities). * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Appendix Table 8. Copper Shocks and Teacher Contracts

	Number of Contract Hours	Number of Contract Hours	Number of Contract Hours	Number of Teachers	Number of Teachers	Number of Teachers
	(1)	(2)	(3)	(4)	(5)	(6)
Normalized Copper Shock	-4.4	7.8	6.3	0.04	0.4***	0.4***
	(3.6)	(5.2)	(4.9)	(0.1)	(0.2)	(0.1)
Normalized Copper Shock* Private School		-64.9***			-2.0***	
		(20.0)			(0.6)	
Normalized Copper Shock* Baseline School Price in Quintile 2			56.2			1.1
			(61.2)			(1.9)
Normalized Copper Shock* Baseline School Price in Quintile 3			-5.1			-0.2
			(26.1)			(0.8)
Normalized Copper Shock* Baseline School Price in Quintile 4			-72.8***			-2.3***
			(20.2)			(0.6)
Normalized Copper Shock* Baseline School Price in Quintile 5			-73.9***			-2.1***
			(23.6)			(0.6)
Year Fixed Effects	X	X	X	X	X	X
School Fixed Effects	X	X	X	X	X	X
Mean of Dependent Variable	611.1			19.4		
	[589.8]			[17.6]		
Observations	46737	46737	46323	46737	46737	46323
Sample	All (School-level)			All (School-level)		

Notes

Normalized Copper Shock is defined as the product of the normalized municipality-specific elasticity of income with respect to copper prices and the log copper price (denominated in 1998 USD). Number of Contract Hours refers to the per week number of contracted teacher hours at a given school. Regressions are clustered at the municipality level (there are 270 municipalities). * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Appendix Table 9. Lead Structure of Copper Shocks

	Log School Price	Log School Price	Log School Price	Log School Price	Log School Price	Log School Price	Log School Price	Log School Price	Log School Price
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Lead Copper Shock (t+1)	0.0018				0.0018	0.0012	0.0044	-0.0005	
	(0.0040)				(0.0040)	(0.0038)	(0.0039)	(0.0047)	
Lead Copper Shock (t+2)		0.0008				0.0009	-0.0029	-0.0050	
		(0.0075)				(0.0076)	(0.0075)	(0.0106)	
Lead Copper Shock (t+3)			-0.0004				-0.0009	-0.0003	
			(0.0055)				(0.0056)	(0.0083)	
Lead Copper Shock (t+4)				0.0086				0.0088	
				(0.0063)				(0.0064)	
Aggregate Future Shock									0.0074
									(0.0114)
Year Fixed Effects	X	X	X	X	X	X	X	X	X
School Fixed Effects	X	X	X	X	X	X	X	X	X
Observations	12004	10176	8375	6583	12004	10176	8375	6583	6583
Sample	All (School-level)		All (School-level)		All (School-level)		All (School-level)		

Notes

The sample includes all private schools. Lead Copper Shock coefficients are estimated separately for lead years 1-4 in Columns (1)-(4). Columns (5)-(8) include lead terms jointly and Aggregate Future Shock is defined as the mean of lead copper shocks 1-4. Regressions are clustered at the municipality level (there are 270 municipalities).

Appendix Table 10. Copper Shocks and Teacher Performance

	Teaching Evaluations	Teaching Evaluations	Teaching Evaluations
	(1)	(2)	(3)
Normalized Copper Shock	-0.005	-0.005	-0.007
	(0.015)	(0.014)	(0.015)
Normalized Copper Shock* Private School		-0.004	
		(0.053)	
Normalized Copper Shock* Baseline School Price in Quintile 2			0.044
			(0.115)
Normalized Copper Shock* Baseline School Price in Quintile 3			-0.016
			(0.065)
Normalized Copper Shock* Baseline School Price in Quintile 4			0.010
			(0.086)
Normalized Copper Shock* Baseline School Price in Quintile 5			0.046
			(0.063)
Year Fixed Effects	X	X	X
School Fixed Effects	X	X	X
Mean of Dependent Variable	2.204		
	[0.249]		
Observations	24582	24582	24333
Sample	All (School-level)		

Notes

Normalized Copper Shock is defined as the product of the normalized municipality-specific elasticity of income with respect to copper prices and the log copper price (denominated in 1998 USD). Teaching Evaluations is the school-level mean of individual teaching evaluations which are scaled from one to four. These evaluations are based on a pedagogical statement written by the respondent and on a classroom video recording. Regressions are clustered at the municipality level (there are 270 municipalities). * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Appendix Table 11. Copper Shocks and Student Grade Point Averages

	Grade Point Average		
	(1)	(2)	(3)
Normalized Copper Shock	0.0092**	0.0103**	0.0093*
	(0.0044)	(0.0048)	(0.0050)
Normalized Copper Shock* Within-municipality Parental Education Quartile 2		-0.0012**	0.0002
		(0.0006)	(0.0004)
Normalized Copper Shock* Within-municipality Parental Education Quartile 3		-0.0006	0.0018***
		(0.0007)	(0.0005)
Normalized Copper Shock* Within-municipality Parental Education Quartile 4		0.0003	0.0028***
		(0.0009)	(0.0007)
Within-municipality Parental Education Quartile 2		0.1892***	0.1431***
		(0.0044)	(0.0030)
Within-municipality Parental Education Quartile 3		0.3051***	0.2219***
		(0.0048)	(0.0052)
Within-municipality Parental Education Quartile 4		0.4328***	0.3023***
		(0.0068)	(0.0078)
Year Fixed Effects	X	X	X
Grade Fixed Effects	X	X	X
Municipality Fixed Effects	X	X	
Lagged School-Municipality Fixed Effects			X
Mean of Dependent Variable	5.663		
	[0.866]		
Observations	17,059,576	10,562,118	9,822,852
Sample	All (Student-level)		

Notes

Normalized Copper Shock is defined as the product of the normalized municipality-specific elasticity of income with respect to copper prices and the log copper price (denominated in 1998 USD). Grade point average is measured on a scale from zero to seven. Regressions are clustered at the municipality level (there are 270 municipalities). * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Appendix Table 12. School Price and Enrollment Responses to Copper Shocks (Santiago Region)

	Log School Price	Log School Price	Log School Price	Number of Students	Number of Students	Number of Students
	(1)	(2)	(3)	(4)	(5)	(6)
Normalized Copper Shock	0.0076	-0.0133***	-0.0142***	-17.1	13.8	13.8
	(0.0049)	(0.0031)	(0.0031)	(12.1)	(15.0)	(14.9)
Normalized Copper Shock* Private School		0.0418***			-61.7***	
		(0.0083)			(17.8)	
Normalized Copper Shock* Baseline School Price in Quintile 2			0.0029			49.3
			(0.0078)			(72.4)
Normalized Copper Shock* Baseline School Price in Quintile 3			0.0179*			-40.4
			(0.0098)			(26.9)
Normalized Copper Shock* Baseline School Price in Quintile 4			0.0322***			-55.0***
			(0.0118)			(18.4)
Normalized Copper Shock* Baseline School Price in Quintile 5			0.0868***			-114.0***
			(0.0118)			(26.0)
Year Fixed Effects	X	X	X	X	X	X
School Fixed Effects	X	X	X	X	X	X
Mean of Dependent Variable	10.738			577.9		
	[0.294]			[499.3]		
Observations	12531	12531	12531	12531	12531	12531
Sample	Santiago Region (School-level)			Santiago Region (School-level)		

Notes

Normalized Copper Shock is defined as the product of the normalized municipality-specific elasticity of income with respect to copper prices and the log copper price (denominated in 1998 USD). Regressions are clustered at the municipality level (there are 52 municipalities). * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Appendix Figure 1: School Price Distribution

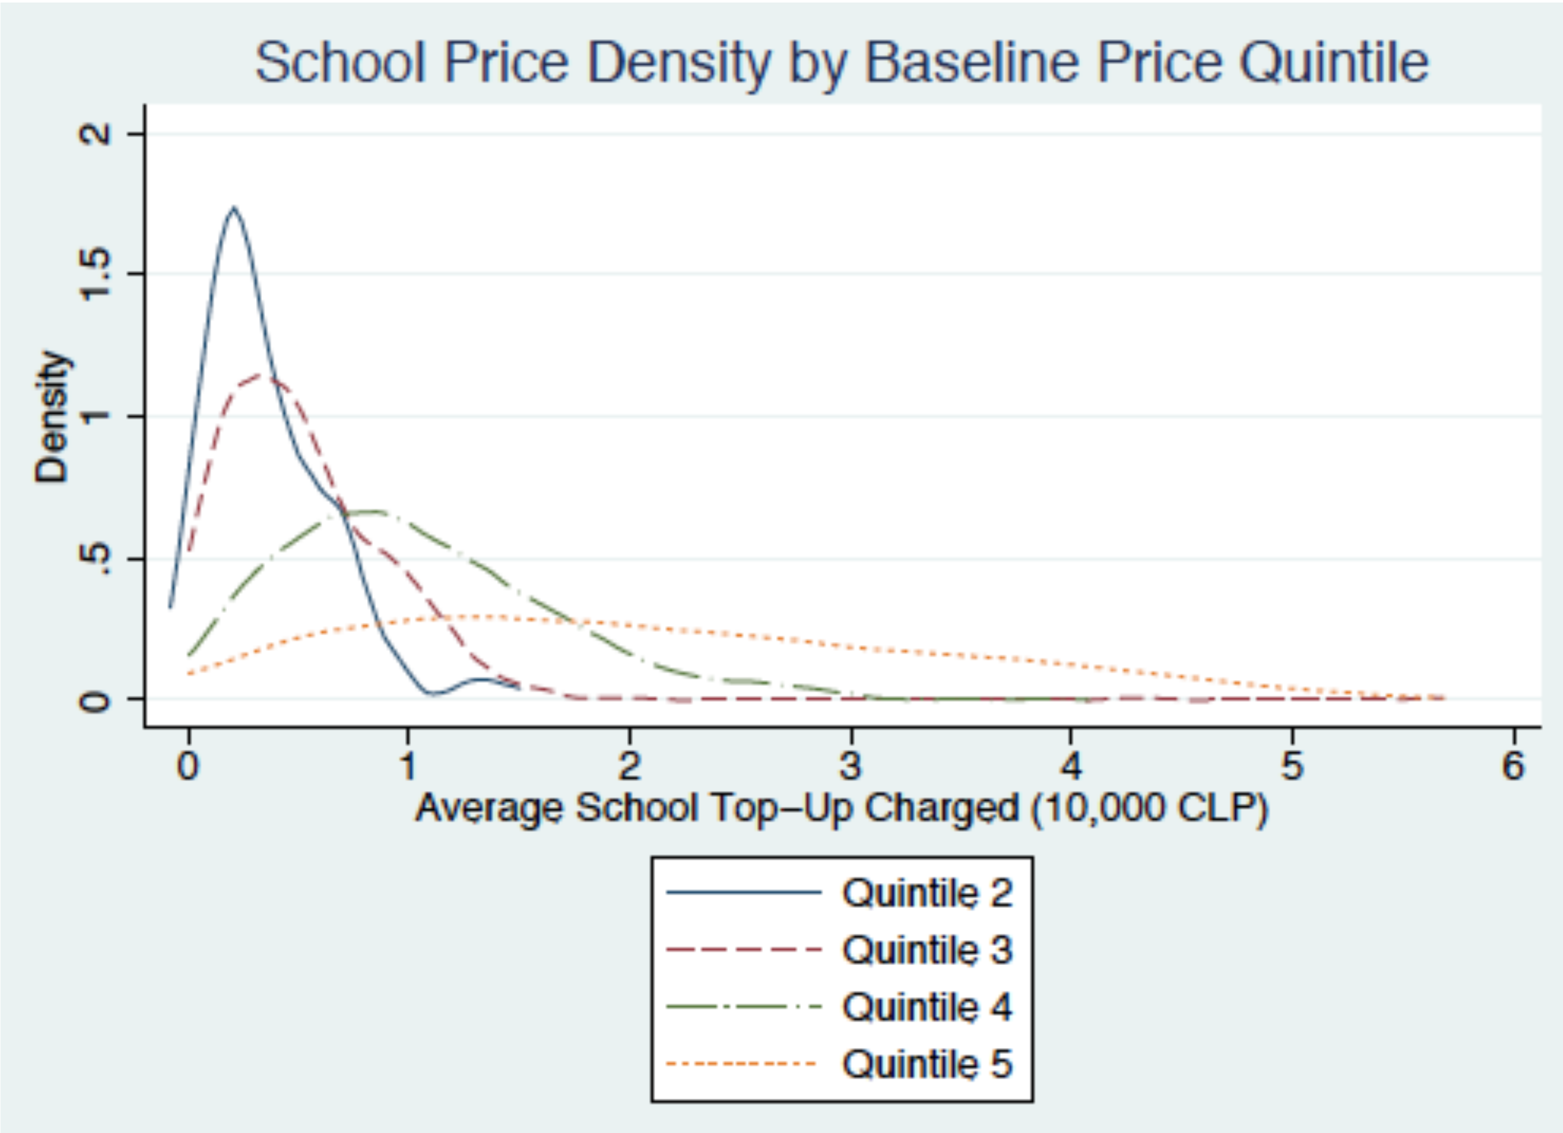

Note: Quintile 1 is excluded because 99.8% of quintile 1 schools do not charge any top-up.

1 Mathematical Appendix

Based on the first-order condition from the school's profit maximization and the expressions for demand and expected willingness to pay of parents that are presented in Section 2, I arrive at a set of three equations characterizing equilibrium prices, enrollment, and expected willingness to pay of parents whose children attend school j:

$$p^* - [\sigma + c + \alpha\tau V^* \frac{q^*}{n}] = 0 \tag{1}$$

$$q^* - \int_{\underline{v}}^{\bar{v}} \Gamma_j(\underline{s}|\alpha, v, p^*, s^*) N f(v) dv = 0 \tag{2}$$

$$V^* - \int_{\underline{v}}^{\bar{v}} v \frac{\Gamma_j(\underline{s}|\alpha, v, p^*, s^*)}{\int_{\underline{v}}^{\bar{v}} \Gamma_j(\underline{s}|\alpha, v, p^*, s^*) f(v) dv} f(v) dv = 0 \tag{3}$$

Here, $p^*, q^*,$ and V^* represent equilibrium prices, enrollment, and expected willingness to pay of parents whose children attend school j, respectively.

I can then apply the implicit function theorem to find expressions for $\frac{dp^*}{d\alpha}$ and $\frac{dq^*}{d\alpha}$. To do so, I construct the inverse of the Jacobian and multiply it by negative one times the vector of the partial derivatives of Appendix Equations (1)-(3) with respect to α. The Jacobian matrix of the partial derivatives of Appendix Equations (1)-(3) with respect to the equilibrium values of the three endogenous variables (p, q, and V) is as follows:

$$\text{J}=\begin{bmatrix} 1 & -\frac{\alpha\tau V^*}{n} & -\frac{\alpha\tau q^*}{n} \\ \frac{q^*}{\sigma} & 1 + \frac{\alpha\tau V^* q^*}{\sigma n} & 0 \\ \frac{V^*}{\sigma} & \frac{E[v^2]\alpha\tau}{\sigma n} + \frac{V^*}{q^*} & 1 \end{bmatrix}$$

The inverse of the Jacobian can then be expressed as:

$$J^{-1} = \frac{1}{1+2\frac{\alpha\tau V^* q^*}{\sigma n} + \frac{\alpha^2\tau^2 q^{*2}}{\sigma^2 n^2}(V^{*2} - E[v^2])} \begin{bmatrix} 1 + \frac{\alpha\tau V^* q^*}{\sigma n} & -\frac{\alpha^2\tau^2 E[v^2] q^*}{\sigma n^2} & \frac{\alpha\tau q^*}{n} + \frac{\alpha^2\tau^2 V^* q^{*2}}{\sigma n^2} \\ -\frac{q^*}{\sigma} & 1 + \frac{\alpha\tau V^* q^*}{\sigma n} & -\frac{\alpha\tau q^{*2}}{\sigma n} \\ \frac{\alpha\tau E[v^2] q^*}{\sigma^2 n} - \frac{\alpha\tau V^{*2} q^*}{\sigma^2 n} & -\frac{\alpha\tau E[v^2]}{\sigma n} - \frac{V^*}{q^*} - \frac{\alpha\tau V^{*2}}{\sigma n} & 1 + \frac{2\alpha\tau V^* q^*}{\sigma n} \end{bmatrix}$$

Finally, the vector of the partial derivatives of Appendix Equations (1)-(3) with respect to α is:

$$P = \begin{bmatrix} \frac{dEq1}{d\alpha} \\ \frac{dEq2}{d\alpha} \\ \frac{dEq3}{d\alpha} \end{bmatrix} = \begin{bmatrix} -\tau V^* \frac{q^*}{n} \\ -\frac{(\underline{s}-\tau\frac{q}{n})q^* V^*}{\sigma} + \frac{N}{\sigma}\int_{\underline{v}}^{\bar{v}} v \frac{exp\frac{\alpha v s(\underline{s})-p(\underline{s})}{\sigma}\int_{\hat{\underline{s}}} exp\frac{\alpha v s(\hat{\underline{s}})-p(\hat{\underline{s}})}{\sigma} s(\hat{\underline{s}})g(\hat{\underline{s}})d\hat{\underline{s}}}{(\int_{\hat{\underline{s}}} exp\frac{\alpha v s(\hat{\underline{s}})-p(\hat{\underline{s}})}{\sigma} g(\hat{\underline{s}})d\hat{\underline{s}})^2} f(v)dv \\ -\frac{(\underline{s}-\tau\frac{q}{n})E[v^2]}{\sigma} + \frac{1}{\sigma}\int_{\underline{v}}^{\bar{v}} v^2 \frac{exp\frac{\alpha v s(\underline{s})-p(\underline{s})}{\sigma}\int_{\hat{\underline{s}}} exp\frac{\alpha v s(\hat{\underline{s}})-p(\hat{\underline{s}})}{\sigma} s(\hat{\underline{s}})g(\hat{\underline{s}})d\hat{\underline{s}}}{(\int_{\hat{\underline{s}}} exp\frac{\alpha v s(\hat{\underline{s}})-p(\hat{\underline{s}})}{\sigma} g(\hat{\underline{s}})d\hat{\underline{s}})^2 \int_{\underline{v}}^{\bar{v}} \Gamma_j(\underline{s}|\alpha,v,p,s)f(v)dv} f(v)dv \\ + \int_{\underline{v}}^{\bar{v}} v \frac{\Gamma_j(\underline{s}|\alpha,v,p,s)\int_{\underline{v}}^{\bar{v}} \frac{v(\underline{s}-\tau\frac{q}{n})}{\sigma} \frac{exp\frac{\alpha v s(\underline{s})-p(\underline{s})}{\sigma}}{\int_{\hat{\underline{s}}} \frac{\alpha v s(\hat{\underline{s}})-p(\hat{\underline{s}})}{\sigma} g(\hat{\underline{s}})d\hat{\underline{s}}} f(v)dv}{(\int_{\underline{v}}^{\bar{v}} \Gamma_j(\underline{s}|\alpha,v,p,s)f(v)dv)^2} f(v)dv \\ - \int_{\underline{v}}^{\bar{v}} v \frac{\Gamma_j(\underline{s}|\alpha,v,p,s)\int_{\underline{v}}^{\bar{v}} \frac{exp\frac{\alpha v s(\underline{s})-p(\underline{s})}{\sigma}}{(\int_{\hat{\underline{s}}} \frac{\alpha v s(\hat{\underline{s}})-p(\hat{\underline{s}})}{\sigma} g(\hat{\underline{s}})d\hat{\underline{s}})^2} \int_{\hat{\underline{s}}} \frac{v s(\hat{\underline{s}})}{\sigma} exp\frac{\alpha v s(\hat{\underline{s}})-p(\hat{\underline{s}})}{\sigma} g(\hat{\underline{s}})d\hat{\underline{s}} f(v)dv}{(\int_{\underline{v}}^{\bar{v}} \Gamma_j(\underline{s}|\alpha,v,p,s)f(v)dv)^2} f(v)dv \end{bmatrix}$$

The intractability of the vector of partial derivatives with respect to α in turn implies that the expressions for price and enrollment comparative statics, which are equal to the first and second entries of the matrix $-J^{-1} * P$, are uninterpretable. To gain insight into the nature of price and enrollment responses to changing aggregate demand, I conduct simulations of the model (as detailed in Section 2).