

Teacher Labor Markets, School Vouchers and Student Cognitive Achievement: Evidence from Chile

Michela M. Tincani¹

September 16, 2020

¹Tincani: University College London, 30 Gordon Street, London, WC1H0BE, UK, m.tincani@ucl.ac.uk. This paper was part of my Ph.D. dissertation at the University of Pennsylvania. I thank my advisers, Ken Wolpin, Petra Todd and Holger Sieg, for their guidance and encouragement. I am grateful to Dionissi Aliprantis, Orazio Attanasio, Jere Behrman, Alberto Bisin, Flavio Cunha, Mariacristina De Nardi, Jan Eeckhout, Hanming Fang, Nirav Mehta, Antonio Merlo, Aureo de Paula, Suphanit Piyapromdee, Shalini Roy and Giorgio Topa for helpful feedback. I also thank seminar participants at various institutions. I am grateful to the Chilean *Agencia de Calidad de la Educacion* for access to some of the data used in this research. The views reported here are those of the author and do not necessarily reflect views at the *Agencia*. Support from the Spencer Foundation and from the ERC (ERC-2015-CoG-682349) is gratefully acknowledged.

Abstract

I use administrative and survey data from Chile and a structural model to evaluate teacher policies in a market-based school system. The model accommodates equilibrium effects on parental sorting across school sectors (public or private), on the self-selection of individuals into teaching and across school sectors, and on teacher wages in private schools. I use the estimated model to simulate a reform that is planned to be implemented in Chile in 2023. Tying public school teacher wages to teacher skills and introducing minimum competency requirements for teaching is predicted to increase student test scores by 0.30 standard deviations and decrease the achievement gap between the poorest and richest 25 percent of students by a third. These impacts are ten times as large as the impact of a flat wage increase in public schools, and over twice as large as the impact of only introducing minimum competency requirements. The key driver of policy outcomes is an improvement in the pool of teachers, amplified by equilibrium effects on teacher wages in private schools. The equilibrium effects are large, accounting for 70 percent of estimated policy impacts.

1 Introduction

Teachers are one of the most important determinants of student achievement (Rivkin, Hanushek, and Kain 2005). Two key questions in education policy are how to attract good teachers into the teaching profession in a cost-effective way; and what impacts we can expect teacher policies to have on student achievement. The answers depend on the structure of the labor market for teachers, on how teacher quality combines with other inputs to produce achievement and, I argue in this paper, on the structure of the market for education.

I use administrative data from Chile and a structural model to empirically quantify the key forces behind teacher policy effectiveness in a market-based school system. Chile has a large-scale school voucher program and a large publicly subsidised private school sector. In this context, changes to teacher contracts in public schools generate equilibrium effects in private school wages and in the teacher-student match across school sectors. For example, suppose that teacher wages in the public sector become more tied to teacher skills. This changes the options available to individuals who are considering whether to teach, and in which school sector to teach. Private schools understand this and may decide to offer better wages to the best teachers, in response to the competitive pressure from public schools in the market for teachers. In turn, this equilibrium effect on wages feeds back into the labor supply decisions of potential teachers. Additionally, in an education system where parents choose with their feet, they may respond to the new allocation of teacher quality across sectors by changing their school choice. If the goal of policy evaluation is to estimate the impacts on student achievement, this parental response cannot be ignored.

I provide an estimable equilibrium model to analyse the response to teacher policies of private schools, potential teachers and parents. In counterfactual experiments using the estimated model, I quantify the relative importance of demand and supply side factors in determining teacher policy effectiveness. Finally, I provide an ex-ante evaluation of a merit-based teacher reform whose implementation is planned to complete in 2023, and compare it to a flat increase in public school wages which does not reward merit. I present predicted impacts on student achievement, and a welfare analysis of the reform. The advantage of estimating a structural model is that it can quantify equilibrium effects in impact and welfare calculations. By doing so, it identifies drivers of policy effectiveness that are not specific to the Chilean context, but apply more generally to any market-oriented school system.

In the model, parents choose between the municipal and the private subsidised school sectors. They care about consumption, which is lowered by tuition payments to private

schools, about their child’s achievement, and they have a direct preference for a school sector independent of its quality. Student achievement depends on student and household characteristics, on sector-specific school inputs, and on teacher quality. The latter is endogenously determined by the teacher labor supply, which is modeled as a Roy model, a workhorse theoretical framework in labor economics. Specifically, individuals with a college degree decide whether to teach, and if so in which sector, whether to work in the non-teaching sector, or whether to stay at home. They care about the wage and non-pecuniary aspects of the occupation.¹ In public schools, the wage is determined by rigid unionized wage formulae, while in private schools it is the product of the individual’s teaching skills and the price of those skills. The latter, together with tuition fees (up to a legal cap), is endogenously determined within the model by a profit-maximising representative private school. The rules that determine public school contracts are determined ex-ante and taken as given by all agents in the model. Therefore, the public sector can be thought of as the first mover in the model.

When evaluating how teacher policies affect teacher quality across schools, many existing studies rely on teacher value added models to estimate teacher quality. This framework typically restricts teacher quality to be additively separable from student ability in the production of test scores. This assumption is problematic in a system with large school choice, because the teacher-student match is a potentially important margin of policy response. If this match matters for test scores, ignoring this interaction would bias aggregate predictions on test scores. In this paper, I do not assume additive separability. As a result, I cannot use test score data alone to estimate teacher quality, like in a teacher value added approach. Instead, I identify teaching skills using the Roy model and the labor supply part of the data. The estimation algorithm then plugs the inferred teaching skills into the cognitive achievement production function, and uses the test score data to identify the complementarity between inferred teaching skills and student ability.² The novelty of this equilibrium approach is that the estimated parameters that determine teaching skills and the complementarity between teaching skills and student ability *simultaneously* rationalize the teacher choices and the parental choices as utility maximizing.

The model determines endogenously various equilibrium objects that can be matched to the micro-level data: the distribution of student achievement, the parental school sector choices and tuition payments, and accepted wages and occupational choices of potential

¹Non-pecuniary job characteristics have been found to be important determinants of teacher labor supply (Boyd, Lankford, Loeb, and Wyckoff (2005), Bonhomme, Jolivet, and Leuven (2016)).

²This approach is similar to studies of the impact of teacher wages on test scores (see, for example, Loeb and Page (2000)).

teachers. I use a number of Chilean datasets from 2006. Data on the pool of potential teachers, including their characteristics, occupational choices and wages, come from the CASEN (*Encuesta de Caracterización Socioeconómica Nacional*), a representative sample of all Chileans, and from the ELD (*Encuesta Longitudinal Docente*) a teacher survey. Data on students come from the SIMCE (*Sistema de Medición de Calidad de la Educación*), which provides administrative test scores and background information on 4th and 10th graders.³

The model is estimated on multiple markets. There are two steps. The first one uses the Method of Simulated Moments (McFadden 1989, Pakes and Pollard 1989) to estimate the technology, wage and preference parameters of households and of potential teachers, and to recover the skill prices that determine observed private school wages. The second step uses Nonparametric Simulated Maximum Likelihood (Laroque and Salanie 1989, Fermanian and Salanié 2004) to estimate the private school objective function parameters that rationalize the recovered skill prices as equilibrium prices. The good fit of the model helps build confidence about the lessons that we learn from the structural estimates and counterfactual experiments I perform.

The first set of results comes from the estimation of the structural parameters. First, the unionized public school wages overvalue degrees and certifications with respect to their impact on teaching skills. Second, outside options matter: worse teachers are found in markets with better non-teaching wages, and this is especially true for public sector teachers. Third, public schools attract less skilled teachers than private schools, confirming previous findings (Behrman, Tincani, Todd, and Wolpin 2016). Additionally, there is a pool of potentially highly skilled teachers not currently employed in teaching. Taken together, these results indicate that making public school wages more reflective of skills has the potential to improve not only the quality of teachers in public schools, but also the quality of the pool of teachers overall.

With regard to the demand for education side, parents care about education quality, but they also have a direct preference for the private voucher sector, independent of quality, that accounts for roughly a fifth of the enrolment share in private voucher schools. Finally, there is a student-teacher interaction in the production of test scores: the impact of teacher skills varies across students of different ability levels (unobserved types in the model). This empirical finding confirms the need for a policy evaluation that

³SIMCE administers each year standardized tests in Mathematics and Spanish that all students of selected grades are required to take. The school's average test results are published annually and parents can compare the performance of locally available schools. See Hastings and Weinstein (2008) for findings on the importance of information in parental school choice.

accommodates the reaction of both teacher and parental sorting across school sectors.

The first counterfactual experiment simulates the planned 2023 implementation of the merit-based reform of the teaching profession introduced by President Bachelet. Specifically, teacher wages in public schools are tied to their skills (so as to induce a 30 percent increase in wages on average), and individuals can enter the teaching profession only if their score on the PSU, the national university admission exam, is in the top 30 percent. The reform is predicted to increase test scores by 0.30 standard deviations (sd) on average, and to decrease the achievement gap between the richest and poorest 25 percent of students by a third.

The impact drivers are an improvement in the pool of teachers, a considerable reduction in the teacher quality gap between public and private voucher schools, and, to a lesser extent, the sorting of parents across school sectors. Specifically, the skills of teachers across school sectors are affected by two factors: first, merit-based pay in public schools attracts skilled individuals from outside of teaching and cream-skims the best teachers from the voucher sector. Second, private schools respond by increasing their wage offers, thus limiting their loss of teacher skills to public schools and also attracting skilled individuals from outside of teaching. These supply-side factors contribute more to policy impacts than parental sorting: they account for 70 percent of the treatment effects. A key lesson that we learn is that the existence of a large school choice program amplifies the positive impacts of merit reforms in the teacher labor market, thanks to equilibrium wage adjustments in the large non-public school sector which result in an even better self-selection of individuals into teaching compared to the selection that would have occurred in the absence of equilibrium effects.

The welfare analysis indicates that the reform increases the average utility of households, with larger gains accruing to the poorer households, who benefit from larger improvements in teacher quality in their schools of choice. Moreover, the policy pays for itself in the long run, because more highly skilled teachers in the public sector attract students back into public schools. This reallocates public expenditures in the form of vouchers back into the public sector, reducing the need for additional funds to cover public school running costs. Compared to the pre-reform system, the merit-based reform is more financially sustainable.

Finally, I simulate a flat increase to public sector wages that results in the same 30 percent average public sector wage increase as under the reform. The impacts on test scores and inequality are one tenth of those of the merit-based reform, because the flat bonus is not as successful at improving the pool of teachers. An experiment that introduces only

the competency requirement component of the merit-based reform achieves just below half of the predicted positive impacts of the reform. Therefore, both the wage and the minimum competency component of the planned merit-based reform are important to improve student outcomes.

The closest paper to this one is Behrman, Tincani, Todd, and Wolpin (2016), which focuses on the dynamics of the teacher labor supply in Chile, but abstracts from equilibrium effects and from the demand side of the market.⁴ By incorporating the demand side of the market, this paper can predict policy impacts on student test scores, the most relevant outcome for policy analysis.

This paper is related to an emerging literature using (static) market equilibrium frameworks to study education policy at the primary and secondary levels. Neilson (2017) models demand and supply of education quality to analyze the equilibrium implications of a Chilean targeted voucher reform. He models competition within the private sector, abstracting from competition across sectors. In this respect, it uses a complementary approach to this paper, which focuses on the response of private schools to competitive pressure from public schools in the market for teachers. While the objective of study in Neilson (2017) is different and, therefore, he does not model the teacher labor market, the conclusions regarding the drivers of policy impacts are similar: the largest margin of policy response comes from the supply side, with private schools improving their quality to attract larger voucher revenues. In particular, correlations between his estimates of school quality and observed measures of teacher quality suggest that it is precisely teacher quality that improves in response to policy, adding to the plausibility of the findings in this paper.

Another related stream of the literature analyzes the sorting of teachers between schools with different compensation regimes. Podgursky (2008) finds that private and charter schools in the U.S. set wages similarly to Chilean voucher schools, and in turn this allows them to hire better teachers, which is in line with my findings. Biasi (2018) develops a model of teacher demand and supply, and estimates it using data from Wisconsin.⁵ Like in this paper, she models teacher labor supply as a (static) Roy model. She uses the model to study teacher composition under unionized wage regimes and under

⁴ Behrman, Tincani, Todd, and Wolpin (2016) estimate large costs of switching occupational sectors in Chile, indicating that the static model of teacher labor supply is an acceptable approximation in this context. See Stinebrickner (2001a) and Stinebrickner (2001b) for dynamic structural models of teacher labor supply estimated on data from the United States, and Rothstein (2015) for simulations from a dynamic model of teacher labor supply using parameters calibrated to the United States.

⁵With respect to equilibrium studies of the teacher labor market, see also Boyd, Lankford, Loeb, and Wyckoff (2013), who provide a matching model of teachers to jobs

individual wage negotiations. My model differs in at least two ways. First, it includes the decision to enter the teaching profession. This allows me to study not only the stratification of teachers across sectors, but also how the composition of teachers reacts to changes in wage regimes. In line with previous work (Behrman, Tincani, Todd, and Wolpin 2016), I show that this is an important margin of response to major reforms. Second, teacher labor demand in my model is derived from the endogenous parental demand for education, which presents an elasticity to teacher skills. Analysing this elasticity is perhaps more relevant in a market-based school system with school choice, like the Chilean one, than in the one studied by Biasi (2018).

Finally, my paper is related to studies of performance pay. Dee and Wyckoff (2015) evaluate a performance pay intervention in Washington D.C. and find positive impacts on student achievement of incentives tied to teacher performance. Springer, Ballou, Hamilton, Le, Lockwood, McCaffrey, Pepper, and Stecher (2011) found no statistically significant impacts of financial bonuses tied to students' performance in Nashville.⁶ Sundararaman and Muralidharan (2011) performed a large scale field experiment in India that randomized performance pay for teachers, and found positive impacts on student achievement. These studies focus on the short-term impacts of performance pay driven by teacher effort (Mehta 2018), and do not analyze longer-term effects that may arise once teacher selection reacts to performance pay. The contribution of my paper to this literature is to study this margin of policy response.

The rest of the paper is organised as follows. Section 2 describes the Chilean school system and the data used in the analysis. Section 3 introduces the model, and section 4 discusses its key features and limitations. Section 5 describes the estimation technique and the identification strategy. Section 6 presents the model fit, and it is followed by the empirical results, in section 7. Section 8 concludes. The appendices follow.

2 Institutional Background and Data Description

2.1 Institutions

In 1981, Chile introduced a nationwide school voucher plan. Under the plan, each school-aged child receives a voucher that can be spent toward full coverage of tuition fees in

⁶Other studies analyze group-based incentives (see, for example, Imberman and Lovenheim (2015)), which are not directly comparable to the individual-based ones analyzed in my paper.

a municipal (public) school or coverage (partial or full) in a private subsidized school.⁷ The value of the voucher was CLP 27,391.903 (\sim \$50) per month in 2006, the sample year. The voucher cannot be used in private unsubsidized schools, from which this paper abstracts. These schools enroll 6% of students and cater to the wealthiest families. Private voucher schools in the sample year were allowed to charge a fee that exceeds the value of the voucher, up to a legal cap of CLP 54,018.768 per month (\sim \$100).

Some children are eligible for a *beca*, a fellowship for private education, that partially or fully covers the tuition fees in excess of the voucher. According to the SIMCE dataset, in 2006 around 60% of all Chilean children enrolled in private subsidized schools received a fellowship. As a result of government guidelines for fellowship assignment, children of lower socioeconomic status and from larger families are eligible for larger fellowships.⁸

Teachers' wages in the municipal sector are determined by rigid formulae that are negotiated between the government and the national teachers' union, the *Colegio de Profesores*. Wages are subject to seniority increments and other adjustments, such as allowances for professional training and for working in difficult conditions. Teacher assignment to schools is centralized nationally. Public schools, therefore, do not have control over the quality of the incoming pool of teachers. Teachers in private schools, on the other hand, are subject to the Private Labor Code, and their wages can be individually negotiated with private schools. Private schools are allowed to tie wages to teacher quality to attract a high-quality pool of teachers.

2.2 Data Description

2.2.1 Data Sources

I combine three data sources from 2006, the only year for which detailed information on both students and teachers in primary and secondary schools is available. I use the *Encuesta de Caracterización Socioeconómica Nacional* (CASEN) dataset to identify the pool of potential entrants into the teaching profession through a representative sample of individuals holding a college degree, a requirement for teaching.⁹ The CASEN survey is a

⁷Municipal schools can receive additional funds if the voucher revenues are not sufficient to cover their costs. Therefore, in practice they are not exclusively funded through the vouchers. This is fully taken into account in the Government cost calculations in the model simulations.

⁸The value of the voucher and the cap on private school tuition can be found in the *Decreto con Fuerza de Ley N° 2, De Educación, De 20.08.98* and in the law on shared financing, *Financiamiento Compartido, Ley N° 19.532*. The guidelines for fellowship assignment can be found in articles 24 and 27 of the *Ley de Subvenciones, Decreto con Fuerza de Ley N° 2, 20.08.98*.

⁹Individuals who want to become teachers must obtain a teaching certification. Although the teachers' statute, *Estatuto Docente, Ley N° 10.070*, allows for four ways to become certified to teach, according to

nationally representative survey of the general population from which I extract a sample of 3,520 individuals holding a college degree, tracking their occupational choices, accepted wages, and characteristics.

To augment the sample of teachers, I use a sample of 3,195 teachers from the *Encuesta Longitudinal Docente* (ELD) dataset. I extract from ELD the same set of individual characteristics obtained from CASEN, as well as the choice of school sector and accepted wages. From ELD and CASEN, I drop individuals who live in the remote Aisén and Tarapacá areas, for sample size reasons.

On the students' side, I randomly select a sample of 100,000 students from the *Sistema de Medición de Calidad de la Educación* (SIMCE) dataset, which contains information on all 4th and 10th graders in the country.¹⁰ The dataset contains administrative information on students' test scores in Mathematics and Spanish, used to measure achievement, as well as information on the students' household, tuition fee payments net of financial aid, and choice of school.

The model is estimated on 18 local labor and education markets. The market boundaries were determined so as to strike a balance between sample size within markets, number of markets, and market closed-ness (see Online Appendix C1). Markets are closed, with 98.8% of teachers working in the market in which they reside, and 99.0% of parents choosing a school in the market in which they reside. Nationally, the voucher sector accounts for 52.99% of student enrollment and 45.16% of teacher employment. However, there is across-market variation in these shares in part due to different local market conditions affecting demand for private education and teacher supply.

2.2.2 Descriptive Statistics

In private schools, there are students with higher socioeconomic status (SES) and less experienced teachers with higher measures of cognitive skills.

Children in the top 25 percent of the income distribution score, on average, 0.68

the 2006 teacher census (*Idoneidad Docentes*), 95% of all teachers (100% of all teachers in this paper's sample) get certified through one of two channels: i) a college degree in education, ii) a college degree in another area and a special degree in education (2-4 semesters). Importantly, anyone with a college degree can become a teacher, as long as they receive training in education if their college major was not education. Because in CASEN I do not observe the college major, in the model I let the non-pecuniary preference for teaching depend on an individual's unobserved characteristics. This captures, in a reduced form way, the fact that to accept an offer from the teaching sector, a college graduate without a major in education must pay the (financial and time) cost of obtaining training in education. Therefore, everything else being equal, the utility cost of teaching is higher for this individual.

¹⁰The sample size is approximately one third of the population size. Selecting a sample was necessary for computational tractability.

Table 1: Household Characteristics by School Sector

Household's characteristic	Public	Voucher	p-value difference
Avg parents' educ (yrs)	9.84	11.55	0.0000
Monthly income (CLP)	208,123	341,145	0.0000
Head of hh not working	7.86%	3.62%	0.0000
Size (n. individuals)	4.06	3.85	0.0000

Source: SIMCE 2006. 1 USD = 604.8 CLP. N=100,000.

standard deviations (sd) higher than children in the bottom 25 percent. There is also a sizable test score gap between public and voucher school students. The difference in test score means is equal to 0.36 sd, which is larger than the gap between charter and traditional public schools in the U.S. A third of this gap remains after controlling for student characteristics.

As documented also in previous studies (e.g., McEwan, Urquiola, Vegas, Fernandes, and Gallego (2008), Hsieh and Urquiola (2006), Urquiola (2005)), in the Chilean education system there is considerable school-sector stratification by students' SES. This is true also in my dataset. Table 1 shows average household characteristics by school sector. Parents of students in private subsidized schools earn 64 percent more than parents of students in municipal schools. Similar patterns are present among virtually all the household characteristics available in the data.

Teachers in the private voucher sector are on average 8.2 years younger and have 9.0 fewer years of teaching experience than teachers in the municipal sector. They score 0.27 standard deviations higher on the PAA test, the Chilean equivalent of the SAT, a measure of cognitive skills.¹¹

Private school wages correlate with cognitive skills, while public school wages do not. In the ELD dataset, an estimated panel data regression of log wages in public schools on teaching experience, teaching experience squared, nonteaching experience, and standardized PAA scores gives an insignificant coefficient on the PAA score (p-value=0.169). The same regression estimated for voucher schools indicates that a one standard deviation increase in the PAA score is associated with 4.0 percent higher wages in private schools

¹¹Private voucher school teachers also score 0.19 standard deviations higher on the AEP test (*Asignación Excelencia Pedagógica*), which evaluates teaching skills. However, taking the AEP test is voluntary and only 1.5 percent of teachers take it. Thus, this figure must be taken with caution. Further, to compare teachers to college graduates who are not teaching, back-of-the-envelope calculations using university entry exam scores standardized in the population of all exam takers indicate that Voucher school teachers have the highest entry exam scores (1.024), followed by Municipal school teachers (0.855), and followed by the pool of college graduates who either work in the non-teaching sector or stay at home (0.800).

Table 2: Average Monthly Teaching Wages by Teaching Experience and School Sector (2006 CLP)

Teaching Experience (years)	Public	Voucher	N	p-value difference
≤ 10	291,752	301,022	1,143	0.5494
11-20	344,566	387,305	1,033	0.0265
21-30	406,076	398,480	970	0.7668
31+	479,017	434,480	642	0.2196

Source: ELD 2006. 1 USD=604.8 CLP.

(p-value=0.009). Similar results have been reported in Bravo, Flores, Medrano, and Santiago (2010), who, additionally, show that teacher PAA scores are positively correlated with student test scores in Chile. This suggests that the individually negotiated private school wages reward teaching skills.

Wages of less experienced teachers are higher in private voucher schools: for example, Table 2 shows that teachers with experience between 11 and 20 years earn 12% more in voucher schools. The sign of this wage difference reverses, and the wage difference becomes statistically insignificant, for more experienced teachers.

Finally, non-teaching wages are on average 62.3% higher in the non-teaching sector for equally educated individuals. A college graduate employed in a non-teaching occupation earns monthly, on average, CLP 777,396 (\sim \$1,550), while a college graduate employed in teaching earns on average CLP 479,041 (\sim \$960). A wage difference persists at all ages, reaching peaks of over 80% for individuals younger than 45. In terms of hourly wages, the gap reduces to 18.7%, reflecting the fact that individuals in the non-teaching sector work more hours. Perhaps because of the higher flexibility of the teaching time schedule, around 70 percent of teachers in Chile are women.

2.3 Reduced-form evidence

2.3.1 Labor Supply

In a related paper (Behrman, Tincani, Todd, and Wolpin 2016) we show that the existence of the private voucher sector in Chile draws higher-productivity individuals into the teaching profession.¹² Therefore, the choice to enter the teaching profession is an important margin to study, in addition to the sorting across school sectors within teaching.

¹²Productivity is measured by the model's unobserved types that receive higher wage offers. A higher proportion of these types is drawn into teaching (in either sector) compared to a scenario in which the Voucher sector does not exist.

The CASEN dataset allows me to identify the pool of college graduates and examine the self-selection into teaching and non-teaching occupations.¹³

As shown in Table C2 in Online Appendix C2, females are 25 percentage points (p.p.) less likely to work, however, conditional on working, they are 47 p.p. more likely to choose teaching over non-teaching. Females with children are less likely to work than females without children and than men with children, however, conditional on working, they are more likely to work in teaching. Given that females represent 70 percent of the teacher labor force in Chile, this indicates that it is important to jointly study the decision to enter the labor force and to enter the teaching profession.

Focusing on the choice of school sector within teaching, older individuals are more likely to work in public schools, likely because of the steep returns to seniority in the public sector. There are no other significant differences in terms of observed individual characteristics. However, there are significant differences across markets. Probit regressions augmented with market dummies indicate that the residual difference in the probability of choosing teaching conditional on working ranges across markets from -22 p.p to +18 p.p.. This may reflect differences across markets in the relative remuneration of teaching and non-teaching jobs, due to differences in local labor market conditions as well as in the demand for teachers. The structural model disentangles teacher labor demand from teacher labor supply factors.

2.3.2 Wages

Wage regressions indicate substantially larger unexplained wage variation in private schools than in public schools. The residual variance is 22.5 percent higher in private schools, where individual negotiations limit the wage compression observed in public schools. In general, unobserved heterogeneity plays an important role in explaining wage variation. For this reason, it is incorporated into the structural model.

Graduate degrees and certifications appear to be valued more in voucher schools (and in non-teaching) than in public schools according to the reduced-form estimates in Online Appendix Table C3. This contradicts what would be expected, given that the rigid negotiated formulae in public schools explicitly reward professional training, while the individually negotiated wages in private schools (and in the non-teaching sector) need not.

¹³Studying the selection into teaching is a point of departure from Biasi (2018), who studies sorting of teachers across schools under different counterfactual experiments abstracting from teacher entry. She finds that a large scale individual wage negotiation policy for teachers would not improve teacher quality, but acknowledges that this result may change in a model that allows for teacher entry.

However, these reduced-form coefficients do not necessarily reflect true wage offer functions. The structural model aims to uncover the true wage offer functions by accounting for unobserved heterogeneity and self-selection into occupations. Indeed, the structural wage parameters are of the expected signs.

Age, a (noisy) proxy for experience, is positively correlated with wages in both teaching sectors. In public schools, this reflects the explicit rewards to seniority. In private schools, this may reflect a positive correlation of experience with skills. While some studies suggest that only the first few years of experience are correlated with teaching skills, Wiswall (2013) finds that experience improves teaching skills even in later career stages.

Estimates in Table C3 in Online Appendix C2 seem also to indicate that females' log-wages are disproportionately penalised in the non-teaching sector as compared to the two teaching sectors, with a negative dummy coefficient (-0.38) that is three times as large as in teaching (-0.10 and -0.14). However, these coefficients may be biased because of self-selection. The structural model explicitly accounts for self-selection into occupations that require different but potentially correlated skills to examine whether there exists a penalisation for women in wage offers, or whether other non-pecuniary factors explain the higher propensity of women to choose teaching.

Finally, like with occupational choices, there is substantial residual variation in sector-specific wages across markets. For example, market fixed effects in log-wage regressions in voucher schools range from -0.39 to $+0.23$. This could reflect different patterns of self-selection into private school teaching, and/or different wage offer functions. Because the wage offered in private schools is an equilibrium object that depends on the alternatives available to potential teachers and on the demand for private school education in the local market, variation in both labor and education market conditions imply variation in wage offers across markets. The reduced-form parameters are uninformative on the distinct contributions of these sources to wage variation. One of the goals of the structural model is to disentangle how much of the residual variation in wages across markets is due to variation in self-selection and in skill prices (i.e., wage offers).

2.3.3 Demand for Public vs. Private Education

Confirming the well documented school sector stratification in Chile, Table C4 in Online Appendix C2 shows that parents with higher income and with higher education are more likely to choose the voucher sector. The public sector is selected more often at the primary level and in rural areas. In terms of variation across markets, probit regressions with

market dummies indicate that the average (residual) probability of choosing the public sector ranges across markets from -22 p.p. to $+25$ p.p.. This may indicate different relative qualities of public versus private education due to, for example, differences across markets in relative teacher quality or in other aspects of school quality. One of the goals of the structural model is to estimate the elasticity of the parental school sector choice to teacher quality.

2.3.4 Achievement

Table C5 in Online Appendix C2 shows reduced-form estimates of the production of achievement as a function of individual student characteristics. First, if observed and unobserved student characteristics are correlated within sectors, these reduced-form coefficients are biased. Second, this simple model leaves substantial unexplained variation in outcomes. By accounting for unobservables of both teachers and students, the structural model identifies how variation in teacher and school quality and in student unobservables contribute to this unexplained variation in achievement.

3 Model

There are two stages in the model. In the first stage, wage rates (i.e., skill prices) and tuition fees in the private school sector are determined. In the second stage, the demand for public and private education and the supply of teachers to the two school sectors are determined. The model endogenously determines wages and fees in private schools, the supply of teachers of various skill levels across sectors, the allocation of students across sectors (and the resulting financial aid received by students in voucher schools), student achievement in both sectors, and Government costs.

Importantly, the private schools, parents, and individuals making occupational decisions (i.e., college graduates) take the public school sector policies as given. These are: the unionized teacher wages, the tuition cap for private school fees, and subsidies to households for private education. One could think of a stage zero in the model where such policies are determined. The counterfactual experiments simulate various public sector policies in stage zero, and the endogenous responses in stages one and two of all other agents in the model (potential teachers, parents and private schools).

The economy is comprised of M closed markets, in each market the model is solved separately.

3.1 First Stage: Tuition Fees and Wage Determination in the Private School Sector

There is a representative private school in each market which, following the existing literature, is assumed to maximise profits.¹⁴ The private school takes as given wage offer functions in the public school sector and in the non-teaching sector, and it chooses teacher wages and tuition fees to maximize profits. Tuition fees p are subject to a legal cap \bar{p} . Wages are of the standard linear pricing type (Ben-Porath 1967): the private school chooses a wage rate/skill price r such that if a teacher possesses s_i units of teaching skills, his/her offered wage is equal to rs_i . The private school acts as a monopolist and monopsonist. This assumption is discussed in Section 3.3.1.

In the first stage, the private school anticipates the behavior, in the second stage, of parents and potential teachers, who are price-takers. Therefore, it takes as given the student enrollment and teacher supply functions: $E(p, r; v)$, $T(p, r; v)$ and $NT(p, r; v)$, where $E(p, r; v)$ is the mass of students who enroll in the private voucher school, $T(p, r; v)$ are the total teaching skills supplied to the school, and $NT(p, r; v)$ is the mass of potential teachers who choose the private school. These demand and supply functions depend on prices p and r , they are indexed by v , the voucher amount, and they depend also on other variables omitted to simplify notation (for example, on the characteristics of the other options available to parents and to potential teachers). As will be clear in the discussion of the second-stage of the model, enrollment depends on the rental rate r because parents care about the skills of the teachers in private schools, which is a function of r . On the other hand, thanks to the assumption that teachers' utility does not directly depend on student identity (see Section 3.3), the expressions for teacher labor supply simplify to $T(r)$ and $NT(r)$.

The profit function depends on school revenues and costs. Intuitively, revenues per student should be determined simply by the per-student tuition fees, p , and total revenues are the product of per-student revenues and enrollment, $p \cdot E(p, r; v)$. In terms of accounting, the school should receive the voucher v from the Government, and the difference between tuition p and the voucher v from the family. However, there are two adjustments to per-pupil revenues that are mandated by law and that concern fellowships

¹⁴Existing formal models of the Chilean market for education are in Urquiola and Verhoogen (2009) and Neilson (2017). Both studies assume that private schools are profit maximizing. Using data from the Ministry of Education, Elacqua (2006) documents that 73.60% of all Chilean voucher schools are for-profit. Finally, also McEwan, Urquiola, Vegas, Fernandes, and Gallego (2008) identify Chilean voucher schools as mostly profit maximizing.

for disadvantaged students and voucher amount adjustments.¹⁵ First, private schools are required to offer fellowships to eligible students towards coverage of the portion of the tuition fee p that exceed the per capita voucher v . These fellowships are co-financed by the private schools and the Government. Since the private school contributes to financing these fellowships, the per-pupil revenues are lower than they would be if the fellowships were entirely financed by the Government. The private school's contribution to the financial aid budget depends on its student composition, and it is determined by a Government formula. Second, the amount of voucher subsidy v received by private schools is adjusted according to a formula that penalises higher tuition charges. This adjustment reduces the rate at which the gross per-capita revenues increase as per-capita tuition fees increase. I incorporate both Government formulae into the calculation of profits and denote the adjusted total revenues by $\tilde{R}(p, r, E(p, r; v))$. The exact formulae can be found in Online Appendix C3. Total costs are determined by the sum of a variable cost quadratic in enrolment, the total teacher wage bill, and other operating costs. Formally, the problem of the private subsidised school is:

$$\begin{aligned}
\max_{(p,r)} \Pi = & \underbrace{\tilde{R}(p, r, E(p, r; v))}_{\text{Adjusted Revenues}} - \underbrace{((c_1 + \epsilon_{cost})E(p, r; v) + c_2E(p, r; v)^2)}_{\text{Variable Cost}} \\
& - \underbrace{rT(r)}_{\text{Teacher Wage Bill}} - \underbrace{c_3 \frac{E(p, r; v)/NT(r)}{45}}_{\text{N Classes per Teacher}} - \underbrace{OC}_{\text{Fixed Operating Costs}} \\
\text{s.t.} [& E(p, r; v) \quad T(r) \quad NT(r)]' = [E^*(p, r; v) \quad T^*(r) \quad NT^*(r)]' \quad \forall p, r \\
& p \leq \bar{p}.
\end{aligned}$$

The first constraint indicates that schools take as given the teacher labour supply function and the student enrolment function which result from utility maximizing behavior in the second stage of the model. The variable cost is subject to a shock ϵ_{cost} which is distributed according to a truncated log-normal distribution with parameters 0 and σ_{cost}^2 , where the truncation guarantees that profits are non-negative, and that there is a private school in the market.¹⁶ The total wage bill is the total amount of teaching skills supplied to the sector, $T(r)$, multiplied by the unit price of those skills, r . The term $c_3 \frac{E(p, r; v)/NT(r)}{45}$, where $NT(r)$ is the number of teachers, is a cost proportional to the minimum average number of classes taught by the same teacher. This term allows the model to capture a key data feature: schools do not hire very few highly skilled individuals and assign them to a large

¹⁵See Article 25 of the *Decreto con Fuerza de Ley N° 2, De Educacion, de 20.08.98*.

¹⁶I choose a log-normal distribution to restrict the cost shock to be positive. The parameter c_1 is restricted to be non-negative. The fixed operating costs are normalised to zero in estimation.

number of students.¹⁷ Finally, if $E = 0$ or $NT = 0$, the private school does not operate.

At the estimated parameter value the profit function is well approximated ($R^2 = 0.96$) by a polynomial that admits only one maximum. Refer to Online Appendix C4 for details of this approximation.

3.2 Second Stage: Demand for Education and Supply of Teachers

In the second stage, parents choose a school sector and individuals with a college degree make labour supply decisions to maximize their utility. Parents choose between the public/municipal (M) and the private voucher (V) sector. Individuals with a college degree choose between private schools teaching (V), public school teaching (M), the non-teaching sector (NT), and home production (H).

3.2.1 Demand for Education

Parents, indexed by h , differ in terms of characteristics that are observed and unobserved to the econometrician. In estimation, the unobserved characteristics are modeled as types, in the spirit of Heckman and Singer (1984) and Keane and Wolpin (1997). A household's type k_h can be one of K , with type population proportions given by π_1, \dots, π_K . The type captures the match of a student's unobserved ability with a school (both in terms of the constant and the coefficients in the achievement production), and the parental direct preference for a school sector.

Parents care about consumption and their child's achievement. Moreover, they have a direct preference for a school sector that is independent of its effect on student achievement. For example, parents may sort on the basis of other amenities that aren't positively correlated with educational gains (Rothstein 2006). Formally, the choice-specific utilities of household h in market m are:

¹⁷The legal cap on class size is 45 in Chile, so if there are x students per teacher, those students must be split into at least $x/45$ classes. An alternative way to match the student:teacher ratio in the data is to assume that either achievement or teachers' utility or both depend on class size in the second stage of the model. However, these modelling options introduce social interactions and a fixed point problem in the second stage that would substantially complicate the numerical tractability of the model.

$$\begin{aligned}
u_{hmM} &= \tau(k_h)\ln(c_hM) + a_{hmM} + \eta(k_h) + \eta_1\text{primary}_h + \eta_2\text{rural}_h + \nu_{hm}^{pref} & (1) \\
u_{hmV} &= \tau(k_h)\ln(c_hV) + a_{hmV},
\end{aligned}$$

where k_h is the household's type, a_{hmj} is achievement in sector $j = M, V$, and the η parameters and idiosyncratic ν_{hm}^{pref} shock capture the direct sectoral preferences.¹⁸ To capture key features of the data, the latter vary by the education level of the child (primary or secondary), and by whether the family resides in a rural or urban area. The shock ν_{hm}^{pref} is a preference shock distributed as $N(0, \sigma_{pref}^2)$. As is standard in discrete choice models, some normalizations are required. The coefficient on achievement is normalized to one because it is not separately identified from σ_{pref}^2 . Moreover, because only the difference in utilities across choices is identified, in the voucher sector the direct preferences are normalized to zero and the shock is normalized to be a degenerate random variable equal to zero.

The utility from consumption is equal to $\tau(k_h)\ln(c_hj)$, where $\tau(k_h)$ is a parameter that measures the trade-off between consumption and child achievement, and it determines parental willingness to pay for private education. Consumption is equal to household income Y_h if parents select a free public school, and it is equal to income net of tuition payments if they choose a private school. Tuition payments are given by the tuition charged by the school p , minus the voucher subsidy v , minus a fellowship f if the student is eligible for one according to Government guidelines. Formally:

$$c_{hj} = \begin{cases} Y_h & \text{if } j=M \\ Y_h - (p - v - f(Z_h)) & \text{if } j=V \end{cases}$$

where Z_h are household characteristics determining the amount of fellowship the student is eligible for (which can be anywhere between 0 and $p - v$). The fellowship formula is:

$$f(Z_h) = b_o + b_1\text{primary}_h + b_2\text{fam_size}_h + b_3\text{rural}_h + b_4Y_h$$

where fam_size_h is the family size.¹⁹ To account for the fact that parents in the sample

¹⁸The shock is uncorrelated with achievement gains, while the constant $\eta(k_h)$ depends on gains due to its dependence on the type. This formulation allows for the direct sectoral preference to be correlated with achievement gains. This captures the idea that the direct preference of parents for a school sector and the ability of their child can be correlated.

¹⁹In estimation, I allow for the fellowship to be measured with error: $\tilde{f} = f + me$ with $me \sim N(0, \sigma_{me}^2)$.

are never observed choosing the private sector when their income is lower than the tuition fees, I assume that parents are credit constrained. The utility from the voucher school is equal to $-\infty$ when tuition is above income, and it is never chosen.

Student achievement is determined by student characteristics, including type, and by the expected skills of the teachers in the school sector (\bar{s}_{mj}).²⁰ There is no distinction between a household's and a student's type because their distributions would not be separately identified. Formally, for $j = M, V$:

$$a_{hmj} = \beta_{0j}(k_h) + \beta_{1j}(k_h)\bar{s}_{mj} + \beta_{2j}(k_h)peduc_h + \beta_{3j}(k_h)y_h + \beta_{4j}y_h^2 + \nu_{hj} \quad (2)$$

where $peduc_h$ is parental education in years (average between mother's and father's education) and y_h is household monthly income divided by household size. The productivity shocks ν_{hM} and ν_{hV} are distributed as independent mean-zero random variables with variances $\sigma_{\nu_{hM}}^2$ and $\sigma_{\nu_{hV}}^2$. They are independent of the preference shock ν_{hm}^{pref} , therefore, correlation between unobserved student ability and sectoral choice is captured by the type. The impact of parental characteristics on achievement may vary by household type and by school type, to allow for different propensities of observationally identical households to invest in educational inputs, and to capture sector-specific educational expenditures.

Household h chooses the utility-maximizing sector: $e(\tilde{X}_h, k_h, \nu_h) \in \{M, V\}$, where $\nu_h = [\nu_{hM}, \nu_{hV}, \nu_{hm}^{pref}]$ is the vector of preference and technology shocks, k_h is the household's type, and \tilde{X}_h includes all household characteristics that are observed by the researcher and that affect the household's utility and fellowship.

3.2.2 Supply of Teachers

Individuals who make labor supply decisions, indexed by i , differ in terms of characteristics that are observed and unobserved by the econometrician. Like in the demand side of the model, in estimation the unobserved characteristics are modeled as types. An individual's type l_i can be one of L , with type population proportions given by ψ_1, \dots, ψ_L .

The labour supply part of the model is a Roy model of occupational choice augmented with a non-work option and with non-pecuniary preferences. First, the non-work option is important to capture the labour supply decisions of individuals with a high utility from not participating in the labour market, like, for example, women with young children. This is especially relevant in the study of the labour market for teachers in Chile,

This shock is independent of all other model shocks.

²⁰Assuming that parents hold rational expectations, the expectation of teaching skills formed by parents is equal to the true equilibrium mean teaching skills.

given the high share of women in the teaching profession. Second, it is important to consider non-pecuniary preferences because the teaching profession differs from non-teaching occupations in terms of job attributes such as flexibility of working arrangements, and because individuals may have an idiosyncratic taste or distaste for teaching. Moreover, within teaching, the private and public sectors differ in terms of non-pecuniary attributes such as job security.²¹ Formally, the choice-specific utilities of individual i in market m from each choice $j \in \{M, V, NT, H\}$ are:

$$\begin{aligned}
u_{imM} &= \ln(w_{imM}) + \mu_{0M}(l_i) + \mu_{0Teach}female_i \\
u_{imV} &= \ln(w_{imV}) + \mu_{0V}(l_i) + \mu_{0Teach}female_i \\
u_{imNT} &= \ln(w_{imNT}) \\
u_{imH} &= \mu_{0H}(l_i) + \mu_1female_i + \mu_2female_i * nk_i + \\
&\quad \mu_3age_i + \mu_4nk_i + \mu_5nk0 - 2_i + \mu_6nk3 - 6_i + \mu_7age_i^2 + \epsilon_{iH}^{pref}
\end{aligned} \tag{3}$$

where w_{imj} is the wage offer from sector j to individual i in market m , the μ parameters capture choice and type specific non-pecuniary preferences, $female_i$ is a gender dummy, $nk_i, nk0 - 2_i, nk3 - 6_i$ are, respectively the number of children, the number of children aged [0-2] and the number of children aged [3,6], age_i is age, and $\epsilon_{iH}^{pref} \sim N(0, \sigma_H^2)$ is a preference shock to the home option. The non-pecuniary term for the non-teaching sector has been normalized to zero because only the differences in non-pecuniary values are identified.

Each individual, including those not observed teaching in the data, is endowed with a certain level of teaching skills, s_i , which would raise the achievement of students if they chose to teach. Teaching skills are a function of an individual's characteristics (both observed and unobserved by the econometrician):

$$s_i = \exp(a_0(l_i) + a_1'X_i + \epsilon_i^{tech}) \tag{4}$$

where ϵ_i^{tech} is a technological shock, and $X_i = [age_i, age_i^2, female_i, cert_i, grad_deg_i]$ where $cert_i$ is a dummy for whether the individual holds professional certificates and $grad_deg_i$ is a dummy for graduate degrees (master's or Ph.D.).

²¹In a related paper (Behrman, Tincani, Todd, and Wolpin 2016), we fully specify a dynamic programming discrete choice model of teacher labor supply, including different future lay-off probabilities. In this paper, the non-pecuniary terms capture in a reduced form way the present discounted value of an occupational choice, including features of the dynamic environment such as lay-off probability.

The wage offers from the three work options depend on the individual's characteristics. Wage offers in public schools are determined by rigid governmental formulae, which are mostly seniority-based with some additional adjustments. Private school wages are a linear function of skills, as in the standard Ben-Porath/Roy model framework (Ben-Porath 1967, Roy 1951). Formally, wage offers in the three working sectors are:

$$\begin{aligned} w_{imM} &= \exp(\alpha_{0mM}(l_i) + \alpha'_M X_i + \epsilon_{iM}) \\ w_{imNT} &= \exp(\alpha_{0mNT}(l_i) + \alpha'_{NT} X_i + \epsilon_{iNT}) \\ w_{imV} &= r_m \exp(a_0(l_i) + a'_1 X_i + \epsilon_i^{tech}) \end{aligned} \tag{5}$$

where r_m is the price of teaching skills in market m (the wage rate), determined in the first stage of the model. The constant in the public school log-wage equation depends on an individual's type l_i to capture those individual characteristics entering the rigid wage formulae that are not observed in the data. The wage shocks $\epsilon_i = [\epsilon_{iM}, \epsilon_i^{tech}, \epsilon_{iNT}]'$ are i.i.d., independent of the preference shock, and distributed as $N(0, \Sigma)$, where Σ is a diagonal matrix with elements $\sigma_M^2, \sigma_V^2, \sigma_{NT}^2$.²² Non-teaching skills are not assumed to be identical to teaching skills, nor do I impose any correlation between them. Any correlation between teaching and non-teaching skills depends on the distribution of types and the values of the type-specific constants in the teaching skill and non-teaching wage equations.²³ These correlations determine how the wage elasticity of the teacher labor supply is affected by non-teaching opportunities.²⁴

Potential teacher i chooses the utility-maximizing occupation: $d(\tilde{X}_i, \psi_i, \epsilon_i) \in \{M, V, NT, H\}$, where $\epsilon_i = [\epsilon_{iH}^{pref}, \epsilon_i^{tech}, \epsilon_{iM}, \epsilon_{iNT}]$ is the vector of preference, technology and wage shocks, ψ_i is the potential teacher's type, and X_i includes all potential teacher's characteristics that are observed by the researcher and that affect his/her utility.

²²By letting public school wages depend on a shock independent of teaching skills and on the type, I allow for them to be correlated with teaching skills, but I do not impose this correlation. Estimated type-specific parameters $\alpha_{0mM}(l_i)$ that differ significantly across types are evidence of non-zero correlation.

²³Under linear pricing, $\alpha_{0mNT}(l_i)$ absorbs the log-price of non-teaching skills and the constant of log-non-teaching skills. In estimation I do not impose any correlation between the random parameters $\alpha_{0mNT}(l_i)$ and $a_0(l_i)$ (the randomness comes from their dependence on type l_i). Therefore, I am not imposing any form of correlation between teaching and non-teaching skills.

²⁴For example, if skilled teachers are also skilled in non-teaching occupations, they obtain high wage offers from outside of teaching and schools must offer higher wages to attract them.

3.3 Equilibrium

Definition 1. *An equilibrium in market m is $(r_m^*, p_m^*, e^*(\cdot), d^*(\cdot))$ such that:*

- (a) r_m^* and p_m^* are a solution to the private school's problem;
- (b) given r_m^* and p_m^* , $e^*(\tilde{X}, k, \nu)$ is an optimal enrollment decision for every \tilde{X}, k, ν ;
- (c) given r_m^* , $d^*(X, \psi, \epsilon)$ is an optimal occupational choice for every X, ψ, ϵ .

An equilibrium is attained when all parents and all potential teachers choose the option that maximizes their utility, and the private school chooses prices that maximize its profits conditional on the tuition cap. Proposition 1 in Online Appendix C5.1 provides an existence argument. Further, it can be shown that the equilibrium of the second stage subgame (parental sorting across school sectors and potential teachers' labor supply) is unique (see Proposition 2 in Online Appendix C5.1). The uniqueness of the first stage, the profit maximization problem, is not proven analytically because the profit function does not admit a closed-form. However, Online Appendix C4 shows that it is well approximated by an interpolant function that admits a unique maximum.

3.3.1 Discussion

The choice of school sector that parents make depends on teacher labour supply directly and indirectly. First, parental school choice depends directly, among other things, on the expected skills of teachers in the two school sectors. This is endogenously determined within the model (see Online Appendix C5 for a derivation of equilibrium teaching skills by sector). Second, the wage elasticity of potential teachers determines the amount of resources needed by voucher schools to attract teachers of certain skills, therefore, it affects the tuition fees charged by private schools. This, in turn, affects parental choice of sector. This chain of effects is endogenously generated within the model. Simulations of counterfactual teacher policies capture not only how the teacher labor supply varies across sectors, but also how parental school choice responds. It is important to consider this margin of policy response, especially when teaching skills and student ability are not assumed to be additively separable in the production of achievement. In this general case, the policy impact on test scores depends on the teacher-student match induced by the policy.

Conversely, the choice of occupation in the labor supply part of the model does not directly depend on the characteristics of students in the two school sectors. Because of this assumption, the teacher supply in the profit function in section 3.1 can be simplified as a function of the wage rate r only, and not also of tuition fees p . This assumption greatly

simplifies the model solution and estimation. However, it still allows for occupational choices to be affected by student characteristics indirectly, through the price mechanism, because the wage rate r adjusts endogenously with school composition. The advantages of this assumption are that, first, it guarantees the existence and uniqueness of an equilibrium in the second stage of the model.²⁵ Without this restriction, solving the second stage of the model would require solving a fixed point problem. Not only existence of a solution would not be guaranteed in this case, but there would be the additional concern of multiplicity of the equilibria, which poses considerable challenges in estimation. Second, this assumption keeps the model tractable numerically. This assumption is less restrictive than it would initially appear, because the price mechanism captures the dependence of teacher labor supply on school composition. For example, a policy that increases the allocation of lower-SES students to private education would reduce the total net revenues of private schools through the financial aid mechanism, and this in turn would lower the price r that the private schools are able to offer to teachers. This would reduce the likelihood of teachers to choose private schools with lower SES students, a pattern uncovered in some previous studies (Hanushek, Kain, and Rivkin 1999).

4 Discussion of Model Features

The model makes a number of assumptions to obtain a reasonable and tractable approximation to how students and teachers choose school sectors.

The model abstracts from competition between private schools in a market and parental sorting across them. There are three reasons for this. First, understanding and tackling stratification across school sectors is a first-order objective of Chilean education policy (Epple, Romano, and Urquiola 2015). Second, this abstraction makes the equilibrium model tractable for two reasons. First, the private schools' pricing strategies can be obtained through the maximization of the profit function of a single monopolistic and monopsonistic private school per market.²⁶ This allows the model to abstract from the large set of potential strategies of competitors within each market, while still letting it

²⁵The second stage is equivalent to a sequential game where teachers move first. An equilibrium is attained when all parents and all potential teachers choose the option that maximizes their utility. An equilibrium exists because each potential teacher and each household have at least one most-preferred choice by construction: utilities are well defined. Moreover, because technology and preference shocks are continuously distributed in the population, no agent is indifferent between two options. Therefore, by backward induction the equilibrium is also unique.

²⁶ Neilson (2017) rejects perfect competition among private schools in Chile because of local market power.

capture the economic forces driving prices. The estimated prices implied by the model are found to correlate in a reasonable way with aggregate market characteristics not used in estimation (see Section 7.1). Second, this assumption keeps the choice sets of parents and of potential teachers small. Third, a school level analysis of teacher wage determination would be uninformative because only few teachers per school are observed.²⁷ While the estimation results reported in Section 7.1 are not affected by the assumptions made on the market structure, the counterfactual results are likely different from those that would be obtained by a model that allows for some degree of competition between different private schools in a market. In Section 7.2.1 I discuss how they would differ.

I use a Roy-model framework for the teacher supply part of the model because it allows me to predict the teaching skills of all individuals in the sample, including those who are not observed teaching in the data. The key parameters to identify teaching skills are those that enter equation (4) (the teaching skill production function) and that determine observed private school wages. Once these parameters are identified, teaching skills can be predicted for all potential teachers in the sample through equation (4). The estimation of a Roy model is made possible by the unique features of the Chilean labor market, where a large scale private school sector with individually negotiated wages co-exists alongside the more traditional public school sector with unionized wages.²⁸

Teaching skills are assumed to be uni-dimensional and private school wages are assumed to be linear in skills. These assumptions have two advantages: first, in line with existing work (Heckman, Lochner, and Taber 1998, Lee and Wolpin 2006), they allow me to separately identify skills from their prices. Second, they simplify the profit maximization problem of private schools, who, in the model, choose only a scalar skill price r rather than a vector of skill prices or a pricing function.

This paper focuses on teaching skills possessed by an individual, and it ignores teacher effort, following a stream of the teacher quality literature (Stinebrickner 2001a, Stinebrickner 2001b, Rothstein 2015, Hanushek, Rivkin, Rothstein, and Podgursky 2004, Biasi 2018, Rivkin, Hanushek, and Kain 2005). The reason is lack of data on teacher effort. Recent findings that separately identify the impact of teacher skills and effort suggest that skills have a considerably larger impact on student achievement than effort (Macartney,

²⁷ Neilson (2017) provides a model of competition among private schools in Chile, abstracting from competition across school sectors. Since he does not study teachers, he does not encounter the same sample size issues in his school-level analysis.

²⁸Exploiting similar data variation, Biasi (2018) estimates a Roy model of teacher labor supply using data from Wisconsin, where the passage of Act 10 in 2011 changed teacher pay determination from collective bargaining (like in public schools in Chile) to individual negotiations (like in voucher schools in Chile).

McMillan, and Petronijevic 2016, Biasi 2018). Therefore, it is an important margin to focus on as a first step. Complementing this study with endogenous teacher effort choices would be a valuable extension if adequate data become available.

The achievement production function does not impose a specific type of complementarity between student ability and teaching skills: teacher effects are allowed to be heterogeneous across students of different unobserved characteristics. This can be seen in Equation (2), where teacher skills \bar{s}_{mj} are multiplied by a type-dependent coefficient $\beta_{1j}(k_h)$, where k_h is the student type. This is a distinction with teacher value added models that measure teacher effects as an additive component of a test score residual (Hanushek 1971, Murnane 1975, Chetty, Friedman, and Rockoff 2014). An advantage of this approach is that it does not impose restrictions on the technology that would drive predictions of the impacts on test scores of policies that alter the student-teacher match. A limitation of this approach compared to teacher value added models is that the unbiased estimation of teacher effects relies on the correct specification of the Roy model of teacher labor supply, because teacher skills are recovered from data on merit-based teacher wages and teacher labor supply.

The achievement function is related to a commonly used specification that includes lagged test scores as one of the determinants of current test scores (Hanushek 1996, Hanushek 2003, Krueger 2000). The objective of such value added specifications is to account for the unobserved history of inputs as well as unobserved permanent student characteristics such as innate ability.²⁹ Because of data limitations, lagged test scores are not included as inputs in the achievement function in this paper. However, the achievement model includes student types, unobserved by the researcher, to capture any unmeasured determinants of test scores.

A limitation of this specification is that it does not include peer effects. This modelling choice guarantees numerical tractability and uniqueness of the equilibrium.³⁰ A second reason for this restriction is that the identification of peer effects poses additional demands on the data which are not clearly met by the dataset (Manski 1993). For similar reasons, existing models of the Chilean market for education equally abstract from peer effects (Urquiola and Verhoogen 2009, Neilson 2017). On the other hand, most studies of school choice that allow for peer effects abstract from teacher labor supply (see Altonji, Huang,

²⁹Todd and Wolpin (2003) show that this specification delivers consistent estimates of input effects only under certain assumptions on the true technology.

³⁰Had there been peer effects, the solution to the parents' problem would have required solving for a fixed point with no guarantee of equilibrium existence or uniqueness (Brock and Durlauf 2001). This is the same reason why school composition does not enter the direct sectoral preference of parents, as already explained in section 3.3.

and Taber (2015) and Epple and Romano (2008) for a structural approach and Dills (2005) for a non-structural one), and most studies of teacher labor supply abstract from parental school choice and peer effects.³¹ Because there currently does not exist a modelling framework that simultaneously studies parental school choice, teacher supply and private school pricing with peer effects, the model presented here could be seen as a reasonable first step. Future studies could build on it to include peer effects.

Finally, the model assumes that residential sorting is exogenous. While a literature on location choices and public goods exists (Epple and Sieg 1999, Nechyba 2000, Ferreyra 2007), there is not yet a well developed literature on two-sided equilibrium models with two-sided residential sorting. The paper that is closest to this one in terms of modelling matching patterns of schools and teachers is Boyd, Lankford, Loeb, and Wyckoff (2013). As in this paper, they estimate their model on multiple markets and treat the allocation of teachers and schools to markets as exogenous.³²

5 Identification and Estimation

5.1 Identification

5.1.1 Teaching Skills and Market-Specific Teaching Skill Prices

Teaching skills depend on parameters $a_0(l_i)$ and a_1 from private school wages (Equation (4)). Self-selection of potential teachers into occupations based on unobservables could bias estimates of wage parameters. To account for this, I frame teacher sorting within a Roy model (Roy 1951) of self-selection into occupations, a workhorse model in labor economics.³³ For a formal proof of identification of the wage and non-pecuniary preference parameters in this class of models see, for example, Heckman and Honore (1990). Intuitively, self-selection bias is accounted for through an adjustment term generated by

³¹Some studies examine the role that school choice programs play for teacher effectiveness, entry, turnover, and salaries (see, for example, Jackson (2012), Hensvik (2012) and Behrman, Tincani, Todd, and Wolpin (2016)), without explicitly analyzing parental school choice or peer effects.

³²In Biasi (2018), the district choice of teachers coincides with their labor supply choice. That is, teachers choose among districts that differ in terms of whether wages are unionized or individually negotiated. While the initial residential location of teachers is assumed to be exogenous, the labor supply choice generates a positive correlation between districts that use individual wage negotiation and teacher skills.

³³It has been used to study self-selection in a number of contexts, for example, immigration and residential choice (Borjas 1987), occupational and industrial choice (Heckman and Sedlacek 1985, Heckman and Sedlacek 1990), optimal taxation with self-selection (Rothschild and Scheuer 2012) and employment in the private and public sectors (Borjas 2002).

the structural model, which measures average unobserved skills conditional on an occupational choice, in the spirit of a control function approach (Heckman and Navarro 2004).³⁴ However, for the case of private school wages, unexplained wage variation is not only due to self-selection based on unobservables, but also to variation in market specific skill prices (which are unobserved). The self-selection correction method cannot separately identify log-prices ($\ln(r_m)$) from unobserved teacher skills ($a_0(l_i)$) in the constant of log-wages. For this, an additional identification strategy is needed.

Specifically, I exploit across market variation in private school wages, following closely the identification strategy in Heckman and Sedlacek (1985) (see note 17 in that paper). Markets in this paper play the role of years in Heckman and Sedlacek (1985). A key identifying assumption is that teaching skill prices r_m vary across markets, but not within markets. Imagine estimating a regression model of log-wages in private schools which includes market dummies. The market fixed effects, α_{mV} , are the sum of two terms: average unobserved skills of private school teachers in that market, $E[a_0(l_i) + \epsilon_i^{tech}|V, m]$, a market-specific self-selection term, and the unobserved log-price of teaching skills in that market, $\ln(r_m)$. The self-selection control function identifies unobserved skills in private schools up to scale, off of within-market data variation. Under the identifying assumption of exogenous residential location of potential teachers, normalizing the price of skills in one market sets the scale of the self-selection terms in all markets. The log-prices in the other markets are identified by subtracting the market-specific self-selection term from the market fixed effect: $\ln(r_m) = \alpha_{mV} - E[a_0(l_i) + \epsilon_i^{tech}|V, m]$.³⁵

5.1.2 Demand for Private Education and Unobserved Student Ability

Self-selection of parents into school sectors could bias the parameters of the achievement production technology and of the preferences for a school sector. First, the structural model explicitly accounts for this source of bias, and in so doing, it corrects for it parametrically in the spirit of a control functions approach. In the model, self-selection into school sectors is governed by the price elasticity of parents, which is determined by parental willingness to trade off child achievement for consumption ($\tau(k_h)$ in equation (1)). This is correlated with students' unobserved ability through the unobserved type k_h . Therefore,

³⁴Exclusion restrictions are not needed for parametric identification. However, in the model fertility variables affect the labour supply decision but not the wage conditional on working.

³⁵Teaching skills are identified only up to scale. This does not affect any of the counterfactual experiments presented, because the choice of normalizing constant does not affect the estimated impact of teaching skills on achievement. To see why, notice that the latter is given by the product $\beta_{1j}(k)\bar{s}_{jm}$ in equation (2), which is identified as a residual in the test score equation and, therefore, its identification is independent of the choice of scale for the average skills, \bar{s}_{jm} .

private and public school students may have different unobserved abilities.

While in estimation I rely on parametric model assumptions to obtain efficient estimates, in the Chilean setting there are naturally occurring exclusion restrictions which identify parental price elasticity for private education in the absence of parametric restrictions. The assignment of fellowships towards the payment of tuition fees is mandatory by law and depends on Government guidelines. This generates variation in tuition fee requests across parents that is not dictated by price discrimination on the part of schools and, therefore, it does not mirror student unobservables. Intuitively, fellowships generate observations where families with identical preferences and expected gains from private and public education make different school choices because they face different tuition fee requests. As long as at least one of the variables entering the fellowship formula is uncorrelated with unobserved student ability, (non-parametric) identification of price elasticities is possible. I find that two of the variables in the formula correlate significantly with the fellowship received by the student, but not with his/her test scores. Specifically, in an achievement regression that controls for income per capita, income per capita squared and parental education, the coefficient on the school level has a p-value of 0.396, and on rurality of 0.134. Conversely, these variables are highly significant in predicting fellowship amounts (the p-values are 0.000).³⁶

This is not surprising, because national guidelines limit the ability of schools in Chile to perfectly price-discriminate through the fellowship. For this reason, other studies use Chilean school fellowships for identification (see Anand, Mizala, and Repetto (2009)). Similar variation in tuition discounts due to family size has been used in the U.S. setting to identify price elasticities of parents and unobserved student ability (Dynarski, Li, and Gruber 2015, Altonji, Elder, and Taber 2005).

5.1.3 Productivity of teaching skills

An important question in this paper is how teaching skills and student characteristics combine to produce achievement in the two school sectors. To identify the parameters of the production of achievement, one must observe variation in teacher skills and in student characteristics and relate them to variation in student outcomes. However, not all teacher and student characteristics are directly observed. To overcome this, I model unobserved heterogeneity on both sides of the education market, and estimate both sides simultaneously on a linked dataset. The key insight is that I use private school wage data

³⁶Sargan and Basmann overidentification tests fail to reject the validity of these excluded variables as instruments for the private school in achievement regressions.

(and an appropriate correction for self-selection on unobservables) to build a measure of teaching skills that combines observed and unobserved teacher characteristics. This measure of teaching skills is then inserted into the achievement production function in estimation.

Specifically, teaching skills are identified for all individuals in the sample, including those not observed teaching, because the parameters of the teaching skill equation (4) are identified, as explained in section 5.1.1. The estimated parameters from this equation can be used to predict teaching skills for all individuals. At each parameter iteration in the estimation algorithm, an inner loop solves for the supply of teaching skills to each school, \bar{s}_{mj} , and inserts these simulated teaching skills into the achievement production function. Therefore, at each parameter iteration, it is *as if* teaching skills were observed. As a result, it is possible to estimate their impact on achievement and their effect heterogeneity across students ($\beta_{1j}(k_h)$). The latter would not be possible if teacher effects were estimated as a residual from test score data which is additively separable from student characteristics. Details of the algorithm can be found in appendix B.

5.1.4 Profit function parameters

The profit function parameters are identified from variation in skill prices r_m across markets. From the first order condition of the private school's problem, skill prices r_m are a function of the supply of teachers to private schools and demand from parents for private schools. These are functions of market characteristics that are identified in the first step of the estimation. To simplify the exposition of the identification argument, I present a linear version of the model. To further simplify, I do not include tuition fees p_m in the following equations, because this variable does not vary across markets. Consider the first order condition for skill prices:

$$r_m = \gamma_0 + \gamma_1 SV(r_m; \alpha_m) + \gamma_2 DV(r_m; \beta_m) + \epsilon_m, \quad (6)$$

where ϵ_m is the cost shock. The functions $SV(r_m; \alpha_m)$ and $DV(r_m; \beta_m)$ represent the supply of teachers to and demand from parents for voucher schools. To simplify further, assume they are linear: $SV_m = \alpha_{0m} + \alpha_1 r_m$ and $DV_m = \beta_{0m} + \beta_1 r_m$. The constants, indexed by m , represent idiosyncratic market conditions. For example, in the full structural model, markets differ in terms of the population distribution of parental characteristics and of potential teacher characteristics, moreover, wage rates in the non-teaching sector may vary across markets. The parameters of these functions are identified and estimated in the first

step of the estimation. Plugging these estimates into equation (6) and rearranging:

$$r_m = \frac{\gamma_0}{1 - \gamma_1 \hat{\alpha}_1 - \gamma_2 \hat{\beta}_1} + \frac{\gamma_1}{1 - \gamma_1 \hat{\alpha}_1 - \gamma_2 \hat{\beta}_1} \hat{\alpha}_{0m} + \frac{\gamma_2}{1 - \gamma_1 \hat{\alpha}_1 - \gamma_2 \hat{\beta}_1} \hat{\beta}_{0m} + \epsilon_m. \quad (7)$$

From this equation, it is clear that the identifying assumption is that the private schools' cost shock ϵ_m is independent of the idiosyncratic market conditions $\hat{\alpha}_{0m}$ and $\hat{\beta}_{0m}$. For example, as wage rates in the non-teaching sector vary, private school wage rates r_m are allowed to strategically respond. However, there should not be a correlation between wage rates in the non-teaching sector and the private school cost shock. When this is the case, variation in wage rates in the non-teaching sector generates variation in the supply of teachers to private schools that helps identify the private school's demand for teachers.³⁷

5.2 Estimation

Estimation is performed on multiple markets, each of which is assumed to be in equilibrium. There are two estimation steps. The first step estimates the parameters from the second stage of the model, θ_{II} . These are the preference and technology parameters of households and of individuals making labor supply decisions and the fellowship formula parameters. This step recovers also the market-specific prices of teaching skills, r_m , $m = 1, \dots, M$. The second estimation step treats the recovered skill prices as observations, and it uses them to estimate the parameters of the profit function in the private school sector, θ_I . The estimated θ_I parameters rationalize the prices p_m (tuition fees) and r_m (skill prices) as profit maximizing. Separation of the estimation in two steps is possible because the equilibrium of the second stage of the model depends on the profit function parameters only through their effect on tuition fees p_m and wage rate r_m .³⁸

5.2.1 Step one of the estimation

The parameter vector θ_{II} is estimated by the method of simulated moments (MSM) (McFadden 1989, Pakes and Pollard 1989). The method minimizes the distance between observed outcomes and outcomes simulated from the model. The outcomes are occupational choices of potential teachers, school choices of parents, wages of potential teachers,

³⁷Labor market opportunities outside of teaching have been found to be important determinants of teacher labor supply (Corcoran, Evans, and Schwab 2004a, Dolton and Klaauw 1999).

³⁸This method is similar in spirit to Moro (2003), who uses the first step to estimate an equilibrium object, and the second step to estimate the model parameters that rationalize the equilibrium.

test scores of children, and fellowship amounts. A list of the conditional moments used can be found in Online Appendix C6.

Because I use multiple data sources, I adjust the criterion function of the estimator and the parameter standard errors to account for the relative sizes of the samples and of their populations of reference. I follow the method developed in Bhattacharya (2005), details of which can be found in Online Appendix C7. This Online Appendix includes also technical details of the estimation, the asymptotic properties of the estimator, as well as details of the estimation of the asymptotic variance of the estimates.


5.2.2 Step two of the estimation

This step estimates the voucher sector cost parameters $\theta_I = [c_1, c_2, c_3, \sigma_{cost}]$. Estimation is by Nonparametric Simulated Maximum Likelihood (NPSML) (Laroque and Salanie 1989, Fermanian and Salanié 2004). All private schools are observed pricing at the tuition cap, hence, there is no variation across markets in this variable. Therefore, I use variation across markets in the wage rates r_m to estimate θ_I . The observation that the price cap is binding in all markets is treated as an over-identifying restriction.

The likelihood function does not have a closed-form expression, because the profit maximizing r_m in each market is an unknown function of the random cost shock, which does not admit a closed form. For this reason, the density of r_m , which enters the likelihood function, cannot be analytically derived. The NPSML method approximates the unknown likelihood function with a kernel-based nonparametric estimator based on simulations of the choice variable r_m . Under regularity conditions the estimator is consistent, asymptotically normal and asymptotically efficient when the number of simulations and observations go to infinity and the bandwidth goes to zero. For a formal description of how the NPSML method is implemented here, refer to Online Appendix C8.

Given a realisation of the cost shock, the model holds predictions not only for the skill price r_m in each market, but also for tuition fees p_m , which are not used in estimation. I use this predictions to test the model's first stage. Simulated tuition fees p_m at the estimated parameter values are binding at the tuition fee cap, as observed in the data. Therefore, the model is able to match a key feature of the data that is not matched by construction. In other words, the first step of the model is falsifiable but not falsified by the data, a desirable property for a model.

Figure 1: Model fit: parental sorting



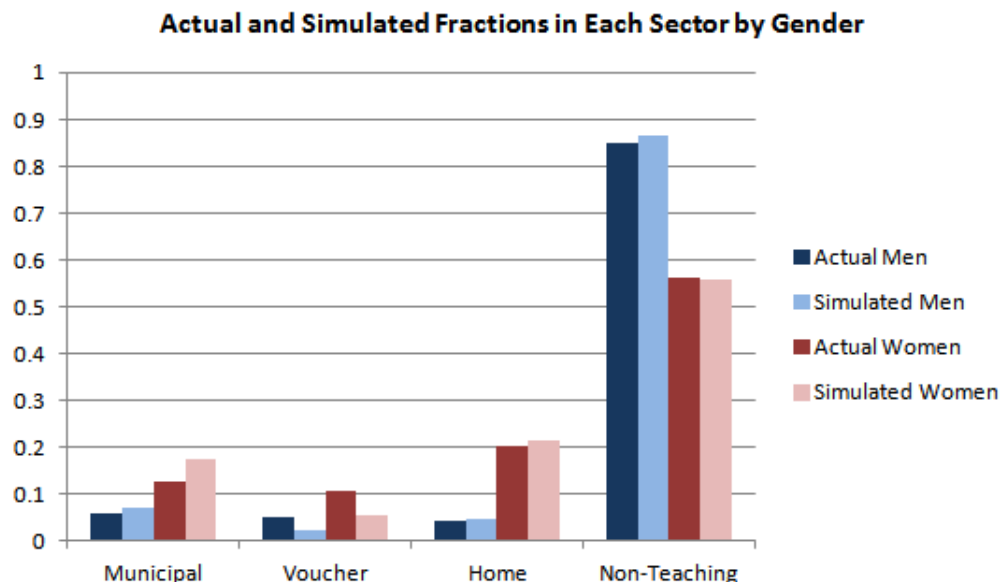
6 Model Fit

Table 3 shows the model fit. Simulations of the choice distributions of parents and potential teachers are very close to the data, within, respectively, 0.9 and 2.8 percentage points. Figures 1 and 2 show visually how accurate the model predictions are for the choice distributions of parents (by income) and potential teachers (by gender). The fit is similarly good when conditioning on other characteristics beyond gender and income. Figure C1 in Online Appendix C2 shows the accuracy of the model in predicting tuition payments (net of the voucher and fellowship) in private schools, which depend on the endogenous selection of households into private schools because different households are eligible for different fellowship amounts. Teaching wages simulated from the model are within 5 percent of actual wages, while nonteaching wages are over-predicted by around 17 percent. Mean test scores are under-predicted by 0.025 standard deviations. The simulated test score gap by school type is close to the actual one, and the gap by income is within 6.5 percent of the actual one. Figure C2 in the Online Appendix shows that the distributions of actual and simulated test scores by school type are close, especially for public schools. The fit is similar within markets. An example of a within market fit is presented in Online Appendix Figure C3.

Table 3: Model Fit

	Actual	Model
Parents		
Proportion Enrolled in Voucher Schools		
Overall	52.99%	52.38%
Primary	51.27%	50.99%
Secondary	55.09%	54.08%
Urban	55.06%	54.36%
Rural	29.46%	29.84%
Mean Tuition (1,000 CLP)	15.25	15.52
Mean Test Scores		
Overall	-0.003	-0.028
Municipal Schools	-0.191	-0.219
Voucher Schools	0.164	0.146
Gap Municipal-Voucher	0.355	0.365
Gap by Income (top-bottom quartile)	0.681	0.725
Potential Teachers		
Proportion Employed in		
Municipal Schools	9.48%	10.52%
Voucher Schools	7.81%	4.99%
Non-Teaching Occupations	70.28%	71.29%
Home	12.44%	13.21%
Mean Accepted Wages (1,000 CLP) in		
Teaching	4790	4553
Municipal Schools	5095	4918
Voucher Schools	4415	3786
Non-Teaching Occupations	7774	9077

Figure 2: Model fit: occupational choices



7 Empirical results

7.1 Estimation results

Parameter estimates of wage offers and potential teacher utility are reported in Table A1 and in Online Appendix Table C8. The reduced-form wage parameters suggested, surprisingly, that graduate degrees and certifications are valued more in private schools and the non-teaching sector than in public schools. By accounting for selection bias, structural estimates reveal that the opposite is true. Degrees and certifications are valued more in public schools, where wages are unionized and follow rigid formulae.³⁹ This gives me confidence in the goodness of the selection-correction provided by the structural model. A second finding is that women’s wages in Chile are not penalized more in non-teaching occupations than in teaching, as would be suggested by reduced-form estimates. In fact, the wage penalization for females in non-teaching jobs is similar to that in the public school sector, and smaller than in the voucher school sector (see the coefficients on the female dummy in Table A1). Therefore, the higher propensity of females to choose the teaching profession is not due to wage penalization (e.g. discrimination) outside of teaching, rather, it is due to non-pecuniary aspects of teaching like, for example, more flexible working

³⁹The log-wage structural coefficients on certifications and on graduate degrees in public school, voucher schools, and the non-teaching sector are, respectively: public schools 0.425, 0.403, voucher schools 0.36, 0.27, non-teaching sector -0.031 , 0.11.

arrangements. This is confirmed by the structural non-pecuniary preference parameters, indicating that females have a larger non-pecuniary preference for teaching, as well as for choosing to not work, especially when there are children in the household.⁴⁰ Finally, the non-pecuniary preference for staying at home is large and positive for women of unobserved type 3 and with children. For example, a woman of type 3 with three children has a positive non-pecuniary preference for staying at home that is two orders of magnitude larger than the coefficient of log-wage. This has implications for policy: the wage changes considered under the 2023 policy reform will not be sufficient to lure these women into teaching, changes to the non-pecuniary aspects of teaching jobs may be necessary too.

Focusing on the type-specific utility and teaching skill parameters, a pattern emerges. While non-teaching jobs give the highest direct utility and staying at home the lowest for all types, there is disagreement as to what teaching sector is most preferable. Specifically, type 1 individuals prefer municipal schools to voucher schools, while the opposite is true for types 2 and 3. In terms of teaching skills, type 1 are those with the highest skills and type 2 with the lowest. Therefore, the most highly skilled teachers (type 1) also have a preference for teaching in public rather than private schools.

The structural model separately identifies the price that voucher schools offer for teacher skills, from teacher skills themselves.⁴¹ About half (50.8%) of the unexplained variation in private school teacher wages across markets is explained by variation in skill prices, and the remaining half by variation in unobserved teaching skills in private schools.⁴² Therefore, variation in unobserved teaching skills and in the price that voucher schools offer for teaching skills are both important determinants of wage dispersion.

Estimated skill prices (reported in Online Appendix Table C9) respond to market forces in the expected direction. For example, they correlate negatively with the unemployment rate in the market (correlation coefficient: -0.31), which is a determinant of reservation wages. They also correlate positively (albeit with a smaller magnitude) with aggregate factors that affect the parental demand for private education, such as various aggregate SES indicators of the parents in the market. These correlations are not matched by the estimation nor are they true by construction. They (informally) validate the model's ability to capture key features of the economic forces operating in the teacher

⁴⁰Similarly, using U.S. teacher data and a structural model, Stinebrickner (2001b) finds that women who are married or who have children have a higher non-pecuniary preference for staying at home.

⁴¹See Heckman, Lochner, and Taber (1998) for a discussion of the advantages of separately identifying unobserved skills from unobserved skill prices in wages.

⁴²This is computed by dividing the variance in optimal skill prices across markets in the baseline simulation by the variance in market fixed effects estimated on simulated accepted log-wages in voucher schools.

labor market.

The correlation between non-teaching wage offers and the teaching skills of those employed in teaching is -0.10 for public school teachers and 0.38 for private school teachers. Private schools are better able to retain highly skilled teachers when appealing non-teaching options exist, a finding consistent with evidence from the United Kingdom and the United States.⁴³ More generally, teachers in public schools have lower skills than teachers in private schools, as can be seen in the top left panel of Figure 3. This is consistent with prior evidence in the literature that uses data not used in this paper (Bravo, Flores, Medrano, and Santiago 2010).

Parameter estimates of the achievement production function and of parental preferences are reported in Table A2 and in Online Appendix Table C10. First, the reduced form estimate of the constant in the production of test scores from Online Appendix Table C5 is lower in municipal (-1.184) than in voucher (-1.105) schools. However, these estimates confound the impact of all unobserved inputs, including teacher skills. The structural estimates net out teacher effects through the equilibrium of the model, and yield a lower constant in voucher than in municipal schools for all student types (-1.88 vs. -1.85 for type 1, -4.88 vs. -1.911 for type 2, -3.010 vs. -1.279). Therefore, other than lower teacher skills, which drives the reduced form result, municipal schools do not have worse unobserved inputs than private schools.

Second, there is a teacher-student complementarity: the coefficient on teaching skills varies among student types in a statistically significant way. Hence, to evaluate teacher policies that change the allocation of teachers and/or students across school sectors, it is not sufficient to use a teacher value added approach. Rather, one needs to simultaneously analyse teacher and student allocation across schools, and how the teacher-student match affects achievement.

Third, parents have, on average, a larger direct utility for private schools than for public school (the utility for private voucher schools exceeds that for municipal schools by 1.12 for type 1, 0.367 for type 2, and 1.196 for type 3). Simulations indicate that if parents did not have this direct preference for the private school sector, enrollment in private schools would be lower by 10 percentage points. However, similarly to the U.S. (Ferreyra and Kosenok 2018), public schools are more attractive at the primary school level. A possible reason for this preference is that public primary schools, with an average

⁴³Dolton and Klaauw (1999) show that “higher opportunity wages increase the tendency among teachers to switch careers and leave the profession voluntarily,” while Corcoran, Evans, and Schwab (2004b) find that the rise in employment opportunities for talented women in the United States is responsible for the decline in the quality of the teacher labor force.

size of 74 students, are smaller than voucher primary schools, which on average enroll 223 students. In both sector, students of type 1 perform better than students of type 3 who perform better than students of type 2.⁴⁴ As Online Appendix Table C10 shows, the parents of the lowest-performing type 2 students dislike municipal schools less than the parents of the other two student types, suggesting that in the data there exists a correlation between the parental direct preference for a sector and student’s ability.

Finally, there is heterogeneity in the parental preference for child achievement. Parents of type 3 place a higher weight on consumption utility than on their child’s achievement (5.648 vs. 1), while the opposite is true for parents of type 1 (0.118 vs. 1) and 2 (0.305 vs. 1).

7.2 Counterfactual experiments

7.2.1 Ex ante evaluation of the Chilean merit-based teaching reform

In 2017, the Chilean Government introduced a teaching reform which will be implemented gradually until 2023 (law *N. 20.903*). The reform has two key elements: new hires will be compensated based on their merit (measured through numerous competency assessments), and there will be minimum competency requirements for new teachers. These requirements are set to gradually increase between 2017 and 2023.

This counterfactual simulates the long-term equilibrium of the reform, after the introduction of the 2023 requirements.⁴⁵ It lets wage offers in public schools depend on teaching skills, and it introduces an entry requirement for teachers. It then simulates equilibrium outcomes, utilities and Government costs. Because an important part of the reform is to introduce granular teacher assessments, this experiment assumes that public schools can observe teacher skills.⁴⁶

⁴⁴This can be seen from a simple regression analysis of the simulated data that includes the simulated types as regressors.

⁴⁵In the transition to the new rules, current teachers in Chile are not subject to the minimum competency requirements. In the long-term equilibrium, minimum competency requirements apply to all individuals wanting to teach. It is likely that the effects reported here represent an upper-bound to the policy effects in the short-term, during the transition period to the long-term equilibrium.

⁴⁶A new programme called *Sistema Nacional de Desarrollo Profesional Docente* (National System of Teachers’ Professional Development) will run numerous assessments at different points in time throughout a teacher’s career. The teacher skill upon which wage depends is measured as a combination of teaching experience (5 categories) and the scores on two tests: the “portfolio” test (measuring how well teachers structure classes, how they master classroom dynamics with their students, how they give feedback to the students etc.) and a “knowledge test” (measuring the subject knowledge). These evaluations are composed of exams that the teachers must take, as well as classroom observations. These tests are first scored on a scale from 1.00 to 4.00, and then the “portfolio” test is converted into a 5 letter-grade (A, B, C, D, E) and the knowledge test into a 4 letter-grade (A, B, C, D), and each teacher’s final evaluation

First, in the counterfactual simulations the wage offer in public schools of individual i with teaching skills s_i is set to $r_M s_i$, where r_M solves

$$r_M = \frac{1.30 \frac{\sum_{i \in M_b} w_{iM}}{n_{M_b}}}{\frac{\sum_{i \in M_c} s_i}{n_{M_c}}}.$$

In other words, r_M is such that average teacher wages in public schools would increase by 30 percent under the reform, the stated Government goal (Sanchez 2016).⁴⁷ Second, simulations introduce the minimum competency requirement that will be implemented in 2023, requiring that teachers have a test score on the PSU, the national university entry exam, equal to the 70th percentile or higher.⁴⁸

Simulation results are reported in the second column of Table 4. The reform is predicted to increase test scores on average by 30 percent of a standard deviation (sd), and to decrease inequality in test scores by family income by around a third. Specifically, the difference in test scores between students in the top and in the bottom quarters of the income distribution decreases from 0.725 to 0.487 sd (fourth row of the top panel).⁴⁹

score is given by the combination of these two grades (so there are 20 possible outcomes). See <https://www.cpeip.cl/sistema-de-desarrollo> for a description of the SNDPD system, and <https://www.ayudamineduc.cl/ficha/tramos-y-progresion-en-carrera-docente> for an explanation of teacher career progression and assessments.

⁴⁷At the numerator, there is the average of public school wages at baseline (before the reform), while at the denominator there is the average of teaching skills in public schools under the counterfactual policy. M_b and M_c are the set of municipal sector teachers at baseline and under the counterfactual reform, respectively. n_{M_b} and n_{M_c} are their cardinalities.

⁴⁸I translate the cutoff in terms of the model’s teaching skills, using the known distributions of the university entry exam scores for the entire population of test-takers and for teachers pre-reform. Specifically, in the population of test-takers the 70th percentile corresponds to a score of 558.3. Using the statistics in Bravo, Flores, Medrano, and Santiago (2010), I calculate that among public school teachers before the reform, this cutoff corresponds to the 41st percentile of the distribution of entry exam scores. Assuming that teaching skills are monotonically increasing in the entry exam score, the 41st percentile in the university entry exam corresponds to the 41st percentile in the teaching skills distribution among public school teachers. From model simulations at baseline (pre-reform), I obtain the value of this cutoff in terms of teaching skills, and impose it as an entry requirement for teaching in all sectors in the policy simulations. These calculations are necessary because the percentile that defines the cutoff is defined with respect to the population of all test-takers, which is different from the population of potential teachers. Potential teachers have a college degree and have, on average, higher college entry exam scores than the population of all test-takers, which includes college dropouts.

⁴⁹These simulations assume that public schools can perfectly observe teacher skills. As a benchmark, I run a simulation that assumes that public schools can base merit pay only on observable teacher characteristics such as age and degrees. In this simulation, public school wages reward these characteristics with their marginal product (this is done by setting the wage offer coefficients equal to those in voucher schools). Test scores would not change substantially on average. The gap between the richest and poorest 25 percent of students would reduce from 0.725 sd at baseline to 0.681 sd, which is a 6.1 percent reduction, much smaller than the 32.8 percent reduction achieved when skills are observed. This indicates that the use of accurate teacher assessments to base merit pay have a potentially large benefit on test scores.

Table 4: Simulation of the 2023 reform and of alternative reforms, and analysis of the mechanisms behind the predicted impacts of the 2023 reform

Outcome	Baseline	2023 Reform	Flat bonus	Only competency	2023 Reform without reaction of:	
					Voucher schools	Parents
	(1)	(2)	(3)	(4)	(5)	(6)
Demand for education side						
Mean test scores (standardized)						
Entire population	-0.028	0.275	-0.004	0.104	0.025	0.183
Municipal schools	-0.219	0.272	-0.180	0.008	0.011	0.350
Voucher schools	0.146	0.280	0.172	0.216	0.235	0.019
Gap by income	0.725	0.487	0.704	0.613	0.144	0.480
Enrolment share in M	0.476	0.612	0.483	0.521	0.944	0.476
Supply of education side						
Mean teaching skills (standardized)						
Municipal schools	-0.058	1.209	0.011	0.504	1.092	1.209
Voucher schools	1.441	0.892	1.558	1.457	0.850	0.890
Non-teaching	-0.129	-0.226	-0.139	-0.166	-0.222	-0.226
Home	0.198	0.198	0.198	0.198	0.198	0.198
Accepted wages of teachers (baseline=1)						
Municipal schools	1.00	1.30	1.31	1.03	1.25	1.30
Voucher schools	1.00	1.07	1.11	1.04	1.15	1.08
Offered teaching skill prices (baseline=1)						
Voucher schools	1.00	1.27	1.09	1.02	1.00	1.27

Notes: Columns 2 to 4 simulate alternative reforms. Their impacts are obtained by comparing the simulated outcomes reported in the table with the corresponding ones from column 1, which reports simulated outcomes at baseline. The second column simulates the 2023 reform (i.e. merit-based pay and minimum competency requirement for teaching). The third column simulates a flat bonus to municipal school teachers. The fourth column simulated the same minimum competency requirement as under the 2023 reform, but it does not introduce merit-pay. The last two columns (5 and 6) are informative on the mechanisms behind the 2023 reform predicted impacts, because they simulate outcomes under the 2023 reform when private schools are not allowed to change the wage rates they offer (column 5) and when parents cannot change their school choice (column 6).

The quality of the pool of teachers improves: simulated teaching skills increase on average by 0.639 sd across both sectors.⁵⁰ Moreover, the teacher skills gap reverses, with better teachers in public schools than in private schools, and it reduces in size considerably, by 80 percent. The merit-based wages in public schools attract skilled individuals from outside of teaching, and cream skim good teachers from the voucher sector. As a result, average teaching skills increase in the public sector by 1.267 standard deviations, and they decrease in voucher schools and in the non-teaching sector by 0.549 sd and 0.097 sd respectively. However, average test scores do not decrease in voucher schools because of a pure compositional effect due to the outflow of lower SES students into public schools.⁵¹

Voucher schools face more competitive pressure from the public sector to attract highly skilled teachers. As a result, they increase their wage rate offers on average by 27 percent, but they cannot match the higher wage rate offered in public schools, leading to the aforementioned decrease in teaching skills. Therefore, average accepted wages in private schools increase only by 7 percent (because they are the product of the offered wage rate and the skills of those who accept the wage offer).

The last two columns explore channels of policy effectiveness. Column five shows that ignoring the response of private schools to the reform would result in an underestimation of the teaching skills of teachers in both school sectors, resulting in an underestimation of the positive effects of the reform on average test scores by 0.250 sd. This means that equilibrium wage adjustments in the large non-public school sector result in a better self-selection of individuals into teaching and across school sectors which greatly amplifies the reform's positive effects on test scores. Column six shows that the reaction of parents, too, amplifies the impacts on average test scores, albeit by a smaller amount (0.092 sd). Moreover, it effectively eliminates the test score gaps across school sectors that would exist if parents were not allowed to move. The changes in teacher composition induced by

⁵⁰There is a distinction between teaching skills, an ordinal measure of skills, and the impact of those skills on student test scores. Because teaching skills in the model are allowed to have heterogeneous impacts on test scores across students and schools, model estimation does not yield a simple value-added measure of teaching skills that does not depend on students or schools. For this reason, I present impacts both on the ordinal measure of teaching skills, to analyze teacher quality across schools, and on test scores, which are the outcome of interest. The Table's sixth column can be used to benchmark the ordinal teaching skills in terms of student test scores: a decrease of about half a standard deviation in teacher skills in private schools results in a decrease by about 0.10 sd in test scores, keeping student composition fixed, while an increase in teacher skills in public schools of 1.3 sd results in an increase in test scores of about 0.60 sd, keeping student composition fixed.

⁵¹The reform does not induce any movement of students out of the public sector. New enrollees in public schools come, on average, from the lower tail of the SES-distribution of baseline voucher sector students. As a result, the reform improves the average SES in both school sectors, without changing considerably the gap in SES, as seen in Table C6 in Online Appendix C2.

the reform would lead to a dramatic reversal of the sign of the test score gap across school sectors if parents did not respond to the policy. However, parameter estimates indicate that parents are responsive to the distribution of teaching skills across sectors, and their school choice adjusts in a way that reduces the test score gap across sectors from 0.331 sd (column 6) to 0.008 st (column 2). These counterfactual experiments demonstrate that an equilibrium model is required to understand the full impacts of a large-scale teacher reform. A key lesson that we learn is that the existence of a large school choice program amplifies the positive impacts of merit reforms in the teacher labor market, thanks primarily to equilibrium wage adjustments in the large non-public school sector that lead to a more skilled pool of teachers, and to a lesser extent to parental sorting.

What are the individual contributions of demand and supply factors? Allowing for only the parental response but not the voucher schools' response would explain about 10 percent of the treatment effect on mean test scores (row one, column four), while allowing for only the voucher schools' reaction but not the parental reaction would explain almost 70 percent of the treatment effect (row one, column five). Therefore, the policy reactions on the supply side of the market contribute relatively more to overall impacts than demand side reactions. A lesson that we learn is that in a context with a large subsidised private school sector, private school reactions are an important driver of the effects of public policies in the teacher labor market. This result is in line with Neilson (2017), who finds that also in the case of reforms of the subsidy structure, policy impacts are driven by the supply side response of private schools rather than by parental sorting when there exists a large private subsidised school sector.⁵²

Finally, Table 5 contains a cost-benefit welfare analysis of the reform. There are positive impacts on the average utility of households, with poorer families enjoying higher gains. There are fairly large positive impacts on those college graduates who before the introduction of the reform would have chosen to work in the voucher sector: these are skilled individuals who enjoy teaching. On average, those who would have chosen to work in public schools at baseline are made worse off by the reform. This result is driven by the low skilled public school teachers who, before the reform, enjoyed high wages relative to their skills. Importantly, Government costs are substantially unchanged. At baseline, the Government needs to inject additional funds into public schools to cover their running

⁵²If the market structure was mis-specified and in the true data generating process there was a higher degree of competition between private schools in the same market, it is reasonable to expect that the teaching skill prices in private schools would increase by more in response to the merit-pay reform in public schools than under current model predictions, and that tuition fees may decrease below the price cap. As a result, the predicted policy impacts reported here are likely to provide a lower bound on the benefits of the reform, both in terms of average achievement and the achievement gap.

costs: 35 percent of overall expenditures in these schools is covered by these additional funds. Under the reform, this fraction reduces to 30 percent, so that spending per-pupil approaches the value of the per-pupil voucher.

The results of this counterfactual policy analysis indicate that the merit-based reform introduced in Chile in 2017 and whose implementation is planned to complete in 2023 will increase student test scores by around 0.30 sd, and it will reduce the achievement gap between the poorest and the richest 25 percent of students by a third.⁵³ This result is driven by a stark improvement in the quality of the pool of teachers due to two factors: public school wages that better reflect skills in public schools, and an increase in wage rates in the large private-subsidised school sector in response to the reform. Additionally, the policy will benefit the welfare of households (as measured by their simulated utility functions) by 0.08 sd, with higher positive impacts on the poorer families. Finally, the policy will pay for itself in the long run because it is financially more sustainable than the status quo pre-reform. Specifically, the improved teaching skills of teachers in the public sector attracts voucher revenues back into the public sector. This reduces the need for additional, non-voucher funds to cover costs in public schools.

7.2.2 Two alternative policies: flat wage increase in Municipal schools and competency requirements for teaching

Under the first counterfactual experiment, public schools offer a bonus that is independent of skills in addition to the wages currently being offered. I simulate the equilibrium labor supply of potential teachers, private school wage offers and parental school choices under this alternative compensation rule in public schools. To make this experiment comparable to the merit-reform, the bonus is such that, in equilibrium, it results in a 30 percent increase in average accepted wages in public schools. This is the same increase that the Government aims to achieve under the merit reform.

This policy is less effective at attracting skilled teachers into public schools than a merit-based reform. Figure 3 shows the simulated distributions of teaching skills by

⁵³A *caveat* to this result is that the model does not include the potential effects of class-size. Under the reform, the average size of classes in public schools is projected to increase by 10 from 12, and in private subsidised schools to decrease by 10 from 27. Therefore, the results may be underestimating (overestimating) the positive impacts in voucher (municipal) schools. Estimates of class size impacts among the highest available in the literature (Krueger 2003) indicate that a 10 pupil increase (decrease) would decrease (increase) test scores by roughly 0.10 sd, suggesting that the main conclusion that the reform is predicted to improve mean test scores overall is likely to hold even in the presence of class-size effects. A more accurate quantification of the role of class-size in policy effects requires a model where class-size is endogenously determined.

Table 5: Welfare analysis of the 2023 merit-based reform


Outcome	Baseline	Reform	Effect
Mean parental utility (standardized at baseline)			
Entire population	0.000	0.079	0.079
Above median income	0.471	0.525	0.054
Below median income	-0.452	-0.349	0.103
Mean utility of college graduates in the labour market (standardized at baseline)			
Entire population of college graduates	0.000	-0.002	-0.002
Public school teachers at baseline	-0.740	-1.131	-0.391
Voucher school teachers at baseline	-0.757	-0.209	0.548
Non-teaching sector at baseline	0.162	0.179	0.017
Private school profits			
Baseline=1	1.00	0.79	-0.21
Cost to the Government			
Baseline=1	1.00	1.00	0.00
Breakdown of costs to the Government (billion CLP, per month)			
<i>Expenditures in Municipal schools</i>	39.13	46.08	6.95
Fraction covered by additional, non-voucher funds	0.35	0.30	-0.05
Per-pupil expenditure (relative to per-pupil voucher)	1.53	1.44	-0.09
<i>Expenditures in Private Voucher schools</i>	27.07	19.96	-7.11
Per-pupil expenditure (relative to per-pupil voucher)	1.04	1.04	0.00
Total expenditures	66.20	66.04	-0.16

Utilities are standardized using the baseline mean and standard deviation, therefore, the last column contains policy effects in terms of standard deviations. Because individuals who do not work at baseline do not work under the reform and viceversa, I analyse the welfare impacts on those who at baseline are in the labour market because these are the only college graduates who are affected by the reform.

Impacts on private school profits and Government costs are expressed in percentage terms (e.g., private school profits increase by 23 percent under the reform). Expenditures in municipal schools are voucher plus non-voucher expenditures. The non-voucher expenditures are calculated as the difference between the total wage bill and total voucher expenditure. Expenditures in private voucher schools are voucher expenditures plus fellowship sponsorship.

occupational sectors at baseline, under the merit reform, and under the flat bonus experiment. The merit-based reform causes the mass of highly skilled teachers to increase in public schools, but under the flat bonus the teaching skill distribution in public schools does not experience such a considerable shift to the right. The distribution of skills in the voucher sector remains close to baseline under the flat bonus, indicating that there is no cream skinning of the best voucher school teachers into public schools, which instead occurs under the merit-based reform. Unlike the merit-based reform, highly skilled individuals are not attracted from the non-teaching sector either (the distribution of skills among non-teachers does not shift to the left like it does under the merit-based reform).

Figure 3: Teaching skills by occupational sector: baseline and counterfactual



Further simulation results are reported in column three of Table 4. Average teaching skills of Municipal school teachers increase by 0.069 sd, ten times less than under the merit reform. In private schools, teaching skills improve by 0.117 sd, because of the absence of cream skinning from public schools and because of a 9 percent increase in teaching skill prices in voucher schools. Average teaching skills improve (marginally) in both sectors because the worst voucher school teachers move to public schools, where they are, on

average, more highly skilled than the incumbents.

Test scores increase on average in both sectors, but by a much smaller amount than under the merit reform. Specifically, they increase by 0.039 sd in municipal schools and by 0.026 sd in voucher schools. Overall, test scores increase by 0.024 sd, only 8 percent of the treatment effect of the merit-based reform. The bonus does not improve inequality considerably, reducing the gap by income by only 3 percent, compared to a 33 percent reduction under the merit reform. Impacts on test score levels and inequality are about a tenth of those produced by the merit-based reform. A flat increase to public school wages would be considerably less effective at improving student achievement levels and inequality than the planned 2023 merit-based reform.

Column four of Table 4 reports the results of an experiment that introduces the same competency requirements for teaching as under the 2023 reform, but that does not change teacher wages. Under this experiment, municipal schools would employ individuals with higher teaching skills than at baseline, as can be seen in Figure 3. However, the teaching skills of those employed in municipal schools would be, on average, 0.70 sd below what they would be under the 2023 reform. On the other hand, voucher schools would have a more skilled teaching body than under the 2023 reform, because in the absence of merit-pay in public schools, it is less costly for private voucher schools to attract the best teachers away from municipal schools. This is confirmed also by the muted reaction of equilibrium teaching skill prices in Voucher schools. This latter aspect of the competency-only experiment, coupled with the stratification of students by socio-economic status across school sectors, helps explain why this policy is not as effective as the 2023 reform at reducing achievement gaps between the rich and the poor. To summarize, this counterfactual indicates that the competency requirement alone would achieve just below half of the predicted positive impacts of the 2023 reform: specifically, 43 percent of its positive impact on mean test scores and 47 percent of its negative impact on the test score gap between rich and poor students. Therefore, both components of the 2023 reform are empirically important in driving its impacts

8 Conclusions

Discussions of school choice typically focus on competition in the market for the education output. I show that competition in the market for teachers (an input) is empirically important too. In the presence of a large-scale school choice program, public policies generate equilibrium effects not only on the sorting of parents, but also on private school

wages and on the sorting of teachers across school sectors. A policy maker needs to take these reactions into account to correctly predict policy impacts.

This is one of the first papers to model both sides of the education market and of the market for teachers. The advantage of this approach is that it allows me to quantify the importance of different channels in driving teacher policy impacts. Simulations from the estimated model are used to perform an ex-ante evaluation of a planned teacher reform in Chile. They show that when merit is rewarded more in public schools, private schools have an increased incentive to reward merit. This induces an equilibrium adjustment in private school wages that improves the selection of teachers with respect to the selection that would have occurred in the absence of the adjustment. Empirically, these equilibrium effects in the market for teachers are important, accounting for 70 percent of estimated policy impacts. A lesson that we learn is that, in a market-based school choice system, competition in the market for teachers amplifies the positive impacts on student test scores of merit-based reforms.

A Estimates of selected structural parameters

Table A1: Parameters of Log-Wage Offer Functions

	Public	Voucher	Non-Teaching
Intercept, type 1	0.0245 (0.0330)	0.0642 (0.0002)	1.3800 (0.0000)
Intercept type 2 minus type 1	0.0479 (0.0002)	-1.0400 (0.0002)	-1.1800 (0.0000)
Intercept type 3 minus type 1	-0.0127 (0.0002)	-0.0193 (0.0002)	-2.0400 (0.0000)
Age	0.0399 (0.0001)	0.0863 (0.0001)	0.0101 (0.0002)
Age Squared	-0.0002 (0.0002)	-0.0017 (0.0002)	-0.0004 (0.0002)
Female Dummy	-0.1430 (0.0002)	-0.1710 (0.0002)	-0.138 (0.0002)
Has Professional Certificates	0.4250 (0.0002)	0.3610 (0.0002)	-0.0313 (0.0002)
Graduate Degree	0.4030 (0.0002)	0.271 (0.0002)	0.1190 (0.0002)
$\log(\sigma)$	-1.2200 (0.0002)	-0.8090 (0.0002)	-0.4000 (0.0002)

Standard errors in parenthesis. The type 1 intercepts in Public and Non-Teaching wages and the type 2 and 3 in Non-Teaching wages admit some regional variation: their mean is reported, the standard errors are calculated under the assumption of independence.

B Estimation Algorithm Embedding Equilibrium Restriction

This algorithm refers to the first step of the estimation. For simplicity, I drop the subscript II from θ_{II} .

- Choose an initial guess for the parameter: $\theta^{(0)}$.
- Draw unobserved types and shocks for each potential teacher and student:

$$k_h, l_i, \nu_{hm}^{pref}, \nu_{hM}, \nu_{hV}, \epsilon_{iH}^{pref}, \epsilon_i^{tech}, \epsilon_{iM}, \epsilon_{iNT}.$$

Table A2: Production of Achievement

	Public	Voucher
Intercept, type 1	-1.8500 (0.0343)	-1.8800 (0.0862)
Intercept, type 2 minus type 1	-0.0610 (0.1414)	-3.0000 (0.1399)
Intercept, type 3 minus type 1	0.5710 (0.1344)	-1.1300 (0.1250)
Teachers' skills, type 1	0.3400 (0.1551)	0.2110 (0.0915)
Type 2 minus type 1	0.0374 (0.1522)	-0.1950 (0.1115)
Type 3 minus type 1	-0.2100 (0.1373)	-0.2110 (0.0908)
Parental education, type 1	0.0572 (0.1271)	0.1030 (0.1112)
Type 2 minus type 1	-0.0471 (0.1242)	0.1380 (0.0942)
Type 3 minus type 1	0.0427 (0.0601)	0.0119 (0.1278)
Income (monthly, per capita), type 1	0.1550 (0.1335)	0.9780 (0.1097)
Type 2 minus type 1	0.0063 (0.1440)	0.2330 (0.1231)
Type 3 minus type 1	-0.1480 (0.1453)	0.2810 (0.0915)
Income squared (monthly, per capita)	-0.0467 (0.1506)	-0.2520 (0.1240)
Log of shock standard deviation	-0.0322 (0.8320)	-0.1950 (0.1322)

The intercept parameters for type 1 admit geographical variation: their means are reported, the standard errors are calculated under independence.

- Use $a^{(0)} = [a_0^{(0)}(l_i) \quad a_1^{(0)}]'$ and the technology shock ϵ_i^{tech} to compute teaching skills for each potential teacher i : $s_i(a^{(0)}, \epsilon_i^{tech})$.
- Calculate the optimal occupational choice of each potential teacher in each market and use these individual choices to calculate the mean skills supplied to each school sector $j \in \{M, V\}$ in each market m : \bar{s}_{jm} . This is the non-linear function of $a^{(0)}$ in equation (4). Simulate also accepted wages.
- Plug the values for mean teacher skills into the production functions for achievement in M and V.
- Simulate achievement of each student in each school, and simulate optimal parental choice of school.
- Average over all simulation draws to compute the value of the objective function of the Method of Simulated Moments using simulated and real data.
- Update guess $\theta^{(0)}$ to $\theta^{(1)}$ (using the Generating Set Search optimization algorithm in HOPSPACK) and repeat until objective function is minimized.

References

- ALTONJI, J. G., T. E. ELDER, AND C. R. TABER (2005): “An evaluation of instrumental variable strategies for estimating the effects of catholic schooling,” *Journal of Human resources*, 40(4), 791–821.
- ALTONJI, J. G., C.-I. HUANG, AND C. R. TABER (2015): “Estimating the cream skimming effect of school choice,” *Journal of Political Economy*, 123(2), 266–324.
- ANAND, P., A. MIZALA, AND A. REPETTO (2009): “Using school scholarships to estimate the effect of private education on the academic achievement of low-income students in Chile,” *Economics of Education Review*, 28(3), 370–381.
- BEHRMAN, J. R., M. M. TINCANI, P. E. TODD, AND K. I. 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.
- BEN-PORATH, Y. (1967): “The production of human capital and the life cycle of earnings,” *Journal of political economy*, 75(4, Part 1), 352–365.
- BHATTACHARYA, D. (2005): “Asymptotic inference from multi-stage samples,” *Journal of econometrics*, 126(1), 145–171.
- BIASI, B. (2018): “The Labor Market for Teachers Under Different Pay Schemes,” Discussion paper, National Bureau of Economic Research.
- BONHOMME, S., G. JOLIVET, AND E. LEUVEN (2016): “Job Characteristics and Labor Turnover: Assessing the Role of Preferences and Opportunities in Teacher Mobility,” *Economic Journal*, 126(2), 540–563.
- BORJAS, G. J. (1987): “Self-Selection and the Earnings of Immigrants,” *The American Economic Review*, pp. 531–553.
- (2002): “The wage structure and the sorting of workers into the public sector,” Discussion paper, National Bureau of Economic Research.
- BOYD, D., H. LANKFORD, S. LOEB, AND J. WYCKOFF (2005): “Explaining the short careers of high-achieving teachers in schools with low-performing students,” *American economic review*, pp. 166–171.
- (2013): “Analyzing the Determinants of the Matching of Public School Teachers to Jobs: Disentangling the Preferences of Teachers and Employers,” *Journal of Labor Economics*, 31(1), 83–117.
- BRAVO, D., B. FLORES, P. MEDRANO, AND D. D. SANTIAGO (2010): “¿ Se premia la habilidad en el mercado laboral docente? ¿ Cuánto impacta en el desempeño de los estudiantes?,” *Serie Documentos de Trabajo*.

- BROCK, W. A., AND S. N. DURLAUF (2001): “Discrete choice with social interactions,” *The Review of Economic Studies*, 68(2), 235–260.
- CHETTY, R., J. N. FRIEDMAN, AND J. E. ROCKOFF (2014): “Measuring the impacts of teachers I: Evaluating bias in teacher value-added estimates,” *American Economic Review*, 104(9), 2593–2632.
- CORCORAN, S. P., W. N. EVANS, AND R. M. SCHWAB (2004a): “Changing labor-market opportunities for women and the quality of teachers, 1957-2000,” *American Economic Review*, 94(2), 230–235.
- (2004b): “Women, the labor market, and the declining relative quality of teachers,” *Journal of Policy Analysis and Management*, 23(3), 449–470.
- DEE, T. S., AND J. WYCKOFF (2015): “Incentives, selection, and teacher performance: Evidence from IMPACT,” *Journal of Policy Analysis and Management*, 34(2), 267–297.
- DILLS, A. K. (2005): “Does cream-skimming curdle the milk? A study of peer effects,” *Economics of Education Review*, 24(1), 19–28.
- DOLTON, P., AND W. V. D. KLAAUW (1999): “The turnover of teachers: A competing risks explanation,” *Review of Economics and Statistics*, 81(3), 543–550.
- DYNARSKI, S., D. LI, AND J. GRUBER (2015): “Cheaper by the Dozen: Using Sibling Discounts at Catholic Schools to Estimate the Price Elasticity of Private School Attendance,” *Harvard Business School Working Paper*.
- ELACQUA, G. (2006): “Public, Catholic, and for-profit school enrollment practices in response to vouchers: Evidence from Chile,” *National Center for the Study of Privatization in Education, Teachers College, Columbia University, paper*, 125, 577–601.
- EPPLE, D., AND R. ROMANO (2008): “Educational vouchers and cream skimming,” *International Economic Review*, 49(4), 1395–1435.
- EPPLE, D., R. E. ROMANO, AND M. URQUIOLA (2015): “School vouchers: A survey of the economics literature,” *forthcoming, Journal of Economic Literature*.
- EPPLE, D., AND H. SIEG (1999): “Estimating equilibrium models of local jurisdictions,” *Journal of political economy*, 107(4), 645–681.
- FERMANIAN, J.-D., AND B. SALANIÉ (2004): “A nonparametric simulated maximum likelihood estimation method,” *Econometric Theory*, 20(04), 701–734.
- FERREYRA, M. M. (2007): “Estimating the effects of private school vouchers in multi-district economies,” *The American Economic Review*, pp. 789–817.
- FERREYRA, M. M., AND G. KOSENOK (2018): “Charter school entry and school choice: The case of Washington, DC,” *Journal of Public Economics*, 159, 160–182.

- HANUSHEK, E. (1971): “Teacher characteristics and gains in student achievement: Estimation using micro data,” *The American Economic Review*, 61(2), 280–288.
- HANUSHEK, E. A. (1996): “School resources and student performance,” *Does money matter? The effect of school resources on student achievement and adult success*, pp. 43–73.
- (2003): “The failure of input-based schooling policies,” *The economic journal*, 113(485), F64–F98.
- HANUSHEK, E. A., J. F. KAIN, AND S. G. RIVKIN (1999): “Do higher salaries buy better teachers?,” Discussion paper, National bureau of economic research.
- HANUSHEK, E. A., S. G. RIVKIN, R. ROTHSTEIN, AND M. PODGURSKY (2004): “How to improve the supply of high-quality teachers,” *Brookings papers on education policy*, (7), 7–44.
- HASTINGS, J. S., AND J. M. WEINSTEIN (2008): “Information, school choice, and academic achievement: Evidence from two experiments,” *The Quarterly journal of economics*, 123(4), 1373–1414.
- HECKMAN, J., AND S. NAVARRO (2004): “Using matching, instrumental variables, and control functions to estimate economic choice models,” *Review of Economics and statistics*, 86(1), 30–57.
- HECKMAN, J., AND B. SINGER (1984): “A method for minimizing the impact of distributional assumptions in econometric models for duration data,” *Econometrica: Journal of the Econometric Society*, pp. 271–320.
- HECKMAN, J. J., AND B. E. HONORE (1990): “The empirical content of the Roy model,” *Econometrica: Journal of the Econometric Society*, pp. 1121–1149.
- HECKMAN, J. J., L. LOCHNER, AND C. TABER (1998): “Explaining rising wage inequality: Explorations with a dynamic general equilibrium model of labor earnings with heterogeneous agents,” *Review of economic dynamics*, 1(1), 1–58.
- HECKMAN, J. J., AND G. SEDLACEK (1985): “Heterogeneity, aggregation, and market wage functions: an empirical model of self-selection in the labor market,” *The journal of political economy*, 93(6), 1077–1125.
- HECKMAN, J. J., AND G. L. SEDLACEK (1990): “Self-selection and the distribution of hourly wages,” *Journal of Labor Economics*, 8(1, Part 2), S329–S363.
- HENSVIK, L. (2012): “Competition, wages and teacher sorting: Lessons learned from a voucher reform,” *The Economic Journal*, 122(561), 799–824.
- 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(8), 1477–1503.

- IMBERMAN, S. A., AND M. F. LOVENHEIM (2015): “Incentive strength and teacher productivity: Evidence from a group-based teacher incentive pay system,” *Review of Economics and Statistics*, 97(2), 364–386.
- JACKSON, C. K. (2012): “School competition and teacher labor markets: Evidence from charter school entry in North Carolina,” *Journal of Public Economics*, 96(5), 431–448.
- KEANE, M. P., AND K. I. WOLPIN (1997): “The career decisions of young men,” *Journal of political Economy*, 105(3), 473–522.
- KRUEGER, A. B. (2000): “An economist’s view of class size research,” *The class size debate, Washington, DC, Economic Policy Institute*, 130.
- (2003): “Economic considerations and class size,” *The Economic Journal*, 113(485).
- LAROQUE, G., AND B. SALANIE (1989): “Estimation of multi-market fix-price models: An application of pseudo maximum likelihood methods,” *Econometrica: Journal of the Econometric Society*, pp. 831–860.
- LEE, D., AND K. I. WOLPIN (2006): “Intersectoral labor mobility and the growth of the service sector,” *Econometrica*, 74(1), 1–46.
- LOEB, S., AND M. E. PAGE (2000): “Examining the link between teacher wages and student outcomes: The importance of alternative labor market opportunities and non-pecuniary variation,” *Review of Economics and Statistics*, 82(3), 393–408.
- MACARTNEY, H., R. MCMILLAN, AND U. PETRONIJEVIC (2016): “A Unifying Framework for Education Policy Analysis,” .
- MANSKI, C. F. (1993): “Identification of endogenous social effects: The reflection problem,” *The review of economic studies*, 60(3), 531–542.
- MC EWAN, P. J., M. URQUIOLA, E. VEGAS, R. FERNANDES, AND F. A. GALLEGU (2008): “School Choice, Stratification, and Information on School Performance: Lessons from Chile [with Comments],” *Economia*, pp. 1–42.
- MCFADDEN, D. (1989): “A method of simulated moments for estimation of discrete response models without numerical integration,” *Econometrica: Journal of the Econometric Society*, pp. 995–1026.
- MEHTA, N. (2018): “The potential output gains from using optimal teacher incentives: An illustrative calibration of a hidden action model,” *Economics of Education Review*, 66, 67–72.
- MORO, A. (2003): “The Effect Of Statistical Discrimination On Black–White Wage Inequality: Estimating A Model With Multiple Equilibria*,” *International Economic Review*, 44(2), 467–500.

- MURNANE, R. J. (1975): “The impact of school resources on the learning of inner city children.” .
- NECHYBA, T. J. (2000): “Mobility, targeting, and private-school vouchers,” *American Economic Review*, pp. 130–146.
- NEILSON, C. (2017): “Targeted vouchers, competition among schools, and the academic achievement of poor students,” *Mimeo, Princeton University*.
- PAKES, A., AND D. POLLARD (1989): “Simulation and the asymptotics of optimization estimators,” *Econometrica: Journal of the Econometric Society*, pp. 1027–1057.
- PODGURSKY, M. (2008): “Teams versus Bureaucracies: Personnel Policy, Wage-Setting, and Teacher Quality in Traditional Public, Charter, and Private Schools.,” *In Mark Berends, Matthew Springer, and Herbert Walberg (eds)*.
- RIVKIN, S. G., E. A. HANUSHEK, AND J. F. KAIN (2005): “Teachers, schools, and academic achievement,” *Econometrica*, 73(2), 417–458.
- ROTHSCHILD, C., AND F. SCHEUER (2012): “Redistributive taxation in the Roy model,” *The Quarterly Journal of Economics*, p. qjs076.
- ROTHSTEIN, J. (2015): “Teacher quality policy when supply matters,” *American Economic Review*, 105(1), 100–130.
- ROTHSTEIN, J. M. (2006): “Good principals or good peers? Parental valuation of school characteristics, Tiebout equilibrium, and the incentive effects of competition among jurisdictions,” *The American economic review*, 96(4), 1333–1350.
- ROY, A. D. (1951): “Some thoughts on the distribution of earnings,” *Oxford economic papers*, pp. 135–146.
- SANCHEZ, J. V. (2016): “Sistema de Desarrollo Profesional Docente,” Discussion paper, CPEIP.
- SPRINGER, M. G., D. BALLOU, L. HAMILTON, V.-N. LE, J. LOCKWOOD, D. F. MCCAFFREY, M. PEPPER, AND B. M. STECHER (2011): “Teacher Pay for Performance: Experimental Evidence from the Project on Incentives in Teaching (POINT).,” *Society for Research on Educational Effectiveness*.
- STINEBRICKNER, T. R. (2001a): “Compensation policies and teacher decisions,” *International Economic Review*, 42(3), 751–780.
- (2001b): “A dynamic model of teacher labor supply,” *Journal of Labor Economics*, 19(1), 196–230.
- SUNDARARAMAN, V., AND K. MURALIDHARAN (2011): “Teacher Performance Pay: Experimental Evidence from India,” *Journal of Political Economy*, 119(1).

- TODD, P. E., AND K. I. WOLPIN (2003): “On the specification and estimation of the production function for cognitive achievement,” *The Economic Journal*, 113(485), F3–F33.
- URQUIOLA, M. (2005): “Does school choice lead to sorting? Evidence from Tiebout variation,” *The American Economic Review*, 95(4), 1310–1326.
- URQUIOLA, M., AND E. VERHOOGEN (2009): “Class-size caps, sorting, and the regression-discontinuity design,” *The American Economic Review*, 99(1), 179–215.
- WISWALL, M. (2013): “The dynamics of teacher quality,” *Journal of Public Economics*, 100, 61–78.