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

November 14, 2019

Abstract

This paper studies private school market behavior. If school quality is a normal good and schools have market power, profit-maximizing schools may lower enrollment when incomes rise given parents' preferences for small class sizes and high-SES classmates. To demonstrate the possibility of such a response, I present simulations based on a stylized model of school pricing. Using data on local income variation in Chile, I show that positive income shocks cause private school prices to rise and enrollments to fall. Enrollment declines are concentrated among low-SES students, who do not experience the same test score gains as their higher-SES peers.

JEL: I24, I25, L1, O15

Keywords: Education and Inequality, Human Development, Educational Markets, School Vouchers

1 Introduction

In recent years, parents in developing countries have increasingly turned to 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 private school enrollment expands, the market for private school will play an increasing role in determining educational opportunity for children across the globe. To analyze the effects of this expansion, we must understand how private schools set their prices, how parents respond, and how this affects enrollment and student outcomes in equilibrium.

My paper highlights two key features that make the market for private school distinct from standard markets. First, private schools may have substantial market power due to idiosyncratic parental tastes, strong preferences for geographic convenience, and the important role of reputation, which takes time for new schools to build. Second, school quality is a normal good, and perceived school quality is decreasing in average class size and increasing in perceived peer quality. These features of the private school market raise the possibility that an increase in local incomes and the consequent rise in demand can actually decrease enrollment, because private schools increase prices beyond what would be expected in a typical market. The price increase also tends to increase average family socioeconomic status (SES) in private schools, as lower-SES families are unable to afford the higher prices despite their own increased incomes.[1] This potential private sector response to (short-run) economic growth underscores the importance of building on recent policy efforts intended to improve educational access, including the introduction of differentiated school vouchers in Chile and of private school reservations in India.

In this paper, I first simulate schools' profit-maximizing pricing behavior based on a stylized model of private school supply and demand, adopted from Urquiola and Verhoogen (2009),

[1] It may also be the case that the cost of educating high-SES students is lower than the cost of educating their lower-SES peers. This offers yet another potential factor that could induce profit-maximizing private schools to reduce enrollment.

which incorporates school market power and the effects of enrollment on perceived school quality and willingness to pay.[2] I show that, across a substantial range of parameter values, the profit-maximizing strategy may be to increase prices to the point that the equilibrium private school enrollment declines as students switch to public schools. The finding that increasing 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 effects of aggregate changes in income on private school prices and enrollment, I use administrative school price and enrollment data provided by the Chilean Ministry of Education. Chile provides an ideal setting for studying the pricing behavior of profit-maximizing schools in response to municipality-level income shocks. The advantages of the Chilean setting are two-fold. First, its private school market is well established and covers the whole country. Chile introduced a nationwide school voucher system in 1981, and during the 2005 to 2013 study period, the country allowed private schools to accept vouchers while charging additional fees. Second, the Chilean Ministry of Education collects data on private school enrollment and prices that are not available in many other countries, and the Ministry links these data to student records. This linkage allows me to investigate how school pricing decisions affect students' enrollment responses and academic performance. While my empirical analysis focuses on the Chilean educational market, my theoretical predictions are broadly applicable in any setting where schools set prices, schools have market power, and perceived school quality relates to characteristics of the student body.

Although the Chilean private school sector has grown substantially in recent years, I present descriptive evidence that this trend is not causally related to contemporaneous growth in Chilean household incomes. Specifically, I show that government funding for public and (most) private schools has also been rising rapidly over the same period. Exploiting variation in the rate of income growth across Chilean municipalities, I find a positive association between local income growth between 2005 and 2013 (my study period) and private school prices, but a negative association between local income growth and the market-level private school enrollment share.

[2] Urquiola and Verhoogen (2009) 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.

To identify the causal impact of aggregate income shocks, I exploit plausibly exogenous variation in income across Chilean municipalities driven by the differential responsiveness of local economies to the price of copper. The Chilean economy is reliant on the export sector relative to its neighbors, and Chile is the world's leading copper exporter. Consistent with the canonical "Dutch Disease" phenomenon, increased international copper prices lead the Chilean Peso to appreciate, and this differentially hurts the local markets that are most reliant on non-copper exports. To leverage this source of variation, I use historical global copper prices and pre-period survey data on household income. From these data, I construct municipality-specific measures of income elasticity with respect to global copper prices, derived from the period before educational microdata are available. The interaction of these elasticities with contemporaneous copper prices serves as the aggregate income shock. In the analysis, for ease of exposition, I focus on the effect of "positive copper shocks," which correspond to copper price decreases in municipalities with negative elasticities. In practice, copper prices rise and fall during the 2005-2013 study period, and so estimates reflect average responses to a combination of predicted increases and decreases in local incomes.

I leverage the differential responsiveness of local incomes to international copper prices to examine impacts on local educational markets. I find that private school prices rise and enrollment levels fall in response to copper price-induced positive income shocks. This enrollment decline is inconsistent with a standard, perfectly competitive market in which private schools expand (or enter the market) to meet increased demand and both prices and enrollment increase with aggregate income. Estimated price and enrollment responses are driven by the schools that were most expensive at baseline: a private school in the top quartile of the price distribution increases prices by 0.50-0.57% and decreases enrollment by 0.50% in response to a shock that increases local incomes by 1.0%. Estimates (though imprecise) suggest that revenues are rising for these private schools in response to the income shock. Evidence that the total input costs faced by these private schools fall in response to the income shock implies that profits increase. To rule out alternative interpretations of these findings, I show that copper shocks have a relatively uniform impact across the log income distribution and do not affect public school funding. Importantly, I confirm that the copper shocks constructed based on this methodology are not predictive of private school prices or enrollment levels from previous years.

Across analyses, I focus on reduced form estimates rather than using copper shocks to instrument for local incomes since local income measures are only available for a minority of the years included in the study period. This data constraint limits the strength of the first stage, which is particularly relevant when examining heterogeneous school responses to copper shocks. Nonetheless, I provide instrumental variables estimates for a limited set of key specifications, and I employ weak IV-robust inference when relevant. In addition, I present complementary OLS estimates that also indicate that short-run income increases lead to increases in prices and decreases in enrollment, though standard omitted variable bias concerns limit the interpretability of these findings.

To examine changes in student enrollment decisions in response to changing private school prices, I exploit the availability of unique student identifiers that allow students to be tracked across years and across schools. Grouping students by maternal education at the municipality level, I find that students whose mothers have the least formal education experience the largest changes in school quality in response to copper shocks. These school quality changes for disadvantaged students offset the expected positive direct effect of income on academic performance. Consequently, when incomes are predicted to rise, low maternal education students exhibit negligible test score changes, while students across the rest of the maternal education distribution experience gains. When incomes are predicted to fall, low maternal education students again exhibit negligible test score changes as school quality improvements offset income losses; in contrast, higher maternal education students experience test score decreases as declines in school quality amplify the direct effect of income losses.

To the best of my knowledge, this paper is the first to demonstrate with data how the profit motive may lead rising demand for school quality to increase educational stratification. Data constraints and a lack of plausibly exogenous variation in demand for or supply of private schooling have limited rigorous evidence on demand behavior and supply response within large-scale private school markets.[3] 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.[4] In contrast, my work sheds light on how schools behave strategically *within* private school markets and how school pricing affects students, by examining the impact of aggregate income shocks at the municipality by year level and by exploiting rich educational panel data.[5] In doing so, my research 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.

The remainder of the paper is structured as follows. Section 2 simulates private school price and enrollment responses to rising aggregate income based on a stylized model of private school sup-

[3] One exception is Muralidharan and Sundararaman (2015), 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.

[4] Here, authors are limited by the fact that the voucher system was introduced nationwide in 1981 as part of a larger educational reform. Early studies, such as Mizala and Romaguera (2000), employ OLS regressions of test scores on school type (private versus public) and include demographic controls in an effort to address 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 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.

[5] Andrabi et al. (2017) 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.

ply and demand. Section 3 discusses the data used in the empirical analysis. Section 4 documents institutional details of the Chilean educational market and presents descriptive analyses. Section 5 presents the empirical model and 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 summarizes the analysis and considers implications for policy.

2 Stylized Model and Simulation Results

My paper highlights two key features that make the market for private school distinct from standard markets. First, private schools may have substantial market power due to idiosyncratic parental tastes, strong preferences for geographic convenience, and the important role of reputation, which takes time for new schools to build. Second, school quality is a normal good, and perceived school quality is decreasing in average class size and increasing in perceived peer quality. In this section, I present a stylized version of the model from Urquiola and Verhoogen (2009) that incorporates these two features of the private school market, and I simulate conditions under which private schools maximize profits by decreasing equilibrium enrollment in response to increased aggregate demand.[6] Importantly, in the simulations presented, I find that private schools may maximize profits by reducing enrollment in response to rising aggregate income for a broad range of parameter values. As such, this section is intended to demonstrate the possibility that private schools maximize profits by reducing enrollment when demand rises. The subsequent empirical application is designed to test whether this response to rising demand is observed in practice.

To model demand for school quality, I assume that the indirect utility function of individual i is defined as in the standard vertical differentiation model but with the inclusion of an additional

[6] For the sake of simplicity, the model emphasizes demand for smaller class sizes as the key driver of reduced private school enrollment. In practice, the incentive to reduce enrollment is compounded by parents' preferences for higher average family socioeconomic status. It may also be the case that the cost of educating high-SES students is lower than the cost of educating their lower-SES peers. This offers yet another potential factor that could induce private schools to reduce enrollment.

error term:

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

In this framework, each individual has willingness to pay for quality v_i, where willingness to pay is a monotonically increasing function of household income. Each school j has an associated quality s and price p. Following the existing literature, I assume that schools choose an overall price level but do not 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, σ, can be interpreted as a measure of the extent of within-market school differentiation.

To model schools' operating costs, I follow Urquiola and Verhoogen (2009) and assume that there are three types of costs faced by schools: (1) F_s, a fixed cost of operating the school (2) F_n, a cost of operating each classroom, and (3) c, a per-student variable cost. Then, the profit function for school j can be expressed as follows:

$$\Pi_j = (p_j - c)q_j - n_j F_n - F_s \tag{2}$$

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, school quality s is a function of λ, the (pre-determined) component of school quality associated with the school's baseline reputation, and q. The distribution of λ within the market is exogenously-determined, as in Urquiola and Verhoogen (2009). Enrollment is a key determinant of perceived quality due to parents' preference for reduced classroom size (and for higher peer quality).[10] I consider the sub-

[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).

[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 (2015) 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.

case from Urquiola and Verhoogen (2009) in which there is a fixed number of classrooms in each school and the class size cap is non-binding. The former assumption, which rules out short-run private school increases in the number of classrooms, improves the tractability of the model by reducing the number of choice variables from three to two. As I discuss in more detail in Section 6.3, this assumption appears to be satisfied in response to short-run changes in demand for school quality in the Chilean setting.[11] The latter assumption is justified by the fact that classrooms are rarely at capacity during the study period, as documented in Section 4.1.

Following Urquiola and Verhoogen (2009), I assume that the functional form of the relationship between school quality and enrollment is characterized as follows:

$$s_j = \lambda_j ln(\frac{45}{q_j/n_j}) \tag{3}$$

In this equation, λ represents the (pre-determined) component of school quality associated with the school's baseline reputation and 45 corresponds to the legal maximum class size. The expression implies that perceived school quality is decreasing in classroom size and that a reduction in class size increases quality more at higher λ_j schools.

Given the above expressions for household preferences, school profit, and school quality, I derive comparative statics for school price and enrollment responses to a shift in the willingness-to-pay distribution. Following Urquiola and Verhoogen (2009), I assume that the number of schools in the market is fixed in the short run.[12] I focus on the case of a positive shock to aggregate demand, represented by a proportional increase in willingness-to-pay across the underlying distribution, and I begin by defining the school's profit maximization problem in terms of prices (i.e., as

[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, as discussed in Section 6.3, the most expensive private voucher schools marginally reduce the number of classrooms as they decrease total enrollment, but average classroom size still falls significantly. Public schools increase the number of classrooms as enrollments rise, but class size increases significantly in relative terms (and marginally in absolute terms). In practice, to the extent that willingness to pay is rising as enrollment declines (conditional on class size) due to perceived improvements in peer quality, reducing the number of classrooms may allow profit-maximizing schools to reduce costs while still partly capturing the increase in willingness to pay associated with reduced enrollment. This potential response is, however, beyond the scope of the model.

[12] This assumption is supported by the lack of significant school entry/exit in response to copper price-induced income shocks found in the empirical application (see the discussion in Section 6.3).

$\max_p \Pi(p, q(p, \lambda)))$. Here, equilibrium enrollment $q^*(p, \lambda)$ is determined by parental demand for a school with pre-determined reputational quality λ 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 $s(\lambda, q^*)$.[13]

In the Mathematical Appendix, I include the resultant expressions for demand and expected willingness to pay of parents. Based on these expressions and the first order condition from the school's profit maximization problem, I arrive at three equations characterizing equilibrium prices, enrollment, and expected willingness to pay of parents whose children attend a particular school, as well as expressions for changes in price and enrollment in response to changes in aggregate demand. While the change in private school prices associated with an upward shift in aggregate demand is unambiguously positive under the parameter conditions required for a maximum to exist (as discussed in the Mathematical Appendix), the sign of the enrollment effect is ambiguous.

To characterize comparative statics empirically, I simulate equilibrium price and enrollment changes in a market with one public and two private schools. The first and second private schools are assigned pre-determined school qualities (i.e., λ values) that are equal to two times and three times the (pre-determined) quality of the public school, respectively. I abstract away from the scenario in which class size caps are binding by considering a market with 45 students. Willingness to pay values are drawn from a Pareto distribution to approximate the income distribution. As shown in Appendix Figure 1, conclusions are qualitatively unchanged, however, under varying assumptions regarding the number of students and schools, relative (pre-determined) school qualities, the shape of the willingness to pay distribution, etc. In this market, the public school has a fixed price equal to marginal cost, while the private schools adjust prices and enrollments to maximize profits. In Figure 1, I present simulation results for a range of σ values (characterizing the degree of market differentiation) and take averages over 500 replications. Notably, as derived in Urquiola and Verhoogen (2009), I only arrive at stable, interior solutions for sufficiently high values of σ (which

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

limit the extent to which demand for schools rises as class size falls) relative to school quality and willingness to pay. Due to sampling variation, a given σ value may produce stable solutions for some replications, but not others (as seen in Figure 1). For σ values that produce an interior solution, private school prices are always increasing, as expected. Private school enrollments decline with positive probability across the σ parameter space (where an interior solution is found). The finding that private school enrollments are less likely to fall with rising aggregate demand at higher σ values reflects the fact that schools' incentives to reduce class size and enrollment fall as σ increases. Based on Gallego and Hernando (2008), which analyzes Chilean school choice in a semi-structural framework, I calculate that a σ value of roughly 25 aligns with the study setting. Though this estimate should be interpreted cautiously given the strong simplifying assumptions embedded in the simulation framework, rising aggregate demand causes private school enrollment declines in slightly over 50% of simulation replications in this range of the parameter space.[14]

3 Data

This section describes the Chilean labor force survey data used to measure local income levels and the copper price data used to construct copper price-induced income shocks. In addition, this section describes the rich administrative and survey data collected by the Chilean Ministry of Education and employed in this study.

3.1 Data on Copper Prices and Income

My identification strategy, described in detail in Section 5, requires that I estimate municipality-specific elasticities of income with respect to copper prices. To do so, I collect annual copper prices (denominated in 1998 USD) from the United States Geological Survey within the Depart-

[14] To assess the degree of horizontal differentiation (i.e. σ value) that best matches the empirical setting, I attempt to match the semi-structural Chilean school choice parameter estimates from Gallego and Hernando (2008) to the simulation parameters and the equilibrium price levels I calculate. However, this comparison is highly imperfect given that Gallego and Hernando (2008) incorporates a variety of school attributes and market structures outside the scope of my simulation, relies on data from before my study period, and includes only students in Santiago.

ment of the Interior (U.S. Geological Survey, 2016).[15] Survey data on historical municipality-level incomes comes from the Chilean National Socio-Economic Survey (CASEN). The measure of employment-based income that is used to construct my measure of local average household income is available for the following years: 1990, 1992, 1994, 1996, 1998, 2000, 2006, 2009, 2011, and 2013. The survey provides a repeated cross-section that lists respondents' municipality of residence. The CASEN survey is representative at the national, regional, and, in most cases, municipality level. Previous work, such as Auguste and Valenzuela (2006), has also used CASEN survey data to construct municipality-level socioeconomic variables.

3.2 Educational Administrative and Survey Data

Administrative data from the Chilean Ministry of Education provides a roster of all students enrolled in Chilean public and private schools in each year from 2004 to 2014. 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 municipality of residence. I also obtained data from the Ministry of Education on the mean school prices for fee-charging private voucher schools for the years 2004 to 2014.[16] Public primary schools cannot charge "top-up" fees during this period. The study period encompasses 2005-2013. Data from 2004 are used to construct baseline school price bins and so are excluded from the analysis sample (as discussed in more detail in Section 6.1). I exclude data from 2014 due to the introduction of a series of reforms in 2014 that phased out private schools' capacity to charge "top-up" fees.[17]

The Chilean Ministry of Education also releases annual test score data. Prior to 2005, the

[15] This data series captures the annual average U.S. producer copper price. While findings are similar when a measure of the global annual copper price provided by the Federal Reserve is employed, the Federal Reserve data series notes that the global copper price is actually measured exclusively based on annual Chilean prices.

[16] Based on the literature (see, for instance, Mizala and Urquiola, 2013) and an analysis of supplementary data from 2012-2013 on whether schools charged any "top-up" fees, I assume that private voucher schools with missing price data did not charge fees during the study period.

[17] In school-level specifications, I exclude primary schools serving special needs or adult students and drop 36 schools that either moved municipalities or changed public/private status during the study period.

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. Starting in 2005, they tested fourth graders annually. 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 present results averaged over the language and math components of the fourth grade exam.[18] For each cohort that takes the SIMCE, the Ministry of Education collects detailed survey data from parents. The survey of parents provides information on household demographic characteristics. The included measure of maternal education plays a central role in the student-level heterogeneity analysis detailed in Section 7. Finally, I make use of teacher-level data provided by the Ministry of Education that includes teacher characteristics, such as educational attainment and teaching experience.

4 Education in Chile

This section first provides a brief history of the Chilean voucher system and summarizes differences across schools based on public/private status and voucher system participation. In addition, this section characterizes recent trends in nationwide private school enrollment and exploits variation in the rate of income growth across Chilean municipalities to examine the association between local income growth and private school price and enrollment patterns.

4.1 Background and Summary Statistics

Section 2 demonstrates that the structure of education markets could lead changes in market conditions to have unanticipated effects 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 school pricing behavior and the implications for students'

[18] Analysis is limited to the language and math components because these are the only subject scores that are available in each year from 2005 to 2013. Following the existing literature, test scores are normalized based on the universe of available fourth grade math and language scores from 2005, the base year (see Feigenberg et al., 2018 for a more extensive discussion).

enrollment decisions and academic performance. Specifically, 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 established a system 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).[19] At the time that the reform was enacted, private schools had only a 22% market share in Chile (Gallego, 2013).

Table 1 provides relevant descriptive statistics by school type from 2005, the first year included in the study sample. By 2005, 50.9% of students in kindergarten through grade eight attended public schools, 42.8% attended private voucher schools, and 6.3% attended unsubsidized private schools.[20] School enrollments and class sizes are lower in public and no-fee private voucher schools than in fee-charging private schools due to the relative concentration of public and no-fee private voucher schools in rural areas. Growth over time in the private school sector reflects a combination of intensive and extensive margin changes: approximately 1.5% of private voucher schools exit the market each year and approximately 2.5% of private voucher schools are in their first year of operation during each year in the 2005-2013 study period.

Under the voucher system, all public primary schools are free. In contrast, private voucher schools are permitted to charge a "top-up" that is up to three times the annual voucher amount. As shown in Table 1, roughly half of private voucher schools charged such fees in 2005, with average per student fees of 13,103 CLP ($25 USD) per month (conditional on charging any fees). Voucher values are determined by grade level and length of school instruction (full- or half-day). In 2005, the average base voucher value per student was 36,093 CLP ($70 USD) per month.[21] Table

[19] The reform also significantly reduced government funding for education, decentralized educational decision-making to the local (municipal) level, relaxed curriculum standards, and revoked teacher union contracts.

[20] As in other related work (i.e., Urquiola and Verhoogen, 2009 and Gallego, 2013), unsubsidized private schools are excluded from the primary analyses. The government does not maintain data on prices charged by these schools, and 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. Using enrollment data, I confirm in Section 6.2 that unsubsidized private school enrollment does not increase (and in fact falls) when incomes are predicted to rise.

[21] The Subvención Escolar Preferencial (SEP) program, introduced in 2008, provided additional voucher revenues to participating public and private schools serving primary students in exchange for these schools agreeing to admit eligible (low-SES) applicants without charging "top-up" fees. In the subsequent analysis, I present estimated school price responses that do and do not account for this additional revenue source. For a more extensive

1 highlights the degree of stratification by family background in Chilean educational markets: average maternal years of schooling is four years higher in fee-charging private voucher schools than in public schools and is seven years higher in non-voucher private schools than in public schools. Differences in average test scores across school types follow a similar pattern.

Officially, private schools in Chile can selectively admit students while public schools that are not "at capacity" (i.e. have enrollment levels below 45 students per classroom) 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. In the public and private voucher school sectors, administrative records indicate that less than 2% and less than 5% of all grade levels were "at capacity" during the study period, respectively.

Most private voucher schools in Chile were for-profit during the study period. Indeed, Elacqua (2009) found that over 75% of Chilean private voucher schools were for-profit during this period, 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 schools within 2.5km, 5km, and 10km of a municipality center is 3, 4, and 9, respectively. Additionally, prior work employing home addresses available for a subset of students finds that the mean and median distance travelled by first graders in 2012 were 2.1 and 1.2 km, respectively (Aguirre, 2017). 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). The expansion of existing schools is constrained, at least in the short run, by the capacity of the school's physical plant.

In the empirical analysis, I restrict the sample to schools serving at least one grade level between kindergarten and eighth grade throughout the study period and focus on primary-level out-

discussion of the SEP program, see Feigenberg et al. (2018).

comes. Since primary school students travel shorter distances from home to school, this facilitates the construction of local market boundaries. While students are free to attend school in any municipality, 88% of primary students attend schools in their home municipality due to high cross-municipality commuting times. Previous research, such as Gallego (2013) and Hsieh and Urquiola (2006), has consequently defined education markets by municipality borders; in 2005, the average municipality included 28.5 schools serving students between kindergarten and grade eight.[22] I also focus on primary grade levels because voucher values for secondary school students vary across academic and vocational tracks, which substantially complicates the analysis of school revenue changes given the endogeneity of the track choice.[23]

4.2 Descriptive Analysis

Turning to the evolution of educational markets over the 2005-2013 period, Figure 2 plots nationwide trends in educational expenditures, household incomes and private school market share. The share of students attending private primary school is rising as incomes increase over this period, consistent with standard competitive market models. Importantly, however, public spending per student in public and voucher private schools is rising even more rapidly than household income during these same years, indicating that the revenues associated with operating private schools are increasing accordingly. This increase in public financial support for both public and private voucher schools raises questions about the appropriate interpretation of the positive correlation between annual mean household income and the nationwide private school enrollment share.

Table 2 exploits cross-municipality variation in household income growth over this same period as a step towards isolating the relationship between income and private school enrollment,

[22] Although students typically attend school in their home municipalities, school switching, which represents an important outcome in the analysis, is relatively common: 15% of primary students switch schools each year.

[23] An additional data-related justification is that 69% of schools with secondary-level instruction (and 83% of private schools with secondary-level instruction) also serve primary school students. Tuition data are only available at the school-by-year level and descriptive patterns suggest that tuition levels differ substantially for primary versus secondary school students. Consequently, restricting the analysis to schools with primary-level instruction facilitates comparisons of cross-school tuition level changes since only 17% of schools with primary-level instruction and only 32% of such private schools also serve secondary students.

and estimates suggest that rising private school attendance over this period is taking place *in spite of* increases in income.[24] Indeed, in a regression of the change in the share of students attending private school in each municipality on the municipality-specific percent change in mean household income, I identify a robust, negative relationship that is present across the universe of municipalities with private schools operating at baseline and in a subsample that includes only the 50 largest municipalities (as defined by number of schools in operation at baseline). Table 2 also indicates that average private voucher school per-student revenues rise in response to income increases. While these estimates are subject to standard omitted variables bias concerns, they motivate the subsequent analysis, which exploits a plausibly exogenous source of short-run income changes.

5 Copper Price-Induced Income Shocks

To isolate plausibly exogenous short-run income changes, I first use historical global copper prices and pre-period labor force survey data on household income. From these data, I construct municipality-specific measures of income elasticity with respect to global copper prices, derived from the period before educational microdata are available. I then exploit cross-sectional variation in the municipality-specific elasticities in combination with time-series variation in global copper prices from the study period.This reliance on a pre-determined measure of income responsiveness substantially mitigates identification concerns related to the co-determination of local incomes and educational market outcomes in the subsequent analysis.

Specifically, to produce municipality-specific elasticities, I estimate the following equation for each municipality using municipality-level data from 1990-2000 (the “pre-period”)[25]:

$$I_{mt} = \alpha_m + \mathscr{E}_m * P_t + \lambda_t + X_{mt} + \varepsilon_{mt} \tag{4}$$

[24] The correlations between mean (median) income changes and 25th and 75th percentile income changes are 0.36 (0.56) and 0.80 (0.78), respectively. Thus, mean/median income increases appear to reflect increases in incomes throughout the distribution.

[25] The 2003 CASEN survey is excluded as it does not capture information on employment-based earnings.

In this equation, I_{mt} represents log mean household employment income in municipality m in year t, and P_t is the log of the world copper price in year t, denominated in 1998 USD. X_{mt} represents a series of time-varying municipality-level controls characterizing the age distribution of the surveyed population, the mean literacy rate, marriage rate, urbanization rate and household size, as well as interactions between each of these covariates and survey year fixed effects. α_m and λ_t are municipality and year fixed effects, respectively. $\mathscr{E}_m$ 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. Additional details regarding sample construction are provided in the Data Appendix.

The global price of copper is expected to affect local incomes differently depending on export intensity and whether an area produces copper. For copper producing areas, incomes will likely rise with copper prices. In areas that export products other than copper, a 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.[26] Consequently, as global copper prices rise, export-oriented industries suffer. Chilean exports represent 34% of GDP (compared to, for example, 20% in Argentina), which makes the relationship between copper prices and export industry competitiveness particularly salient (World Bank, 2013). Indeed, during the study period, as copper prices rose dramatically, exporters concerned about currency appreciation advocated for capital controls and the Chilean government invested over $12 billion in a program aimed at weakening the currency (Pica and Wisnefski, 2012).

To provide additional evidence that currency appreciation drives the municipality-level income responses to copper price changes, I examine the relationship between municipality-level income elasticity coefficients and region-level exports as well as the industrial composition of region-level output.[27] Using pre-period data, I find that, conditional on regional log GDP per capita, local

[26] See Figure 3 for the time-series plots of copper prices and the exchange rate.
[27] Data on region-level output by industry are provided by the Chilean Instituto Nacional de Estadísticas and data

elasticities are increasing in log per capita regional exports in the primary copper-producing regions of Chile, but falling in log per capita exports in the remainder of the country.[28] In a complementary specification, I regress local elasticities on the interaction between region-level exports as a share of GDP and the share of exports derived from the copper industry. Here, the coefficient on the level term characterizing region-level exports as a share of GDP is negative, while the interaction term coefficient is positive and over three times as large in magnitude.

Turning to the industrial composition of region-level economic activity, roughly 90% of Chilean exports during the 1990s were derived from either mining activity or industrial manufacturing. I find that the share of region-level economic output derived from manufacturing is negatively associated with local elasticities. While data on municipality-level output are not available, the CASEN survey does ask respondents about the industry in which they are employed. Analyzing municipality-level employment patterns, I correspondingly find that the manufacturing employment share is negatively associated with local income elasticities.[29]

In sum, these findings provide additional support for the hypothesis that negative local income responses to copper price increases are caused by reduced export competitiveness. One concern regarding interpretation is that copper price changes may differentially affect educational market outcomes for residents of copper-producing areas if educational investment responses in these areas are driven by changes in the perceived future returns to employment in the copper mining sector. To isolate income variation driven by export competitiveness, I exclude schools located in the region of Antofagasta, in which the majority of Chilean copper is mined.[30]

Throughout the analysis, I define the copper shock C_{mt} assigned to municipality m in year t

---

on region-level exports are provided by the Chilean Ministerio de Planificación y Cooperación. Unfortunately, measures of economic output and export activity are not available at a more disaggregated geographical level.

28 I rely on data from the earliest pre-period year for which both region-level exports and region-level output were available (1996). There were a total of 13 regions in Chile during the 1990s (there are 16 regions at present).

29 Consistent with the possibility that currency appreciation has a detrimental impact on the tourism industry, I also find a negative association between municipality-specific elasticities and the share of employment in the "business" sector (which explicitly includes hotels and restaurants).

30 3.5% of Chilean students reside in Antofagasta and results are robust to including this region in the analysis. While the analysis could alternatively exploit variation in the Chilean exchange rate directly, the Chilean government has played an active role in fighting currency appreciation during the sample period, which raises endogeneity concerns that are mitigated by the copper shock-based approach (Pica and Wisnefski, 2012).

as being equal to the municipality-specific elasticity $\mathscr{E}_m$ 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" and corresponds to a predicted increase in local incomes.

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 2006, 2009, 2011, and 2013:

$$I_{mt} = \alpha + \beta * C_{mt} + \gamma_m + \lambda_t + \varepsilon_{mt} \tag{5}$$

Here, C_{mt} represents the aggregate copper shock experienced by residents of municipality m in year t, I_{mt} represents log mean household income in municipality m in year t, and γ_m and λ_t represent municipality and year fixed effects, respectively. In this and in all subsequent specifications that employ the copper shock regressor, a two-step bootstrapping procedure is implemented to construct accurate standard errors that account for the fact that elasticities are themselves estimated.[31] The inclusion of year fixed effects in this specification removes the effect of copper price variation on local incomes that is common across municipalities. Consequently, β is identified based on variation in the strength of local economic shocks that results from cross-municipality differences in responsiveness to global copper price variation (i.e. the income elasticities estimated in Equation (4)). In practice, a decrease in the price of copper is expected to lead to a larger positive effect on local incomes in those municipalities with elasticities that are negative and largest in magnitude.

Table 3 reveals the predicted positive relationship between copper shocks constructed using historical data and household incomes during the period for which educational microdata are available. Column (1) implies that a one-unit positive copper shock is associated with a 23.4% increase

[31] Specifically, to correct for inconsistent standard errors produced by standard inference on generated regressors, as discussed in Pagan (1984) and Murphy and Topel (1985), a random sample is first drawn with replacement from the CASEN survey data for the 1990-2000 period that is used to construct municipality-specific elasticity estimates. A sample of municipalities or schools is then drawn with replacement in the "second stage" regression (depending on the particular specification), with block bootstrapping at the municipality level to account for potentially correlated residuals across municipality-specific observations. This process is repeated 500 times and the standard deviation estimated for each "second stage" coefficient is the standard error that I present.

in the mean household income of municipality residents.[32] Columns (2)-(4) show that this estimated relationship is robust to the inclusion of the same controls included in Equation (4) and to the inclusion of municipality-specific weights based on the baseline (2004) number of schools in each municipality. In Columns (5)-(6), I linearly interpolate log mean household income values for missing years and present regression estimates from Equation (5) for 2005-2013. The decline in the magnitude of point estimates occurs since mean income is linearly imputed based only on observed municipality-specific mean income values. As a result, there is, by construction, no scope for imputed income values to be affected by non-linear changes in copper prices for the years for which CASEN data are unavailable.

In Figure 4, I present point estimates and confidence intervals from a Poisson QMLE specification that characterizes impact heterogeneity by income decile.[33] Specifically, I jointly estimate interactions between the copper shock and a set of income decile indicators (while including income decile, municipality and year fixed effects). Estimates indicate that impacts of the copper shock are relatively uniform across the income distribution.[34] Importantly, the lack of an apparent relationship between income decile and estimated copper shock impact suggests that distributional effects of the shock cannot explain the student-level heterogeneity that is investigated in Section 7.

6 Impacts of Local Income on School Prices and Enrollment

This section first presents benchmark estimates of the impact of copper price-induced income shocks on school prices, enrollment levels, and school revenues. For comparison, I discuss OLS

[32] A one-unit positive copper shock implies that copper prices double in a municipality that is at the 95th percentile in the distribution of elasticities. Thus, a positive copper shock of this magnitude 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.

[33] I return here to including data from only the four years for which income data are available.

[34] Below the 30th percentile of the household earnings distribution, a significant share of municipalities have household income measures that are always equal to zero and so sample selection issues emerge. Using an OLS specification that employs log average household income as the dependent variable, I confirm that estimated impacts at the 30th percentile and above appear similarly uniform. When decile-by-municipality fixed effects are included in the specification, point estimates suggest that income gains are larger at the two extremes of the distribution. However, likely due to the large number of resultant fixed effects parameters, estimated coefficients are imprecise and equality of estimated impacts across the income distribution cannot be rejected.

estimates of the relationship between school-level outcomes and local incomes. I next present falsification test results that provide support for my preferred research design. While my analysis is focused on the "reduced-form" relationship between educational outcomes and copper shocks, I provide instrumental variables estimates for a limited set of key specifications. The remainder of the section provides evidence on the robustness of findings, causal mechanisms, and implications for net changes in school profits.

6.1 Effects of Copper-Price Induced Income Shocks: Benchmark Estimates

Identifying the causal relationship between local incomes and educational market outcomes is a complex problem. Schools' simultaneous determination of price and enrollment levels implies that the subsequent analysis must carefully consider joint impacts along these two margins. Moreover, even if all schools share common production and cost functions, differences in baseline reputation may drive heterogeneous responses. As outlined in Section 2, the enrollment response to a change in aggregate income is ambiguous. This section estimates both price and enrollment relationships.

The following benchmark specification is used to estimate average school-level price impacts:

$$P_{smt} = \alpha + \beta_1 * C_{mt} + \beta_2 * C_{mt} * V_s + \gamma_s + \lambda_t + \varepsilon_{smt} \tag{6}$$

Here, 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 base voucher value received by school s in year t for primary-level students.[35] V_s is an indicator variable equal to one for private voucher schools and C_{mt} represents the aggregate copper shock experienced by residents of municipality m in year t.[36] γ_s and λ_t represent school and

[35] Survey evidence reveals limited within-cohort price discrimination based on student characteristics in Chilean private schools. Consequently, in analyzing price changes, school-level specifications are appropriate. The benchmark price measured is calculated using base voucher values rather than actual voucher receipts given that voucher receipts are an endogenous function of lagged attendance rates. While this measure includes only voucher values determined as a function of grade level and a "zone"-specific adjustment based on a school's location, subsequent specifications examine how estimated price responsiveness varies when additional voucher revenue sources are incorporated and when actual rather than predicted voucher revenues are used.

[36] Since the Chilean school year begins in March and since students may still switch schools after the school year

year fixed effects, respectively.[37] In specifications examining enrollment responses, the dependent variable is the number of primary-level students in school s in municipality m in year t. Enrollment is examined in levels rather than logs, given evidence that income increases lead to within-sector movement from smaller to larger schools. This, in turn, generates negative log enrollment estimates in the public and private school sectors (with a more negative estimate in the latter) in spite of the fact that, as I will show, the number of students in the market is non-decreasing.

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 to one if the tuition charged at baseline (in 2004, the first year for which price data are available) by school s falls into quartile q. Quartiles are constructed within regions to account for substantial heterogeneity in private school tuition levels across more versus less urbanized regions of the country.[38] I proceed to estimate the following specification:

$$P_{smt} = \alpha + \beta_1 * C_{mt} + \sum_{q=2}^{4} \beta_q * C_{mt} * P_{sq} + \gamma_s + \lambda_t + \varepsilon_{smt} \tag{7}$$

Quartile one corresponds to public schools, quartile two corresponds to private school that do not charge "top-up" fees at baseline, and quartiles three through four correspond to fee-charging private schools. Research on the Chilean education sector suggests that private voucher school quality is, on average, higher than public school quality.[39] Within the private school sector, price is

has begun, contemporaneous income changes have the potential to affect contemporaneous school- and student-level outcomes. The use of contemporaneous copper shocks in benchmark specifications also facilitates the interpretation of falsification checks based on lead values. Nonetheless, as detailed in Section 6.3, I confirm that findings are robust to alternatively employing a lagged copper shock measure or a measure that averages over contemporaneous and lagged values of the copper shock.

37 The sample excludes schools in municipalities without a sufficient number of CASEN observations to allow for the estimation of municipality-specific income elasticities. This reduces the number of school-by-year level observations by 10.6% and reduces the number of primary students attending included schools by 4.5%.

38 93% of schools in the sample appeared in 2004. Results are robust to defining price quartile for the remaining schools based on the first year that each school was open and only including observations from subsequent years.

39 In a meta-analysis, Drago and Paredes (2011) concludes that private school test scores are approximately one-

highly positively correlated with measures of perceived school quality, such as average test scores and students' average household income. As a result, heterogeneous impacts by baseline school price can be readily interpreted as reflecting impact heterogeneity by baseline school quality.[40]

Column (1) of Table 4 indicates that a one-unit positive copper shock causes an average private school price increase of 3.8% (significant at the 1% confidence level). As expected, public school prices, which are determined entirely by voucher values, are unaffected. Column (2) estimates Equation (7) and reveals that private school price increases are driven by schools in the upper portion of the baseline price distribution; prices rise by 6.3% and 11.8% in quartiles three and four, respectively. In Columns (3)-(4), the dependent variable uses actual base voucher revenues, which are a function of lagged student attendance, and results appear similar to those in Columns (1)-(2).

I next examine corresponding enrollment responses to copper price-induced income shocks. Column (5) reveals a positive and statistically significant increase in average enrollment in response to a positive copper shock. Since primary school enrollment in Chile is close to universal, this enrollment increase results from in-migration rather than reduced school dropout. In supplementary analyses (described in Section 6.3), I show that the market-level enrollment response is concentrated in rural municipalities and that overall patterns of school type-specific price and enrollment responses shown in Table 4 are similar when these rural municipalities are excluded. Columns (6)-(7) indicate that price impact heterogeneity is matched by heterogeneous changes in average enrollment. Public schools increase enrollment significantly and the implied enrollment change for private schools in price quartile two is positive. In contrast, coefficients on the interaction terms for quartiles three through four are indicative of net enrollment declines and are statistically significant at the 1% level. These coefficients imply that enrollment falls by 27.7 stu-

tenth of a standard deviation higher than public school scores conditional on student characteristics.

40 Appendix Table 1 shows the cross-sectional relationship between baseline price quartile and a series of school quality measures: mean maternal education and log household income of classmates, mean test scores, fraction of teachers with college training, and fraction of teachers with one or fewer years of experience. I also present corresponding correlations between school fees and these characteristics and examine changes in school pricing over time. Results are shown for both pairwise and joint regressions given that many of the included school quality measures are highly positively correlated. Estimates indicate that higher student test scores, mean log household income, and mean maternal education of classmates are all positively associated with school prices. In Columns (1)-(2), the finding that the fraction of inexperienced teachers is positively associated with school price quartile is explained by the fact that teacher experience is higher, on average, in the public school sector.

dents (5.9%) per school in quartile three and by 52.2 students (11.6%) per school in quartile four.

Column (8) combines price and enrollment results by estimating Equation (7) while replacing the dependent variable with log total school revenues. Since voucher revenues are an endogenous function of lagged attendance patterns and fixed as well as variable revenue sources, these estimates do not line up precisely with estimated price and enrollment responses. Estimates reveal that revenue changes are negative in public schools and in private schools that do not charge "top-up" fees at baseline due to the previously-noted within-sector movement of students from smaller to larger schools in response to positive copper price-induced income shocks (paired with the inclusion of a logged dependent variable). In contrast, net revenue changes are statistically indistinguishable from zero for fee-charging private voucher schools. Columns (9)-(10) aggregate school-level observations to the baseline school price quartile-by-municipality-by-year level to address the downward effect on estimates induced by student movement from smaller to larger schools. Here, estimates (though in some cases imprecise) provide evidence of rising revenues in fee-charging private schools.[41] In Section 6.3, I analyze the implications of measured revenue changes for school profits in the private sector.

For comparison, Appendix Table 2 regresses school outcomes on local incomes and shows that private school prices rise and relative private school enrollments fall when incomes increase (although the enrollment estimate is imprecise).[42] However, in corresponding specifications that include lead terms, I find that rising future income significantly increases contemporaneous private school prices and reduces private school enrollment. These estimates suggest that differences in income trajectories across markets limit the extent to which OLS specifications can shed light on the causal effects of changing incomes *per se*. For instance, communities experiencing sustained earnings growth may also be experiencing changes in the real or perceived return to schooling that separately influence demand for schooling. Such omitted variable bias concerns highlight the

[41] Given that the voucher system provides the majority of funding for participating schools, the finding that school revenues are either falling or unchanged in public schools and in private schools that do not charge "top-up" fees at baseline suggests that rising attendance in these schools does not appear to be driven by increases in either the per-student or total funding made available to these schools by the central government.

[42] The analysis sample includes the four years for which income data are available (2006, 2009, 2011, and 2013).

relative strength of the copper price-based research design that represents my preferred approach.

To provide additional support for the copper price-based research design, Table 5 presents results from a series of falsification tests. Specifically, Columns (1)-(3) of Table 5 re-estimate Equation (6), characterizing heterogeneous price responses to positive copper price-induced income shocks based on private school status, but add a series of lead terms ($C_{m,t+1}$, $C_{m,t+2}$, and $C_{m,t+3}$) to test whether future copper shocks are correlated with current school prices.[43] Columns (4)-(6) of Table 5 test whether lead copper shocks predict changes in public/private school enrollment. Reassuringly, none of the lead term coefficients in these specifications are statistically significant.

6.2 Instrumental Variables Estimates

While previous estimates examined how local copper shocks affect school prices and enrollment levels, the underlying causal relationship of interest is between local incomes and these educational outcomes. Unfortunately, the estimated relationship between local incomes and copper shocks is not sufficiently strong to allow for instrumental variables estimates to be presented for the full set of heterogeneity analyses included in Table 4. However, examining municipality-level enrollment responses to local income changes requires only a single instrument. Table 6 leverages this fact to present a series of municipality-level IV estimates. Specifically, the second stage equation used to estimate municipality-level IV specifications is of the following form:

$$E_{mt} = \alpha + \beta * \hat{I}_{mt} + \gamma_m + \lambda_t + \varepsilon_{mt} \tag{8}$$

E_{mt} is alternatively the total number of enrolled primary students, the number of primary students enrolled in non-voucher private schools, or the number of primary students enrolled in public schools in municipality m in year t. In the corresponding first stage equation, the copper shock experienced by municipality m in year t, C_{mt}, serves as an instrument for log mean household income,

[43] Given the high degree of serial correlation in the copper price data, contemporaneous copper prices are included when lead copper price elasticities are estimated and the contemporaneous copper shock term is included in Table 5 specifications along with the relevant lead terms.

I_{mt}. The first stage F-statistic of 20.5 for these specifications is sufficient to allow for conventional inference. Importantly, these IV specifications rely on the imputation of local incomes for years for which CASEN data are unavailable (the first stage coefficient of 0.118 is shown in Column (5) of Table 3). Across specifications, Panel B presents reduced form estimates for comparison.

Here, the exclusion restriction requires that copper shocks affect educational market outcomes only through changes in local incomes. While there are many mechanisms that may explain the link between local incomes and educational market outcomes (addressed in detail below), the exclusion restriction seems reasonably likely to hold. The exclusion restriction would be violated if, for instance, short-run copper price fluctuations influenced long-run educational investments by changing expectations of future employment– this seems unlikely to be a concern in practice.

Columns (1)-(4) of Panel A of Table 6 present IV estimates which confirm that school-level findings from Table 4 scale up to reflect significant changes in the municipality-level number of students attending public schools (both with and without a control for total municipality-level enrollment). Estimates also confirm that enrollment declines in private voucher schools in response to rising income are not matched by enrollment increases in the non-voucher private school sector (indeed, we observe parallel declines in private non-voucher school enrollment). This set of results previews the finding (discussed in Section 6.3) that there is limited school entry/exit in response to observed income changes. These municipality-level estimates also demonstrate the strength of the IV approach in the subset of specifications for which only a single instrument is needed.[44]

School price measures cannot be readily scaled up to the municipality level and so Columns (5)-(7) present IV estimates from a series of school-level specifications in which I instrument for the log of mean household income and its interaction with an indicator for private school status using the copper shock measure and its interaction with an indicator for private school status. Here, the strength of the first stage is a concern. Consequently, I present hypothesis test results based on weak IV-robust methods. Specifically, I use projection-based methods to construct the 95% confidence

[44] In municipality-level specifications, the full universe of non-special educational and non-adult students enrolled in a given market in each year is included in the sample. These specifications thus test whether the exclusion of schools that also serve special education or adult students and of schools that do not continuously enroll primary-level students lead school-level and market-level estimates to diverge.

set for the coefficient on each included endogenous variable and examine whether the confidence set allows for the null hypothesis of a zero-valued coefficient to be rejected at the 5% significance level. In addition, I present p-values associated with weak IV-robust tests of the jointly included endogenous variables for each specification. In the specifications in Columns (5)-(6) examining changes in the same school-level price outcomes included in Table 4, I identify increases in private school prices in response to rising local incomes that are statistically significant at the 5% level. These estimates are consistent in magnitude with reduced form estimates when those estimates are scaled up by the change in log mean income associated with a change in the copper shock measure. In Column (7), I correspondingly identify a significant increase in public school enrollment, while a 95% confidence set cannot be constructed for the coefficient characterizing the interaction between log mean household income and private school status. Notably, the p-value on the test for the joint significance of the included endogenous variables in Column (7) is far below 0.01.

6.3 Robustness Checks and Causal Mechanisms

Appendix Tables 3-6 confirm the robustness of benchmark findings to alternative variable definitions and specifications. Specifically, Appendix Table 3 incorporates a revenue measure that accounts for supplementary targeted voucher (SEP) revenues and adds a control for total municipality-level primary school enrollment to the school-level enrollment and revenue specifications from Table 4. To provide evidence that estimated responses to copper price-induced income shocks are not driven by compositional changes in the student population, Appendix Table 4 re-estimates Table 4 specifications for schools in municipalities in which the majority of the population is classified as urban (this includes 72% of municipalities and 78% of sample schools). For this subsample, there is no longer a significant (or large in magnitude) estimated change in average enrollment in response to copper price-induced income shocks, but the patterns of heterogeneous school price and enrollment responses appear quite similar to those found in Table 4. In Appendix Table 5, I replace the contemporaneous copper shock regressor with a lagged measure defined accordingly to confirm that estimates are not sensitive to this choice. To verify that heterogeneity in the relationship between copper shocks and local incomes that is correlated with market structure is not driving findings, Appendix Table 6 shows that the effect of copper shocks on local incomes does not vary by private school market density and that prior price and enrollment estimates are robust to the inclusion of year-by-municipality fixed effects.

If rising aggregate income induces changes in the number of schools in operation in particular municipalities, previous estimates may suffer from selection concerns. To address this issue, I show in Appendix Table 7 that positive copper price-induced income shocks do not have a consistent impact on the number of schools in operation in a given municipality.[45] In Appendix Table 8, I re-estimate Table 4 specifications, but replace the prior copper shock measure with an average shock that is calculated as the municipality-specific elasticity multiplied by the average of the current and first two lagged values of the annual copper price. If Appendix Table 8 point estimates are attenuated relative to Table 4, this would suggest that the enrollment and price effects associated with copper shocks are largely transitory, which could in turn explain the lack of school entry effects identified (as entry presumably entails a substantial fixed cost). In practice, Appendix Table 8 closely mirrors Table 4. This implies that alternative factors, such as the time required for school reputation-building, may contribute to the lack of observed school entry. These results also indicate that even short-run copper price fluctuations have significant effects on the quality of schools attended over a multi-year period.

There are multiple mechanisms that may drive the observed relationship between aggregate income and school prices and enrollment levels, and I next investigate the relative importance of alternative channels other than the demand for school quality that I have emphasized above. Seemingly, the most relevant alternative explanation for the link between aggregate income and school outcomes is that rising aggregate income affects funding for public schools and so affects qual-

[45] I present results in levels and using a Poisson QMLE specification to characterize percentage changes, and I show estimates for both contemporaneous and average shocks (constructed based on current, lagged and twice lagged shocks). Results are shown for the full sample as well as public school and private school subsamples. To account for the high degree of serial correlation in the dependent variable, regressions control for the lagged value of the dependent variable throughout. Estimates are generally negative and are not statistically significant, aside from two negative coefficients corresponding to short-run percentage-wise changes in the number of public schools and longer-run percentage-wise changes in the number of total schools. In any case, the negative point estimates from these specifications are opposite what would be found if rising income induced school entry.

ity of and demand for public schools. Previous research indicates, however, that municipal tax revenues do not significantly affect local educational expenditures in Chile (Auguste and Valenzuela, 2006). Consequently, changes in public school funding would likely come from national government sources. While vouchers are the primary mechanism through which the central government distributes funds, voucher values are determined nationally and I have shown that public school voucher revenues do not increase when incomes rise. The most important remaining source of local education funding is the National Fund for Regional Development (FNDR), which distributes funds to select municipalities to support their public schools. Appendix Table 9 estimates, which are inconsistent in sign and not statistically significant at conventional levels, do not provide evidence that FNDR funding increases with positive municipality-specific copper shocks.[46]

Rising aggregate income may also affect school outcomes by shifting schools' marginal cost curves. The largest component of a school's marginal cost is teacher incomes. Appendix Table 10 shows that changes in teacher hourly incomes in response to positive copper price-induced income shocks are small in magnitude and not statistically distinguishable from zero (teacher contracts are typically defined on an hourly basis). While teacher income changes cannot be estimated at the school level based on available data, Appendix Table 11 reveals that there are no differential increases in private schools in the share of experienced teachers or the share of teachers with college degrees. These characteristics are both positively correlated with teacher incomes.[47] Appendix Table 12 confirms that rental incomes also do not change significantly, which offers further evidence that schools' cost curves are not significantly affected by copper price-induced income shocks.[48]

Previous results indicate that fee-charging private voucher schools likely experience positive (though imprecisely-estimated) revenue changes when incomes are predicted to rise. Additional evidence suggests that these schools reduce costs, implying that profits increase. Specifically,

[46] Moreover, FNDR funding for all public schools in a municipality is only about 10% of the magnitude of corresponding voucher revenues.

[47] Public school teacher salaries are determined by a fixed schedule that incorporates prior experience and training, among other factors. In voucher private schools, administrators must provide a minimum level of compensation but otherwise maintain significant pay-setting discretion (Behrman et al., 2016; Mizala and Romaguera, 2005).

[48] Case study evidence indicates that capital costs, which include infrastructure costs, account for under 25% of average school costs (Ugarte and Williamson, 2012).

Appendix Table 13 examines changes in the number of classrooms in operation and the number of contracted teacher hours.[49] Estimates indicate that fee-charging private voucher schools marginally decrease the number of classrooms in operation when incomes are predicted to rise. If parents value peer quality (conditional on class size) and peer quality rises as enrollment falls, this may allow profit-maximizing schools to reduce costs while still partly capturing the increase in willingness to pay associated with reduced enrollment. In any case, these estimated decreases are small in magnitude (and marginally significant only for the most expensive schools), which explains the large declines in class size in fee-charging private schools shown in Column (10). At the same time, fee-charging private schools do significantly reduce the number of contracted teacher hours and the number of employed teachers, primarily by reducing part-time staff in non-core disciplines (according to available survey evidence). These findings indicate that fee-charging private schools reduce costs, and so increase profits, in response to copper price-induced positive income shocks.[50]

7 Test Scores and School Switching

This section presents student-level analyses that characterize heterogeneous student test score and enrollment responses to copper price-induced income shocks. Student-level estimates shed light on the distributional implications of changing private school enrollment patterns.

7.1 Test Scores

I have shown that positive aggregate income shocks induce movement from private to public schools, with effects driven by elite (high-priced) private schools. I next examine distributional

[49] Case study evidence suggests that personnel costs constitute 87% of total non-capital costs in Chilean schools, and so changes in staffing are the key determinant of whether schools can reduce costs in the short run (Ugarte and Williamson, 2012).

[50] In contrast, public schools significantly increase the number of classrooms when incomes are predicted to rise. As a result, estimated changes in class size in the public sector are positive but small in magnitude compared to net enrollment changes. Importantly, relative class size in the public (versus private) sector, a key determinant of willingness to pay for private school, grow substantially.

effects on academic performance. This analysis can inform our understanding of how large-scale private school markets mediate the welfare gains associated with rising incomes. I focus on normalized SIMCE test scores as a measure of educational achievement.[51] Column (1) of Table 7 estimates the overall impact of a positive copper shock on test scores using the following student-level specification:

$$T_{ismt} = \alpha + \beta * C_{mt} + \gamma_m + \lambda_t + \varepsilon_{ismt} \tag{9}$$

Here, T_{ismt} is the normalized average SIMCE score of student i who attends school s in year t and who resided in municipality m during the previous year (i.e., in year $t-1$). Column (1) reveals that a positive copper shock increases average test scores in a municipality by 0.056 standard deviations. To characterize those students most affected by copper price-induced income changes, 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 maternal education, which has the advantage of being both time-invariant and highly correlated with other measures of student socioeconomic status, such as household income. I define three maternal education terciles at the lagged municipality-by-grade-by-year level.[52] Additional details related to the construction of maternal education terciles are presented in the Data Appendix. To identify heterogeneous impacts based on maternal education, I estimate the following specification:

$$T_{ismt} = \alpha + \beta_1 * C_{mt} + \sum_{k=2}^{3} \beta_k * C_{mt} * X_{ik} + \gamma_{mk} + \lambda_t + \varepsilon_{ismt} \tag{10}$$

Here, X_{ik} is an indicator for whether student i is in maternal education tercile k and γ_{mk} represent municipality-by-maternal education tercile fixed effects. Column (2) of Table 7 estimates Equation (10) and reveals that students in the bottom maternal education tercile do not experience test score

51 SIMCE test scores have been used in many previous studies of the Chilean voucher system, including Hsieh and Urquiola (2006) and Gallego (2013). Scores are averaged across math and language results and are normalized by the 2005 mean and standard deviation.

52 I construct terciles based on lagged municipality of residence since contemporaneous municipality of residence could potentially respond endogenously to realized copper shocks. Municipality-specific terciles account for the fact that educational attainment varies starkly across geography. As an example, only 6% of mothers do not have a college education in Vitacura, one of the wealthiest municipalities in Santiago, while only 14% of mothers have any college schooling in Maria Pinto, one of the poorest municipalities in Santiago.

gains (the estimated impact is negative but imprecise). In contrast, students in maternal education terciles two and three experience test score gains of 0.08 and 0.19 SD, respectively, relative to students in the bottom tercile (although the tercile two estimate is somewhat imprecise).

Column (2) estimates reveal the combined impact of changes in willingness to pay for schooling and any changes in household resources available for non-tuition expenditures, such as tutoring or even food, which may affect test scores. Consequently, the finding of larger test score responses to income changes for higher maternal 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.[53] Importantly, as shown in Figure 4, copper shock-induced income changes are relatively uniform across the underlying distribution and so differential test score effects do not reflect differences in the magnitude of income changes. While I cannot definitively disentangle these two possible explanations for differential test score gains, I next present evidence on differential changes in school quality as a function of student background. These estimates align closely with test score estimates and suggest that school quality changes are an important driver of test score impact heterogeneity.

7.2 School Switching Patterns

While rates of elite private school attendance at baseline are rising in students' socioeconomic status, the marginal students who move between elite private schools and public schools and lower-quality private schools need not be representative of their classmates. To the extent that policymakers are concerned about inequality of educational opportunities, the characteristics of these marginal students determine changes in social stratification and so are of particular policy relevance. Moreover, given that the magnitude of test score changes are increasing in student socioeconomic status, evidence on changes in educational opportunity by student background can shed light on the appropriate interpretation of observed test score impacts.

[53] 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.

To characterize impacts on the quality of schools attended by different subgroups of students, Columns (3)-(4) of Table 7 examine changes in the baseline school price quartile of the school a student attends. Importantly, this outcome is not affected by changing student composition and is highly correlated with alternative measures of school quality (as shown in Appendix Table 1). Columns (3)-(4) reveal that students experiencing a positive copper shock attend schools of lower quality on average, with this effect driven by those students in the lowest maternal education bin.

The fact that increased enrollment in public and low-cost private schools is driven by low maternal education students suggests that these same students will likely experience relative declines in peer SES. To measure changes in peer SES and social stratification, I focus on school-level mean maternal education. Column (5) of Table 7 identifies a positive (0.054) and significant change in the average maternal education of schools attended by students experiencing a positive copper shock.[54] Column (6) re-estimates Equation (10) with the average maternal education dependent variable and reveals that low maternal education students do indeed experience significant relative decreases in school-level mean maternal education. Column (7) re-estimates Equation (7), which examines school-level heterogeneity, using mean maternal education as the dependent variable. The estimates reveal that average maternal education rises dramatically in the top two quartiles of the baseline price distribution, while effects in the bottom two quartiles are notably smaller.

In sum, these findings are consistent with test score estimates. The most disadvantaged students, whose test scores do not change significantly in response to income shocks, experience school quality changes likely to offset the direct effects of income changes. Specifically, when incomes rise, they experience the greatest degree of relative school downgrading, and they experience no change in average peer SES. In contrast, the highest-SES students experience the largest test score gains, attend school of the same (or higher) baseline quality, and experience the largest increases in peer SES in response to positive copper price-induced income shocks.

[54] The possibility of a positive overall effect on mean maternal education may initially appear puzzling given that the overall distribution of maternal educational attainment is not shifting in response to positive copper shocks experienced by municipalities. However, an increase in mean maternal educational attainment can be rationalized by strong ex-ante sorting across schools based on maternal education in combination with higher rates of school downgrading by those who would have been at the bottom of the maternal education distribution within their counterfactual schools.

8 Conclusions

This paper 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 and perceived peer quality are important components of perceived school quality. Simulations based on a stylized model of school pricing reveal that income increases can cause private school enrollment declines across a broad range of parameter values.

To investigate whether the conditions under which aggregate income shocks could 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 enrollment falls 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 expand enrollment to absorb those additional students who would have attended private voucher schools absent the rise in aggregate income.

In the analysis, I find that the effects of aggregate income increases on the achievement of lower-SES children are muted by the offsetting effects on their school quality. This would not be the case for increases in income concentrated at the lower end of the socioeconomic distribution, which would be expected to increase achievement because of both the direct effects of improved household resources and the indirect effects associated with making these children more desirable to private schools. The theoretical predictions I derive are also applicable to other settings in which schools set prices and exhibit market power and in which perceived school quality is related to characteristics of the student body. As such, the evidence I present can help us to understand how educational outcomes may be expected to respond to changing market conditions in a variety of alternative educational settings across the globe, including tertiary schooling markets in which reputational concerns are particularly salient and school differentiation is particularly stark.

References

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

Aguirre, J. (2017). Can Progressive Vouchers Help the Poor Benefit from School Choice? Evidence from the Chilean Voucher System. Working Paper.

Andrabi, T., J. Das, and A. Khwaja (2017). Report Cards: The Impact of Providing School and Child Test Scores on Educational Markets. *American Economic Review 107*(6), 1535–1563.

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.

Behrman, J., M. Tincani, P. Todd, and K. Wolpin (2016). Teacher quality in public and private schools under a voucher system: The case of chile. *Journal of Labor Economics 34*(2), 319–362.

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.

Feigenberg, B., S. Rivkin, and R. Yan (2018). Illusory Gains from Chile's Targeted School Voucher Experiment. *The Economic Journal,* forthcoming.

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 (2008). On the Determinants and Implications of School Choice: Semi-Structural Simulations for Chile. *Economía 9*(1), 197–244.

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

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. (2015). Improving Learning in Primary Schools of Developing Countries: A Meta-Analysis of Randomized Experiments. *Review of Educational Research 85*, 353–394.

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

Mizala, A. and P. Romaguera (2005). Teachers' Salary Structure and Incentives in Chile. In E. Vegas (Ed.), *Incentives to Improve Teaching: Lessons from Latin America*. The World Bank.

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 (2015). The Aggregate Effect of School Choice: Evidence from a Two-Stage Experiment in India. *Quarterly Journal of Economics 130*, 1011–1066.

Murphy, K. and R. Topel (1985). Estimation and Inference in Two- Step Econometric Models. *Journal of Business and Economic Statistics 3*(4), 370–379.

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.

Pagan, A. (1984). Econometric Issues in the Analysis of Regressions with Generated Regressors. *International Economic Review 25*(1), 221–247.

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.

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 (August 10, 2016). 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.

Figure 1: Model Simulations

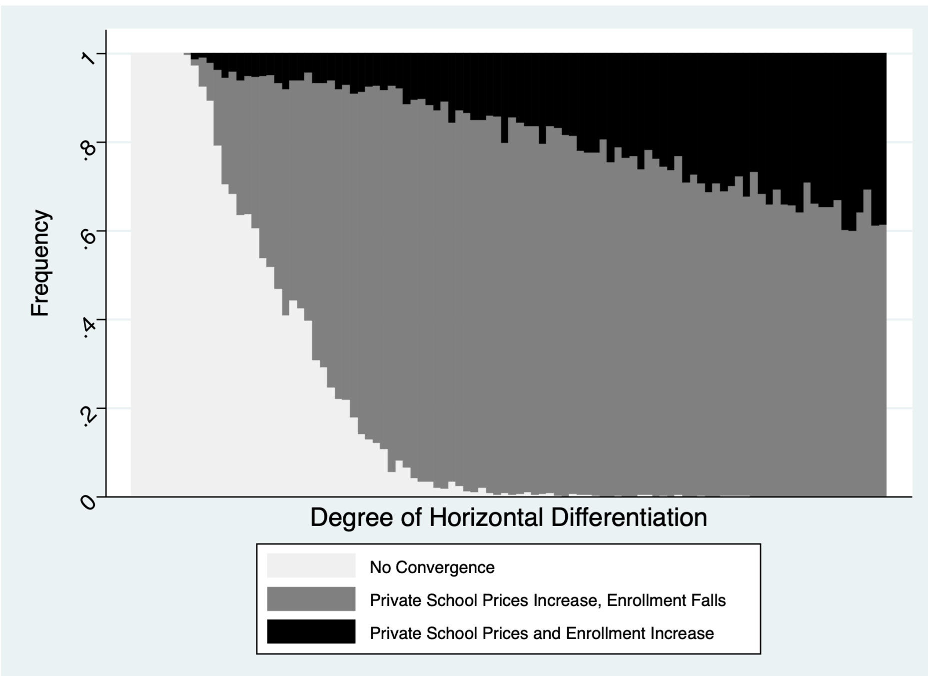

Notes: This figure presents simulation outcomes over 500 replications for each integer value of the horizontal differentiation parameter (σ) between the values of 1 and 100.

Figure 2: National Trends in Income, Enrollment, and Education Spending

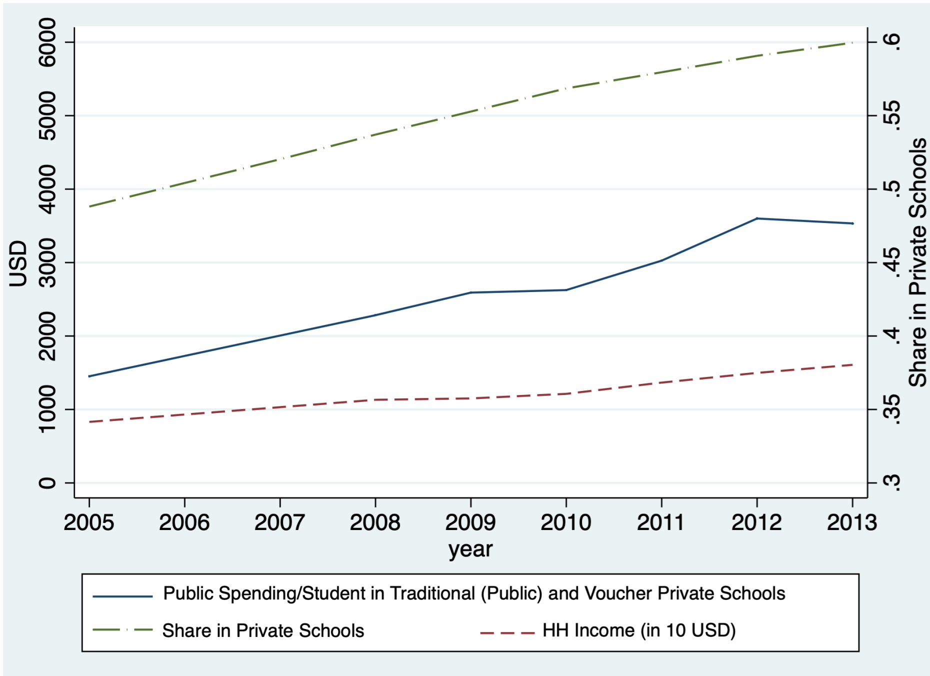

Notes: This figure presents annual changes in mean household income (measured in 10 USD), public spending per student in traditional (public) and voucher private schools (measured in USD), and fraction of students attending private voucher or non-voucher schools. Spending and income values are in real terms.

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

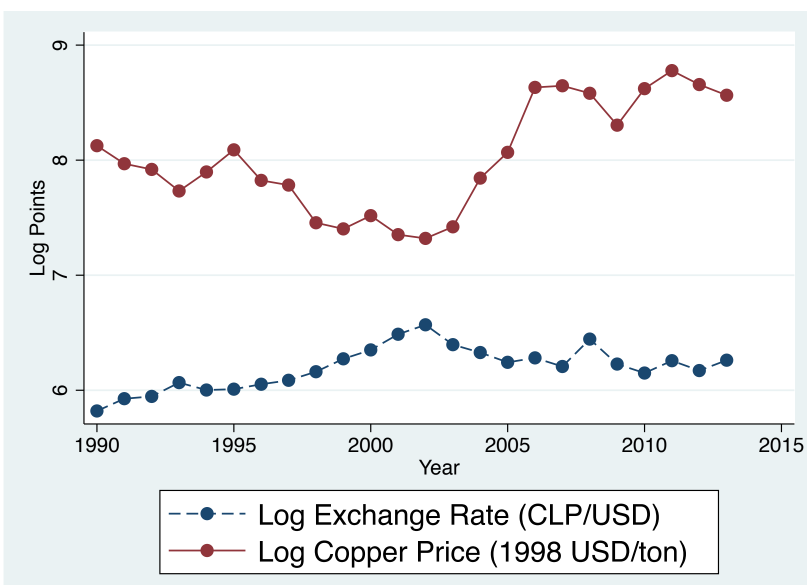

Figure 4: Copper Shock Impacts by Household Income Decile

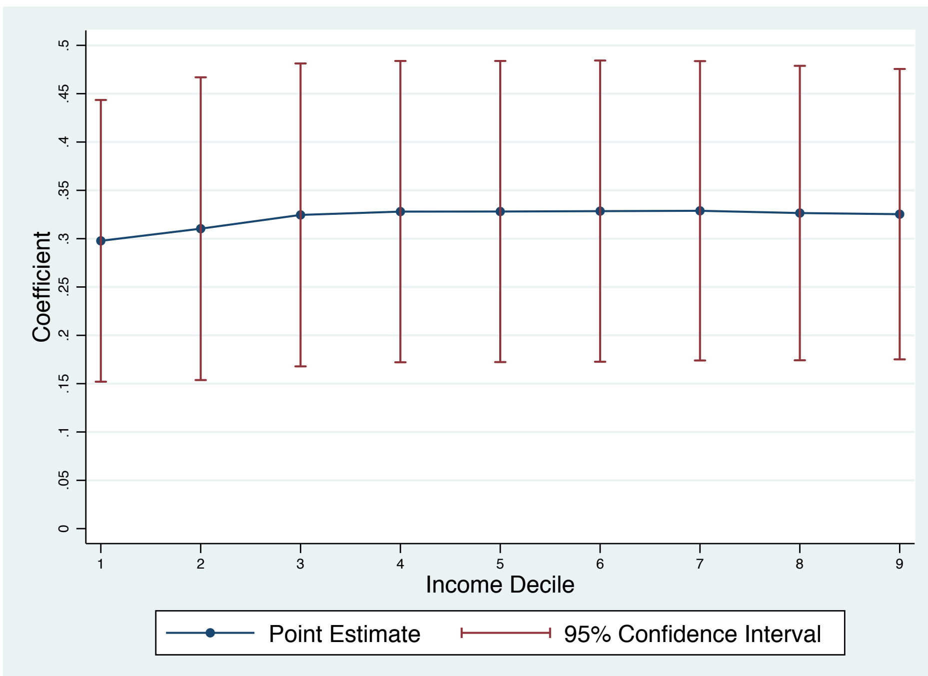

Notes: This figure presents point estimates and 95% confidence intervals from a Poisson QMLE specification at the income decile-by-year-by-municipality level that regresses household income on a set of interactions between the copper shock measure and decile-specific indicator variables. The model includes income decile, municipality and year fixed effects.

Table 1. School-Level Descriptive Statistics

	All Schools	Public Schools	No-Fee Private Voucher Schools	Fee-Charging Private Voucher Schools	Private non-Voucher Schools
	(1)	(2)	(3)	(4)	(5)
Number of Students per School	254.1	227.7	165.8	428.3	260.7
	(319.2)	(302.0)	(235.6)	(384.7)	(299.2)
Average Class Size per School	24.3	22.9	24.7	29.8	17.8
	(10.9)	(10.5)	(11.0)	(9.9)	(9.4)
Mean Maternal Education (Years)	9.6	8.4	8.5	12.5	15.8
	(3.0)	(2.0)	(2.6)	(1.7)	(1.0)
Mean SIMCE Score	-0.17	-0.30	-0.46	0.21	0.70
	(0.59)	(0.49)	(0.58)	(0.43)	(0.50)
Average Per Student Monthly Fee Charged	2527.9	0.0	0.0	13103.0	--
	(7104.8)	--	--	(11143.6)	--
Average Per-Student Base Monthly Voucher Value	36092.3	37186.2	36093.9	32566.1	0.0
	(7355.6)	(7759.7)	(6784.5)	(5150.9)	--
Number of Schools per Municipality	28.5	16.2	5.4	5.1	1.8
	(28.5)	(12.1)	(10.6)	(11.8)	(5.3)
Total Number of Schools	9,583	5,450 (56.9%)	1,818 (19.0%)	1726 (18.0%)	589 (6.1%)
Total Number of Students	2,435,215	1,241,105 (50.9%)	301,341 (12.4%)	739,216 (30.4%)	153,553 (6.3%)

Notes

1 Descriptive statistics (means and standard deviations) are constructed from the base year (2005) and observation numbers reflect the maximum number of observations for each school type. SIMCE Score is normalized at the student level to have a mean of zero and a standard deviation of one. Fee and Voucher Value measures are in 2005 CLP (1 USD=514 CLP in 2005). Information on Average Per Student Monthly Fee Charged is not available for Private non-Voucher Schools.

Table 2. Long-Run Income Changes and Municipality-level Enrollment/Price Responses

	Change in Fraction Attending Private School	Change in Fraction Attending Fee-Charging Private School	Change in Log Base Revenue per Student in Voucher Private Schools
	(1)	(2)	(3)
	Panel A: All Municipalities with Private Voucher Schools at Baseline (Unweighted)		
Change in Mean Household Income (%)	-0.045**	-0.045**	0.061***
	(0.020)	(0.019)	(0.016)
Observations	236	236	227
	Panel B: All Municipalities with Private Voucher Schools at Baseline (Weighted by Number of Schools at Baseline)		
Change in Mean Household Income (%)	-0.054***	-0.062***	0.063***
	(0.019)	(0.017)	(0.014)
Observations	236	236	227
	Panel C: 50 Largest Municipalities (Defined by Number of Schools at Baseline)		
Change in Mean Household Income (%)	-0.082**	-0.078**	0.096***
	(0.040)	(0.033)	(0.033)
Observations	50	50	50

Notes

1 Specifications are at the municipality level. All dependent variables are constructed based on differences between the start of the sample period (2005) and the end of the sample period (2013). Change in Mean Household Income (%) is constructed by taking an average from two end-of-period years (2011 and 2013), subtracting the average of two start-of-period years (2003 and 2006) and dividing by the average of the two start-of-period years. Revenue per Student measure is the sum of base voucher revenue and school fees and Change in Log Base Revenue per Student in Voucher Private Schools is constructed as the log of the (enrollment-weighted) mean in 2013 minus the log of the (enrollment-weighted) mean in 2005.

2 Change in Mean Household Income (%) and Change in Log Base Revenue per Student in Voucher Private Schools measures trim values below the 1st percentile and above the 99th percentile. Robust standard errors are presented in Columns (1)-(3). * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table 3. Municipality-Level Income Responses to Copper Shocks

	Log Mean Income	Log Mean Income	Log Mean Income	Log Mean Income	Log Mean Income	Log Mean Income
	(1)	(2)	(3)	(4)	(5)	(6)
Copper Shock	0.234***	0.159**	0.244***	0.200**	0.118***	0.123***
	(0.063)	(0.066)	(0.087)	(0.083)	(0.031)	(0.041)
Year Fixed Effects	X	X	X	X	X	X
Municipality Fixed Effects	X	X	X	X	X	X
Additional Controls		X		X		
Weighted			X	X		X
Linear Imputation					X	X
Mean of Dependent Variable	11.84	11.84	11.90	11.90	11.84	11.90
SD of Dependent Variable	[0.420]	[0.420]	[0.378]	[0.378]	[0.414]	[0.371]
Observations	1,030	1,030	1,030	1,030	2,307	2,307
Specification	Municipality-level		Municipality-level		Municipality-level	

Notes

1 Copper Shock is defined as the product of the municipality-specific elasticity of income with respect to copper prices and the log copper price (denominated in 1998 USD). Log Mean Income is denominated in 1998 CLP. Linear imputation refers to the imputation of the dependent variable for years in which the CASEN survey was not conducted. Weighted regressions are weighted by the baseline (2004) number of schools in the municipality. Additional controls include controls for the fraction of the adult population that is literate, the fraction of the adult population that is married, the fraction of households residing in rural areas, mean household size, the mean number of working-aged household members, and the interactions between each of these covariates and survey round fixed effects.

2 Standard errors are constructed by a two-stage bootstrap procedure, as described in text. * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table 4. School Price and Enrollment Responses to Municipality-Level Copper Shocks

	Predicted Log School Price (Base Voucher)		Realized Log School Price (Base Voucher)		Number of Students			Log Revenue		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Copper Shock	0.0023	0.0016	-0.0077	-0.0084	11.01**	27.32***	28.02***	-0.064***	-0.037*	-0.012
	(0.0038)	(0.0036)	(0.0062)	(0.0062)	(5.24)	(9.50)	(9.54)	(0.021)	(0.020)	(0.021)
Copper Shock* Private School	0.036***		0.038**			-38.43***				
	(0.012)		(0.015)			(13.91)				
Copper Shock* Baseline School Price in Quartile 2		0.0020		-0.0054			-18.68	-0.003	0.041	-0.039
		(0.0081)		(0.0081)			(11.58)	(0.033)	(0.056)	(0.037)
Copper Shock* Baseline School Price in Quartile 3		0.061***		0.079***			-55.69***	0.0028	0.092*	0.045
		(0.019)		(0.020)			(17.69)	(0.065)	(0.055)	(0.055)
Copper Shock* Baseline School Price in Quartile 4		0.116***		0.142***			-80.23***	0.099*	0.116	0.129**
		(0.035)		(0.044)			(25.51)	(0.055)	(0.080)	(0.056)
Year Fixed Effects	X	X	X	X	X	X	X	X	X	X
School Fixed Effects	X	X	X	X	X	X	X	X		
Municipality-by-Quartile FEs									X	X
Weighted by Number of Schools										X
Mean of Dependent Variable	9.77	9.77	9.59	9.59	244.5	244.5	244.5	17.17	19.49	20.26
SD of Dependent Variable	[0.25]	[0.25]	[0.28]	[0.28]	[298.6]	[298.6]	[298.6]	[1.40]	[1.29]	[1.16]
Observations	58,016	54,180	58,016	54,180	58,016	58,016	54,180	54,180	5,987	5,987
Specification	School-level				School-level			School-level	Municipality-by-Quartile level	

Notes

1 The dependent variable in Columns (1)-(2) is constructed based on average school ``top-up'' fee and the average grade-specific voucher value, adjusted based on the zone-specific multiplier provided by the Chilean Ministry of Education. In Columns (3)-(4), realized rather than average voucher values are used to construct the school price measure (realized voucher payments are a function of lagged attendance rates). Log revenue is constructed based on total realized school revenue (excluding revenue based on realized student/teacher performance). In Columns (9)-(10), the unit of observation is the municipality-by-baseline school price quartile-by-year and the dependent variable is the log of total revenues at this level of aggregation.

2 Standard errors are constructed by a two-stage bootstrap procedure, as described in text. * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table 5. Lead Structure of Copper Shocks

	Predicted Log School Price (Base Voucher)			Number of Students		
	(1)	(2)	(3)	(4)	(5)	(6)
Copper Shock	0.0023	0.0022	0.0023	27.32***	27.13***	27.25***
	(0.0038)	(0.0038)	(0.0038)	(9.39)	(9.38)	(9.37)
Copper Shock* Private School	0.036***	0.035***	0.036***	-38.41***	-38.27***	-38.38***
	(0.012)	(0.012)	(0.012)	(13.31)	(13.31)	(13.31)
Lead Copper Shock (t+1)	-0.0012			0.53		
	(0.0034)			(0.80)		
Lead Copper Shock (t+1)* Private School	-0.0061			-1.48		
	(0.0056)			(1.68)		
Lead Copper Shock (t+2)		-0.00079			-1.60	
		(0.0018)			(1.92)	
Lead Copper Shock (t+2)* Private School		-0.0027			1.29	
		(0.0046)			(3.06)	
Lead Copper Shock (t+3)			-0.0015			-1.71
			(0.0019)			(1.87)
Lead Copper Shock (t+3)* Private School			0.0016			1.35
			(0.0034)			(2.69)
Year Fixed Effects	X	X	X	X	X	X
School Fixed Effects	X	X	X	X	X	X
Mean of Dependent Variable	9.77	9.77	9.77	244.5	244.5	244.5
SD of Dependent Variable	[0.25]	[0.25]	[0.25]	[298.6]	[298.6]	[298.6]
Observations	58,016	58,016	58,016	58,016	58,016	58,016
Specification		School-level			School-level	

Notes

1 Lead Copper Shock coefficients are estimated separately for lead years 1-3. All specifications control for contemporaneous copper shock.

2 Standard errors are constructed by a two-stage bootstrap procedure, as described in text. * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table 6. Price and Enrollment Responses to Municipality-Level Income Shocks: IV Estimates

	Number of Students	# Non-Voucher Private School Students	Number of Public School Students		Predicted Log School Price (Base Voucher)	Realized Log School Price (Base Voucher)	Number of Students
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
			Panel A: Instrumental Variables Specifications				
Log Mean Income	2245**	-678*	4333***	2622***	0.042	-0.041	250.5**
	(1117)	(352)	(1404)	(923)			
P-value based on Weak IV-Robust 95% Confidence Set					--	--	*p<.05*
Log Mean Income*Private School					0.272**	0.304**	-389.8
P-value based on Weak IV-Robust 95% Confidence Set					*p<.05*	*p<.05*	--
Weak IV-Robust p-Value for Joint Test of Endogenous Variable(s)					*0.001*	*0.021*	*0.0002*
First Stage F-statistic	20.5	20.5	20.5	18.4	3.9	3.9	3.9
			Panel B: Reduced Form Specifications				
Copper Shock	265.1*	-80.1**	511.8***	296.2***	0.0026	-0.0060	27.6***
	(154.3)	(39.7)	(195.7)	(104.5)	(0.0041)	(0.0063)	(9.7)
Copper Shock*Private School					0.035***	0.036**	-38.8***
					(0.012)	(0.015)	(14.3)
Year Fixed Effects	X	X	X	X	X	X	X
Municipality Fixed Effects	X	X	X	X			
Control for Total Students				X			
School Fixed Effects					X	X	X
Mean of Dependent Variable	8250.7	590.5	3581.9	3581.9	9.77	9.59	244.7
SD of Dependent Variable	[11502.5]	[1929.8]	[3998.7]	[3998.7]	[0.25]	[0.28]	[298.7]
Observations	2,307	2,307	2,307	2,307	57,925	57,925	57,925
Specification	Municipality-level		Municipality-level		School-level		

Notes

1 In Columns (1)-(4), Copper Shock serves as an instrument for municipality-level Log Mean Income. In Columns (5)-(7), Copper Shock and Copper Shock*Private School serve as instruments for Log Mean Income and Log Mean Income*Private School. In Columns (5)-(7), the F statistic presented is the Kleibergen-Paap F-statistic. Copper shock is defined as the product of the municipality-specific elasticity of income with respect to copper prices and the log copper price (denominated in 1998 USD). All IV specifications rely on the linear imputation of municipality-level log mean income for years in which the CASEN survey was not conducted.

2 Standard errors clustered by municipality are presented in Columns (1)-(4) of Panel (A). In Columns (5)-(7), hypothesis testing employs weak IV-robust projection-based inference that generates coefficient-specific 95% confidence sets ("--" indicates that no such confidence set can be estimated). In Panel (B), standard errors are constructed by a two-stage bootstrap procedure, as described in text. * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table 7. Copper Shocks, Student Test Scores and Student Sorting Patterns

	Normalized Average SIMCE Score		Baseline School Price Quartile		School-Level Mean Maternal Education		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Copper Shock	0.056**	-0.021	-0.046**	-0.103***	0.054*	-0.056	0.099*
	(0.029)	(0.041)	(0.021)	(0.037)	(0.033)	(0.042)	(0.052)
Copper Shock*Maternal Education Tercile 2		0.079		0.076***		0.119**	
		(0.051)		(0.027)		(0.051)	
Copper Shock*Maternal Education Tercile 3		0.19**		0.145***		0.222***	
		(0.080)		(0.047)		(0.082)	
Copper Shock* Baseline School Price in Quartile 2							-0.029
							(0.11)
Copper Shock* Baseline School Price in Quartile 3							0.20**
							(0.087)
Copper Shock* Baseline School Price in Quartile 4							0.25**
							(0.12)
Year*Grade Level Fixed Effects	X	X	X	X	X	X	
Municipality Fixed Effects	X		X		X		
Municipality*Maternal Ed Tercile FEs		X		X		X	
Year Fixed Effects							X
School Fixed Effects							X
Mean of Dependent Variable			2.08	2.08	10.94	10.94	9.67
SD of Dependent Variable			[1.19]	[1.19]	[1.93]	[1.93]	[2.36]
Observations	1,660,902	1,538,865	15,378,722	12,538,404	15,609,989	12,737,959	53,918
Specification	Student-level		Student-level		Student-level		School-level

Notes

1 Copper shock is defined based on lagged municipality of residence. The dependent variable in Columns (1)-(2) is the average of the student's normalized fourth grade math and language scores. The dependent variable in Columns (3)-(4) is constructed based on the "top-up" fee charged by each school in its first year in operation. The dependent variable in Columns (5)-(7) is the average maternal years of schooling measured at the school-by-year level. Maternal education terciles are constructed at the year-by-lagged municipality of residence-by grade level.

2 Standard errors are constructed by a two-stage bootstrap procedure, as described in text. * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.