Essays in Economic Development and Education

by

Benjamin Feigenberg

B.A. Economics and Latin American Studies, Brown University (2006)

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2014

Author ..... Signature redacted ....
Department of Economics
May 15, 2014

Certified by... Signature redacted ....
Esther Duflo
Abdul Latif Jameel Professor of Poverty Alleviation and Development Economics
Thesis Supervisor

Certified by.... Signature redacted ....
Benjamin Olken
Professor of Economics
Thesis Supervisor

Accepted by ...... Signature redacted ....
Michael Greenstone
3M Professor of Economics
Chairman, Departmental Committee on Graduate Studies

Essays in Economic Development and Education

by

Benjamin Feigenberg

Submitted to the Department of Economics
on May 15, 2014, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

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

The second chapter studies the responsiveness of United States-Mexico migration to U.S. border enforcement policy. Spending on border enforcement has risen by 240% in the U. S. in the last decade, and the construction of a fence on the U.S.-Mexico border has become a focal point in the debate over the costs and benefits of increased border security. However, whether and by how much the fence actually reduces migration from Mexico to the U. S. remains an open question. This paper estimates the impact of the fence on migration flows between Mexico and the U.S. and investigates the mechanisms driving observed impacts. To conduct this analysis, I exploit variation in the timing of U.S. government tactical infrastructure investment in fence construction in the period after the passage of the 2006 Secure Fence Act. Using Mexican household survey data and data I collected on border fence construction, I find that construction in a given municipality reduces migration by 39% from that municipality and by 26% from adjacent municipalities. I also find evidence that fence construction reduces migration rates for residents of non-border states with historically low access to smugglers by 38%. Based on these estimates, I calculate that the implied cost of the fence per migrant deterred is $4,800 USD. My findings suggest that fence construction deters migration because the migration costs faced by prospective migrants are sensitive to the particular set of available crossing locations. I derive a simple migration selection model to test this hypothesis and find that a left-censoring of the migration cost distribution, consistent with the disproportionate elimination of low-cost crossing options, best rationalizes evidence on changing migration patterns.

The third chapter of this dissertation (coauthored with Erica Field and Rohini Pande) examines the economic returns to social interaction. For this research, microfinance clients were randomly assigned to repayment groups that met either weekly or monthly during their first loan cycle and then graduated to identical meeting frequency for their second loan. Long-run survey data and a follow-up public goods experiment reveal that clients initially assigned to weekly groups interact more often and exhibit a higher willingness to pool risk with group members from their first loan cycle nearly two years after the experiment. They were also three times less likely to default on their second loan. Evidence from an additional treatment arm shows that, holding meeting frequency fixed, the pattern is insensitive to repayment frequency during the first loan cycle. Taken together, these findings constitute the first experimental evidence on the economic returns to social interaction, and provide an alternative explanation for the success of the group lending model in reducing default risk.

Thesis Supervisor: Esther Duflo
Title: Abdul Latif Jameel Professor of Poverty Alleviation and Development Economics

Thesis Supervisor: Benjamin Olken
Title: Professor of Economics

Contents

Acknowledgements

I am grateful to my advisors David Autor, Esther Duflo, and Ben Olken for their guidance. They have taught me how to conduct research with patience, persistence, and thoughtfulness. They have led by example and instilled in me a passion for discovery and learning that I hope to pass on to my own students.

I am thankful to many other faculty members for the knowledge they have shared with me. I thank Abhijit Banerjee and Paulo Somaini for being so generous with their time and for challenging me to grapple with difficult questions. I have also benefited tremendously from conversations with Daron Acemoglu, Josh Angrist, Melissa Dell, Dave Donaldson, Glenn Ellison, Sara Ellison, Amy Finkelstein, Francisco Gallego, Rachel Glennerster, Michael Greenstone, Rema Hanna, Jeffrey Harris, Nancy Rose, Monica Singhal, Tavneet Suri, Robert Townsend, and Heidi Williams, among others.

I thank Rohini Pande for serving as a mentor from my days as a research assistant through today. I thank her, Amitabh Chandra, Erica Field, Asim Khwaja, Kaivan Munshi, and Nancy Qian for helping me to discover that I wanted to become an economist in the first place. They were kind, they trusted me from the very beginning, and they cared about my own fulfillment and learning in ways that I will not forget.

I am fortunate for the wonderful classmates I have had at MIT. Together, we shared the highs and lows of the last five years. Hopefully their brilliance, passion, and humility have rubbed off on me. I could not imagine another group with whom I would rather have shared this experience, and I count many of them among my closest friends.

Lastly, I thank my family and friends. I am blessed to have a circle of incredible friends from all walks of life who challenge me and motivate me to do things that make the world a better place. I am thankful for my family, who have always been my rock and whom I admire more than words can describe. Finally, I thank my wife, Swapna, for her support. I wake up every day feeling very lucky to have found her.

List of Tables

List of Figures

Chapter 1

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

1.1 Introduction

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

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

In this paper, I propose a model of private school supply and demand which incorporates schools' market power and the fact that school quality (e.g. class size) is a normal good, but adjustments to the number of classrooms in a school are infrequent.[1] Based on this model, I simulate schools' profit-maximizing pricing behavior in a simplified market in which consumers (i.e., parents) experience a positive aggregate income shock. I show that, for a wide range of parameter values, private schools' profit-maximizing strategy is to increase prices to the point that the equilibrium private school enrollment share declines as students switch to public schools. The finding that rising aggregate income could lead private school enrollment to fall in equilibrium with non-trivial probability makes the market for private schooling unusual even among the class of imperfectly competitive markets.

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

Chile provides an ideal setting for studying how the pricing behavior of profit-maximizing schools mediates the impact of municipality-level aggregate shocks that would unambiguously improve outcomes for students in a non-market educational environment. The advantages of the Chilean setting are two-fold. First, its private school market is well-established and covers the whole country. Chile introduced its nationwide school voucher system in 1981 and allows private schools to accept vouchers while charging additional fees. At present, only 7% of Chilean students attend primary or secondary schools which do not participate in the voucher market. Second, the Chilean Ministry of Education collects data on private school enrollment and prices that is not available in other settings, and the Ministry links this data to individual student records. This linkage allows me to investigate how school pricing decisions impact students' enrollment responses and their academic performance.

[1] The model builds on industrial organization demand estimation research, including Berry et al. [1995] and Nevo [2000], that analyzes consumer choice within differentiated products markets. The model is most closely related to Urquiola and Verhoogen [2009], which models education supply and demand in Chile in order to investigate how schools' enrollment and pricing decisions affect estimates of the impact of class size on student outcomes.

Rigorous evidence on demand behavior and supply response within large-scale private school markets more broadly has been limited by data constraints as well as a lack of plausibly exogenous variation in demand for or supply of private schooling that can be isolated and studied.[2] While there is an extensive school voucher literature (see, for instance, Rouse, 1998, Angrist et al., 2002, Krueger and Zhu, 2004, Howell and Peterson, 2006) in which the authors use voucher lotteries to identify the causal impact of gaining access to private school on educational outcomes, voucher experiments have typically taken place in settings in which the group of voucher recipients was small relative to overall private school enrollment. As a result, researchers have been unable to use these experiments to study school price responses and the implications for students' school choices and academic achievement. There is also a large body of research that estimates the causal impact of market competition (i.e., the penetration of private voucher schools) on educational outcomes in Chile using cross-sectional data.[3] By identifying a source of aggregate income shocks at the municipality by year level and by exploiting rich educational panel data, my work sheds light on how schools behave strategically *within* private school markets and how school pricing impacts students.[4] In doing so, this project builds on computational studies of education markets, such as Epple and Romano [1998], Nechyba [2000], and Ferreyra [2007]. In that research, the authors simulate aggregate responses to tuition voucher policies based on varied assumptions about the structure of the market and the determinants of parental demand.

To conduct the reduced-form analysis, I construct plausibly exogenous municipality-level aggregate income shocks using variation in global copper prices and in how income in a given municipality responds to global copper prices. Specifically, I use historical global copper prices and pre-period labor force survey data on household income to construct municipality-specific measures of the elasticity of income with re-

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

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

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

spect to global copper prices from the period before educational microdata is available. For the purposes of exposition, I normalize these municipality-specific elasticities. The interaction of the normalized elasticities with contemporaneous copper prices serves as the aggregate income shock. In the analysis, I focus on the effect of "positive copper shocks," which correspond to copper price increases in municipalities with positive normalized elasticities. Based on this identification strategy, I find that a positive copper shock has a relatively uniform impact across the log income distribution. Importantly, I confirm that the copper shocks constructed based on this methodology do not have any statistically significant impact on private school prices from previous years. I then proceed to estimate the effect of these positive copper shocks on school prices, student enrollment patterns, and student test scores.

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

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

In order to characterize the mechanism driving observed declines in private school enrollment, I proceed

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

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

The remainder of the paper is structured as follows. Section 1.2 presents a model of the private school education market that predicts how private school prices and enrollment levels respond to rising aggregate income. Section 1.3 documents institutional details of the Chilean educational market. Section 1.4 discusses the data used in the empirical analysis, and Section 1.5 outlines the identification strategy. Section 1.6 estimates the impact of positive copper shocks on school prices and enrollment. Section 1.7 investigates the impact of positive copper shocks on student school switching patterns and test scores. Section 1.8 structurally estimates an extended version of the model, and Section 1.9 concludes.

1.2 Model

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

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

$$U_{ij} = \alpha v_i s_j - p_j + \epsilon_{ij} \tag{1.1}$$

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

In this framework, α represents an aggregate willingness to pay shifter, each individual has willingness to pay for quality v_i, and each school j has an associated quality s and price p. As is common in the literature, this expression assumes that schools choose an overall price level but do not have the capacity to price discriminate across students.[7] The error term captures the household's idiosyncratic utility gain from having a child attend school j. The error term has a mean zero Type I extreme value distribution, as in the standard aggregate logit model.[8] The scale parameter of the error term, defined as σ, can be interpreted as a measure of the extent of school differentiation within the market. I assume that willingness to pay has probability density function $f(v)$ and is distributed on the interval $\underline{v}, \bar{v}$. Given these assumptions, the expression for the share of individuals with willingness to pay v who attend school j is defined as:

$$\Gamma_j(\underline{s}|\alpha, v, p, s) = \frac{exp^{\frac{\alpha v s(\underline{s}) - p(\underline{s})}{\sigma}}}{\int_{\hat{\underline{s}}} exp^{\frac{\alpha v s(\hat{\underline{s}}) - p(\hat{\underline{s}})}{\sigma}} g(\hat{\underline{s}}) d\hat{\underline{s}}} \tag{1.2}$$

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

$$D_j(\underline{s}|\alpha, v, p, s) = N \int_{\underline{v}}^{\bar{v}} \Gamma_j(\underline{s}|\alpha, v, p, s) f(v) dv \tag{1.3}$$

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

$$\Pi_j = (p_j - c)q_j - n_j F_n \tag{1.4}$$

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

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

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

8 See, for instance, Nevo [2000].

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

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

satisfied in the Chilean setting (in which I test the predictions of the model).[11] I assume that the functional form of the relationship between school quality and enrollment is characterized as follows:

$$s_j = \underline{s}_j - \tau(\frac{q_j}{n_j}) \tag{1.5}$$

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

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

$$\max_p \Pi(p, q(p, \underline{s})) \tag{1.6}$$

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

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

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

[12] This parameterization of the school quality production function differs from that in Urquiola and Verhoogen [2009]. There, the authors model the impact of classroom size reductions on school quality to be increasing in baseline school quality and proportional to percentage changes in classroom size. My assumption that the impact of classroom size reductions on school quality is independent of baseline quality is motivated by the finding (discussed in Section 1.8) that heterogeneity in parental preferences for reduced classroom size by baseline characteristics is relatively limited. Since wealthier students are more likely to attend schools with higher baseline quality, these students would exhibit stronger preferences for reduced classroom size if associated quality gains were indeed larger in higher baseline quality schools. I assume that quality decreases linearly in classroom size to reduce the incentive for schools to dramatically reduce class size (to near zero) in order to raise prices as this behavior is not consistent with the data. Nonetheless, model simulations look similar if I alternatively assume that school quality changes are proportional to percentage changes in classroom size.

[13] I assume that the number of schools in the market is fixed in the short run. This assumption is supported by the fact that there is no significant change in Chilean school entry/exit in response to aggregate income shocks in the empirical application (see the discussion in Section 1.6.1).

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

prices, enrollment, and expected willingness to pay of parents whose children attend school j:

$$p^* - [\sigma + c + \alpha\tau V^* \frac{q^*}{n}] = 0 \tag{1.7}$$

$$q^* - \int_{\underline{v}}^{\bar{v}} \Gamma_j(\underline{s}|\alpha, v, p^*, s^*) N f(v) dv = 0 \tag{1.8}$$

$$V^* - \int_{\underline{v}}^{\bar{v}} v \frac{\Gamma_j(\underline{s}|\alpha, v, p^*, s^*)}{\int_{\underline{v}}^{\bar{v}} \Gamma_j(\underline{s}|\alpha, v, p^*, s^*) f(v) dv} f(v) dv = 0 \tag{1.9}$$

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

Finally, I apply the implicit function theorem to find expressions for $\frac{dp^*}{d\alpha}$ and $\frac{dq^*}{d\alpha}$. These expressions are derived and presented in the Mathematical Appendix, but they are algebraically complex and provide little intuition regarding the sign or magnitude of price and enrollment impacts. This intractability is a consequence of the fact that all schools adjust prices in response to a change in α. As a result, a school's own price setting decision must take into account the full set of price changes undertaken by other schools in the market. To simplify the problem, I simulate equilibrium price and enrollment changes in a market with one public school, one private school, and 45 students (the maximum classroom size permitted by law). In this market, the public school has a fixed price equal to marginal cost, while the private school adjusts price and enrollment to maximize profits. The private school is assumed to have a baseline quality that is equal to two times the baseline quality of the public school. I present simulation results averaged over 50 replications.

In Figure 1.1, the blue-shaded region characterizes the set of parameter values at which rising aggregate demand leads to a price increase and a decline in enrollment at the private school. The simulations reveal that the model can produce a private school price increase and enrollment decline for a range of values of the classroom size preference parameter (τ) and the horizontal differentiation parameter (σ). Notably, the simulations suggest that both classroom size preferences and horizontal differentiation are necessary to generate a private school enrollment decline. The region in which the simulation produces an enrollment decline is primarily characterized by values of the τ/σ ratio close to 0.1. There are two instances in which simulations do not produce private school enrollment declines. First, when the ratio τ/σ is sufficiently large (denoted by the yellow region of Figure 1.1), the model is unstable as small changes in initial conditions (i.e., the assumed willingness-to-pay distribution) produce very different equilibrium prices and enrollment levels. This is a result of the fact that when classroom size preferences are strong and the market is close to competitive, small changes in class size cause large shifts in enrollment patterns. Second, in the red-shaded region, private school enrollment increases in equilibrium as private school price rises. In the upper portion of the graph, strong differentiation and weak class size preferences prevent students from switching to public school when prices and class size rise. In the bottom portion of the graph, strong class size preferences and relatively weak differentiation constrain the magnitude of price and class size increases such that the net effect on private school enrollment is positive.

1.3 Education in Chile

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

At present, 50% of students in grades one through twelve attend private (voucher) schools, 43% attend public schools, and 7% attend unsubsidized private schools.[15] Under the voucher system, all public primary schools are free, and all public secondary schools can charge at most \$7 USD per month.[16] In contrast, primary and secondary voucher schools are permitted to charge a "top-up" that is up to three times the annual voucher amount. Voucher values are determined by grade level and length of school instruction (full- or half-day), with the voucher value in secondary school depending additionally on whether a student is studying in the vocational or college-preparatory track.[17] In 2011, the average voucher value was \$110 USD, and the average "top-up" charged was \$34 USD. Officially, private schools can selectively admit students while public schools that are not "at capacity" are required to admit all applicants. However, as noted in Gallego and Hernando [2009], school-side screening appears limited based on the following evidence: 93% of parents report that their children attend the parents' preferred school, the average number of schools to which a student applies is 1.1, and only 4% of parents say their child was rejected from at least one school.

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

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

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

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

In Chile, school switching, which represents an important outcome in the analysis, is relatively common: 17% of students switch schools each year, including 12% of students not entering ninth grade. While students are free to attend school in any municipality, high cross-municipality commuting times imply that students' choice set is limited, in practice, to those schools within their home municipality: 83% of all students, and 88% of primary school students, attend school in their home municipality. Previous research, such as Gallego [2013] and Hsieh and Urquiola [2006], has consequently defined education markets by municipality borders.

1.4 Data

1.4.1 Copper Shocks

My identification strategy requires that I estimate municipality-specific elasticities of income with respect to copper prices. To do so, I collect annual global copper prices (denominated in 1998 USD) from the United States Geological Survey within the Department of the Interior [U.S. Geological Survey, July 11, 2013]. Survey data on historical municipality-level incomes comes from the Chilean National Socio-Economic Survey (CASEN) for which data is available for the following years: 1990, 1992, 1994, 1996, 1998, 2000, 2003, 2006, 2009, and 2011. The survey provides a repeated cross-section that lists respondents' municipality of residence. As one of its objectives, the CASEN survey is designed "[t]o characterize the population according to household income, quality of housing, education, participation in the labor market, composition of family income, and other relevant variables" [Ministry of Social Development, 2013]. The CASEN survey is representative at the national, regional, and, in some cases, municipality level. Previous work, such as Auguste and Valenzuela [2006], has also used CASEN survey data to construct municipality-level socioeconomic variables.

1.4.2 Educational Administrative and Survey Data

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

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

The Chilean Ministry of Education also releases annual test score data. Prior to 2006, the Ministry administered the national Educational Quality Measurement System Exam (SIMCE) to one grade level across the country each year, rotating among grades four, eight and ten. After 2006, they tested fourth graders annually and eight and tenth graders biennially, rotating between the two. According to the Ministry, the test is designed to improve educational outcomes by providing an external measure of students' mastery of the curriculum [Agency of Educational Quality, 2013]. In the analysis, I average student scores on the language and math components of the exam and normalize scores at the grade-year level.[19] For each cohort that takes the SIMCE, the Ministry of Education also collects detailed survey data from teachers and parents. The data collected from teachers includes proxies for school quality, such as teacher training, teacher experience, and classroom technology use. The survey of parents provides information on household demographic characteristics, such as household size, approximate household income, and parental education. This measure of parental education plays a central role in the student-level heterogeneity analysis detailed in Section 1.7.

1.5 Identification Strategy

This research examines how private school price and enrollment responses to aggregate income shocks can shed light on demand behavior and supply response in the Chilean voucher market. To generate aggregate

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

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

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

$$Y_{mt} = \alpha_m + \beta_m * P_t + \epsilon_{mt} \tag{1.10}$$

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

Figure 1.2 overlays the estimated municipality-level coefficients onto a map of Chile that includes municipality borders. There does not appear to be any clear pattern in the geographical distribution of coefficient values, although the map does make clear that municipalities without estimated copper coefficients are most likely to be located in the sparsely populated north-south extremes of the country. While the constructed municipality-specific coefficients are only important in relative terms since year fixed effects and municipality fixed effects are included in all subsequent specifications, it may initially seem counterintuitive that the mean municipality-level coefficient is negative-valued. The negative relationship between log incomes and log copper prices during this period appears to be explained by the canonical "Dutch Disease" phenomenon. In the case of Chile, increasing global copper prices are negatively correlated with the Chilean Peso (CLP) to USD exchange rate during the relevant years.[22] Consequently, as global copper prices rise, export-oriented industries suffer. The fact that Chilean exports represent 34% of gross domestic product (compared to, for example, 20% in Argentina) makes the relationship between copper prices and export industry com-

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

[21] β_m is just identified with two observations. However, the 20 municipalities with only two observations are excluded because estimated elasticities for these municipalities are very noisy compared to elasticities estimated based on 3+ observations. Appendix Table 1.1 confirms that results are robust to estimating elasticities at the regional level, in which case data is available for all regions in all years.

[22] See Figure 1.3 for the time-series plots of copper prices and the exchange rate.

petitiveness particularly salient [World Bank, 2013]. Indeed, in the past two years, as copper prices have risen dramatically, exporters concerned about currency appreciation have advocated for capital controls and the Chilean government has invested $12 billion in a program aimed at weakening the currency [Pica and Wisnefski, 2012].

To provide additional evidence that currency appreciation drives the negative income responses to copper price increases calculated for most municipalities, I estimate industry-specific elasticities of income with respect to copper prices using CASEN labor force survey data. The four industries with the negative elasticities largest in magnitude are: timber extraction (-4.5), restaurants (-3.4), construction (-3.3), and "large-scale businesses" (-2.2). The industries with the most positive elasticities are: communications (1.0), public instruction (1.0), garments (2.5), and non-ferrous metals (3.7). It is unsurprising that the non-ferrous metals elasticity is largest in magnitude as this category includes copper production. More generally, the set of industries with negative elasticities is primarily export-oriented or tourism-related (wood products are Chile's second largest export after copper).[23] In contrast, industries with positive elasticities (other than copper) appear to be more domestically-focused. When I estimate the cross-sectional correlation between regional elasticities and regional industry shares, shares of the following industries are positively correlated with local elasticities: agriculture, public administration, and electricity/gas/water.[24] Finance, manufacturing, housing, business, and hotel/restaurant shares are negatively correlated with regional elasticities, while mining is positively correlated with regional elasticities in copper-producing regions and negatively correlated in regions that do not produce copper. These correlations are again consistent with the hypothesis that negative income responses to copper price increases are caused by reduced export competitiveness. Indeed, the value of regional exports relative to total regional economic activity negatively predicts regional elasticities in non-copper producing regions, while the relationship is positive in those regions that produce copper.

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

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

$$Y_{mt} = \alpha + \beta * C_{mt} + \gamma_m + \lambda_t + \epsilon_{mt} \tag{1.11}$$

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

[24] As noted, Appendix Table 1.1 shows that reduced-form estimates generated using regional copper shocks are consistent with estimates from the benchmark specification that identifies impacts using municipality-level variation.

In this equation, γ_m and λ_t represent municipality and year fixed effects, respectively. In this and in all subsequent specifications, standard errors are clustered at the municipality level and calculated by bootstrapping the two-step procedure in which municipality-specific elasticities are first estimated (based on Equation 1.10) and then included in the set of regressors in a second specification (such as Equation 1.11).

Table 1.1 provides evidence that the copper shocks constructed using historical data on incomes and copper prices have a significant impact on household incomes during the period for which educational data is available. Column (1) implies that a one standard deviation positive copper shock is associated with a 2.9% increase in the mean income of municipality residents, and the estimate is statistically significant at the 1% confidence level.[25] Columns (2)-(6) estimate the impact of positive copper shocks on income deciles by replacing the log mean income dependent variable with dependent variables measuring the logs of the 50th through 90th percentiles of the earnings distribution.[26] While estimates differ somewhat across columns, there is no evidence of a monotonic relationship between income decile and estimated copper shock impact that might explain the evidence of heterogeneity in school choice impacts based on family background that is detailed in Section 1.7.

1.6 School Prices and Enrollment

1.6.1 Panel Estimates

When aggregate private school demand rises with aggregate income, equilibrium private school prices and enrollment must adjust in order for the market to clear. As outlined in Section 1.2, the enrollment response to a change in aggregate income is ambiguous. This section estimates both price and enrollment relationships. The benchmark specification used to estimate average school-level price impacts is the following:

$$P_{smt} = \alpha + \beta * C_{mt} + \gamma_s + \lambda_t + \epsilon_{smt} \tag{1.12}$$

P_{smt} is the log mean total price charged by school s in municipality m in year t. This measure is calculated as the log of the sum of the mean "top-up" charged by school s in year t plus the mean voucher value received by school s in year t.[27] The remaining variables are as defined in Equation 1.11, except that municipality

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

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

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

fixed effects are subsumed by school fixed effects, γ_s.[28]

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

$$P_{smt} = \alpha + \beta_1 * C_{mt} + \beta_2 * C_{mt} * V_s + \gamma_s + \lambda_t + \epsilon_{smt} \tag{1.13}$$

Here, V_s is an indicator variable equal to one for private voucher schools. While this specification allows price changes to depend on whether a school is public or private, it does not allow for price changes to vary heterogeneously *within* the private school sector. Since low-end private schools more closely resemble public schools than higher-priced private voucher schools with regards to the students they attract and the prices they charge, it seems likely that average private school price impacts mask substantial heterogeneity. Any such heterogeneity will, in turn, significantly alter the impact of rising aggregate income on students' enrollment decisions. To test for differential school price changes within the private school sector, I define an indicator variable P_{sq} which is equal to one if the tuition charged at baseline (in 2004, the first year for which price data is available) by school s falls into quintile q within municipality m.[29] Then, I estimate the following specification:

$$P_{smt} = \alpha + \beta_1 * C_{mt} + \sum_{q=2}^{5} \beta_q * C_{mt} * P_{sq} + \gamma_s + \lambda_t + \epsilon_{smt} \tag{1.14}$$

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

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

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

[30] Appendix Figure 1.1 plots the distribution of prices by baseline school price quintile for a sample year (2006).

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

[32] In Appendix Table 1.2, I estimate the following specification to provide evidence that this is the case:

$$P_{sgmt} = \alpha + \beta * X_{sgmt} + \gamma_{sg} + \lambda_t + \epsilon_{sgmt} \tag{1.15}$$

Column (1) of Table 1.2 estimates Equation 1.12 and indicates that a one standard deviation positive copper shock causes an average price increase of 0.42% (significant at the 1% confidence level). Column (2) estimates Equation 1.13 and reveals that this small aggregate price change masks a significantly larger 2.74% increase in private school prices. The level term coefficient estimated in Column (2) is not statistically different from zero, reflecting the fact that public school prices do not change in response to copper shocks. Column (3) estimates Equation 1.14 and interaction term coefficients provide evidence that there is substantial heterogeneity in school price impacts based on baseline school price/quality. The coefficients on the interactions between the copper shock measure and price quintiles two and three are small in magnitude. In contrast, the coefficients on the interaction terms for quintiles four and five imply that a one standard deviation positive copper shock leads to a 2.24% increase in private school prices in the fourth quintile of schools and a 4.21% increase in prices in the fifth quintile. These findings indicate that the average private school price impacts identified in Column (2) are driven almost entirely by private schools in the upper half of the baseline price distribution.

Columns (4)-(6) of Table 1.2 re-estimate Equations 1.12-1.14, but replace the dependent variable P_{smt} with E_{smt}, a measure of enrollment in school s in municipality m in year t. Enrollment data is available for the full universe of public and private schools and for a number of years for which school price data is not available. Nonetheless, for the sake of comparability, the school-level sample is restricted to include only those school-year observations for which price data is available. Column (4) estimates mean school-level enrollment impacts and reveals that there is a small but statistically significant increase in average enrollment in response to a positive copper shock (the coefficient of 6.4 is equivalent to a 1.8% enrollment increase). Column (5) reveals that small average enrollment increases mask large enrollment declines in private schools and large increases in public school enrollment. Estimates suggest that average private school enrollment falls by 34.1 students (5.8%) per school and public school enrollment rises by 14.4 students (5.6%) per school in response to a one standard deviation positive copper shock.[33] Column (6) indicates that evidence of price impact heterogeneity is matched by heterogeneous changes in average enrollment. The level term in Column (6) implies that enrollment in quintile one schools rises by 12.9 students (5.2%) per school in response to a one standard deviation positive copper shock. Interaction terms for school price quintiles

where P_{sgmt} is the log price for grade level g in school s in municipality m in year t (defined as the log of the sum of the average school price and the grade-specific voucher value) and γ_{sg} represent school-grade fixed effects. This equation is estimated at the grade four level, as SIMCE data is most frequently available for grade four students. X_{sgmt} represent fourth grade-specific quality measures: mean test scores, mean parental education of classmates, mean log household income of classmates, whether classmates' parents expect them to graduate from college, average years of teacher experience, fraction of certified teachers, fraction of teachers with graduate degrees, and the fraction of teachers that frequently use computers for work. In Appendix Table 1.2, I also estimate the cross-sectional correlation between these characteristics and school prices. Coefficients on student test scores, mean parental education of classmates, mean log household income of classmates, and whether classmates' parents expect them to graduate from college are positive across specifications and are significant at the 1% level in most cases.

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

two and three are positive but not statistically significant, implying that the enrollment impacts in these schools are statistically indistinguishable from those estimated for public schools. In contrast, coefficients on the interaction terms for quintiles four and five are negative and statistically significant. These coefficients indicate that enrollment declines by 26.4 students (4.2%) per school in quintile four schools and by 55.8 students (9.4%) per school in quintile five schools.

Table 1.2 provides evidence that positive copper shocks cause private school prices to rise and enrollments to fall, with the magnitude of impacts increasing in baseline school price. These estimates reflect (scaled) measures of the causal impacts of changes in aggregate income to the extent that copper shocks only affect local educational markets through their impact on aggregate income conditional on year and municipality fixed effects.[34] Even if this exclusion restriction were to be satisfied, however, there are multiple mechanisms that may drive the link between aggregate income and school prices and enrollment levels. In order to gauge the extent to which rising aggregate income impacts school prices by shifting aggregate demand, I investigate the relative importance of alternative channels other than the demand for school quality that I have emphasized here. Seemingly, the most relevant alternative explanation for the link between aggregate income and school prices and enrollment levels is that rising aggregate income affects funding for public schools and so affects quality of and demand for public schools. Previous research indicates, however, that municipal tax revenues do not significantly affect local educational expenditures in Chile [Auguste and Valenzuela, 2006]. Consequently, changes in public school funding in response to copper shocks would likely come from national government sources. While vouchers are the primary mechanism through which central government funds are distributed, voucher values are determined nationally. The most important remaining source of local education funding to consider is the National Fund for Regional Development (FNDR), which distributes funds to select municipalities in order to increase their public school revenues. Appendix Table 1.3 provides evidence that FNDR funding is uncorrelated with municipality-specific copper shocks.

Rising aggregate income may also affect school prices and enrollment levels by changing the marginal cost curves that schools face. The largest component of a school's marginal cost is teacher incomes. In Appendix Table 1.4, I test for teacher income changes in response to positive copper shocks when the data is aggregated up to the regional level. I find that teacher incomes do not change significantly in response to positive copper shocks (coefficients are negative and I can rule out economically significant increases in teacher incomes).[35] While teacher income changes cannot be estimated at the school level based on available data, Appendix Table 1.5 reveals that average teacher experience and the school-level share of certified teachers do not change in response to the aggregate income shock. These characteristics are correlated with

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

35 This analysis is conducted at the regional level given that there are relatively few teachers observed in the CASEN survey data and so municipality-level estimates are quite noisy. As shown in Appendix Table 1.1, price and enrollment results are robust to estimating income elasticities at the regional level.

teacher incomes, and so these findings provide additional evidence that schools' cost curves are not changing significantly in response to aggregate income changes. Appendix Table 1.6 confirms that rental incomes also do not change dramatically in response to positive copper shocks (estimates are marginally significant and smaller in magnitude than estimated average changes in total income). This finding implies that schools' marginal costs are also not significantly affected by changes in the rental prices they face.

A remaining concern is that price and enrollment impacts estimated by including only those schools in operation in a given year may suffer from selection issues if rising aggregate income induces changes in the number of schools in operation in particular municipalities. To assuage this concern, I show in Appendix Table 1.7 that positive copper shocks do not have a statistically significant impact on the number of schools in operation in a given municipality.[36]

Table 1.2 price and enrollment results imply that elite private voucher schools experience revenue losses in response to the aggregate income shock. For these schools' price and enrollment adjustments to be profitable, it must then be the case that reductions in school costs are larger. Case study evidence suggests that personnel costs constitute 87% of total non-capital costs in Chilean schools, and so changes in contracted teacher hours are the key determinant of whether schools can reduce costs in the short run [Ugarte and Williamson, 2012]. In Appendix Table 1.8, Table 1.2 specifications are re-estimated for dependent variables that measure the total number of teacher contract hours per week at a given school and the total number of teachers employed at that school. The results imply that private schools reduce contracted hours by 57.1, on average. These reductions are driven by quintile four and five schools, which decrease contracted hours by 66.5 (8.0%) and 67.6 (8.0%), respectively. On average, private schools decrease the number of teachers employed by 1.6, while private schools in baseline price quintiles four and five reduce the number of employed teachers by 1.9 and 1.7, respectively. In contrast, the average public school increases the number of teachers employed by 0.4. Given that the average teacher contract is for 32 hours, these results indicate that private school cost reductions are explained primarily by reductions in the number of teachers in a school rather than in the number of contracted hours per teacher.[37] Under reasonable assumptions, the magnitude of estimated cost savings in elite private voucher schools implies that school profits are indeed increasing in response to the aggregate income shock.[38]

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

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

38 Private schools in baseline price quintiles four and five reduce revenues by 2.1% and 5.6%, respectively. However, personnel costs are estimated to fall by 8.0% for both sets of schools. If non-personnel costs fall by a similar percentage, then profits in quintile four and five schools rise if markups are lower than 281% and 43%, respectively.

1.6.2 Additional Robustness Checks

Tables 1.3 and 1.4 present results from a series of specifications that test the robustness of price and enrollment estimates to alternative definitions of the aggregate income shock, the sample of schools, and the unit of analysis. Columns (1)-(6) of Table 1.3 re-estimate the specifications from Table 2, but replace the copper shock measure C_{mt} with an average shock that is calculated as follows: First, Equation 1.10 is re-estimated separately using lagged copper values and twice-lagged copper values. Then, the average shock measure is constructed as the average of the current year and two lagged coefficients multiplied by the corresponding current and lagged copper prices. In practice, estimates in Columns (1)-(6) of Table 1.3 closely resemble Table 1.2 estimates (Table 1.3 estimates are slightly larger in magnitude). This suggests that overall price and enrollment impacts are relatively stable across specifications that vary in the time frame over which copper price changes are permitted to induce changes in local school prices and enrollment levels.

Columns (1) and (2) of Table 1.4 re-estimate Equations 1.12 and 1.13, but restrict the sample to exclude the Santiago metropolitan region. These specifications are presented in order to confirm that results are not driven by the country's capital (and largest city), where population density is high and so competition may be strongest. Coefficients are somewhat smaller in magnitude but remain highly significant here and in Columns (3) and (4), which replace price with enrollment as the dependent variable. Column (5) returns to the full sample and confirms that school-level enrollment impacts are robust to the inclusion of a control for the total number of enrolled students in municipality m in year t. Columns (6)-(8) of Table 1.4 estimate enrollment impacts at the municipality level to confirm that school-level estimates scale up to reflect significant changes in the municipality-level share of students attending public school. The specification employed is the following:

$$E_{mt} = \alpha + \beta * C_{mt} + \gamma_m + \lambda_t + \epsilon_{mt} \tag{1.16}$$

Here, E_{mt} is an enrollment measure reflecting either the total number of enrolled students or the number of students enrolled in public schools in municipality m in year t. The remaining variables are as defined in Equation 1.11. In Column (6), the dependent variable is the total number of enrolled students in municipality m in year t. The reported estimate suggests that the impact of positive copper shocks on total enrollment is not statistically significant (although the magnitude of the coefficient is consistent with the average school-level enrollment increase estimated in Column (4) of Table 1.2). In Columns (7) and (8), the dependent variable is the number of students enrolled in public schools in municipality m in year t. The estimated coefficients reflect an increase in public school enrollment of 5.1 to 6.3%.

Finally, as a falsification exercise, Appendix Table 1.9 re-estimates Equation 1.12 with the full sample of private schools, but adds a series of lead terms ($C_{m,t+1}$, $C_{m,t+2}$, $C_{m,t+3}$, and $C_{m,t+4}$) to test whether future copper shocks are correlated with current school prices. None of the lead term coefficients presented in the table are statistically significant or large in magnitude.[39]

[39] When both public and private schools are included in the sample, a subset of the $C_{m,t+1}$ coefficients are statistically

1.7 School Switching and Test Scores

1.7.1 School Switching Patterns

The evidence on average enrollment impacts reveals that aggregate income shocks induce movement from high-cost private schools to low-cost private schools and public schools. However, these results only tell half of the story. While rates of elite private school attendance at baseline are rising in students' socioeconomic status, the marginal students who move from elite private schools to public schools and low-quality private schools need not be representative of their classmates. To the extent that policymakers are concerned about inequality of educational opportunities, identifying what the characteristics of these marginal students tell us about changes in social stratification is of particular policy relevance.

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

$$\begin{aligned} Y_{ismt} = \alpha + \beta_1 * C_{mt} + \sum_{k=2}^{4} \beta_k * C_{mt} * X_{ik} \\ + \sum_{k=2}^{4} \delta_k * X_{ik} + \gamma_m + \lambda_t + \epsilon_{ismt} \end{aligned} \tag{1.17}$$

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

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

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

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

42 One important distinction between the school-level and student-level specifications is that the municipality corresponding to each observation in the student-level specifications reflects the municipality in which the students lives in the first year in which he appears in the administrative records. Consequently, the structure of fixed effects is complicated by the fact

k.

Column (1) of Table 1.5 uses the specification outlined in Equation 1.17 to estimate the impact of positive copper shocks on the baseline school price quintile of the school that a student attends. Here, estimates indicate that students in parental education quartiles two and three experience the largest relative decreases in school quality, while students in quartile four experience smaller but still statistically significant relative declines.[43] Since middle and high parental education students attend more elite schools at baseline, these estimates may be explained by the fact that these students are simply more likely to end up in worse schools because they are more likely to be at the margin of attending schools that raise their prices significantly. To estimate average changes in school quality in year t conditional on being exposed to the *same* price increase, Column (2) adds lagged school by municipality fixed effects to the specification estimated in Column (1). These fixed effects are defined such that, for an observation in year t for individual i, a fixed effect is included for the school that individual i attended in year $t-1$ in combination with the municipality in which individual i lived in year $t-1$. In Column (2), the interaction terms for parental education quartiles two and three remain negative and statistically significant but are only 10% as large as Column (1) estimates, while the interaction term for quartile four is no longer statistically significant and is even smaller in magnitude (only 2% of the size of the corresponding Column (1) coefficient). These estimates suggest that the larger average price increases that middle parental education students face (relative to low parental education students) explain nearly all of their differential reduction in school quality.[44] In sum, middle parental education students are most affected by copper shocks because high parental education students are more willing to pay to attend elite schools even when tuitions rise, while low parental education students did not consider attending elite private schools in the first place.

Observed declines in school quality can be driven by either an increase in school "downgrading" or a decrease in school "upgrading," where downgrading is defined by whether a student attends a school in year t that has a lower baseline price quintile measure than the school she attended in year $t-1$ and upgrading

that school fixed effects no longer subsume municipality fixed effects. This distinction will be discussed in more depth in the subsequent analysis of student test score impacts.

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

44 Nonetheless, interaction coefficients remain negative in Column (2). This likely reflects the fact that, even conditional on attending the same school in year $t-1$, higher parental education quartile students are ex-ante more likely to attend higher baseline price schools in year t and so there is a greater margin for reducing school quality.

is defined by whether a student attends a school in year t that has a higher baseline price quintile measure than the school she attended in year $t-1$. Columns (3) and (4) of Table 1.5 examine the impacts of positive copper shocks on indicator variables for whether a student upgrades and downgrades, respectively. Estimates imply that reductions in school quality are driven by a combination of reduced upgrading and increased downgrading (although estimated reductions in upgrading are somewhat larger in magnitude).

The fact that increased enrollment in public and low-cost private schools is driven by middle parental education students suggests that these same students will likely experience the largest relative declines in peer quality. To measure changes in peer quality and social stratification, I focus on changes in school-level mean parental education. Column (5) of Table 1.5 estimates Equation 1.17 and reveals that middle parental education students do indeed experience significant relative declines in school-level mean parental education. However, estimates also reveal that all subgroups experience overall increases in this measure. Column (6) re-estimates Equation 1.14 using mean parental education as the dependent variable, and the results reveal that all students experience an average increase in peer quality because all segments of the market experience non-decreasing average parental education (with parental education rising significantly in quintiles four and five).[45] Estimated impacts on school-level mean parental education suggest that changes in social stratification are small relative to average increases in peer quality in response to rising aggregate income. However, the fact that school quality is rising in peer quality but falling in classroom size suggests that school quality may nonetheless be falling in a subset of schools.

1.7.2 Test Scores

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

Column (1) of Table 1.6 estimates the overall impact of a positive copper shock on test scores using the

[45] The possibility of a positive overall effect on mean parental education may initially appear puzzling given that the overall distribution of parental educational attainment is not shifting in response to positive copper shocks experienced by municipalities. However, an increase in mean parental educational attainment can be rationalized by strong ex-ante sorting across schools based on parental education in combination with higher rates of school downgrading by those who would have been at the bottom of the parental education distribution within their counterfactual schools. As a simplified example, suppose that school A is composed half of students from parental education quartile one and half from quartile two, while school B is composed half of students from quartile three and half from quartile four. Then, if all school downgrading that occurs is driven by quartile three students moving from school B to school A, mean parental educational attainment in both schools will rise.

following student-level specification:

$$T_{ismt} = \alpha + \beta * C_{mt} + \gamma_m + \lambda_t \tag{1.18}$$

Here, T_{ismt} is the normalized SIMCE score of student i who lives in municipality m and attends school s in year t. Column (1) estimates reveal that a positive copper shock increases average test scores in a municipality by 0.020 standard deviations. Column (2) examines whether test score impacts vary across parental education quartiles by re-estimating Equation 1.17 with normalized test scores as the dependent variable. Column (2) estimates reveal that test score gains are largest for students from the lowest and highest parental education quartiles. While impacts remain positive for all subgroups, the magnitude of test score increases is approximately 20% lower for students in parental education quartiles two and three relative to quartile one students. Students in parental education quartile four also experience (smaller) marginally significant test score declines relative to quartile one students.

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

To disentangle these two possible explanations for differential test score gains, I return to the specification from Column (2) of Table 1.5 that adds lagged school by municipality fixed effects. Estimates, presented in Column (3) of Table 1.6, provide evidence that the impact of positive copper shocks on school choice drives the differential test score gains that are estimated in Column (2) of that table. The coefficient on the positive copper shock level term of 0.025 is statistically significant at the 1% confidence level and similar in magnitude to the corresponding coefficient in Column (2). However, the coefficient on the parental education quartile two interaction term is no longer statistically significant, while the interaction terms for quartiles three and four are now positive and statistically significant at the 10% and 5% levels, respectively. Column (4) replaces lagged school by municipality fixed effects with school by municipality fixed effects for the school attended in year t and estimated coefficients are similar.

The school that an individual attends is presumably correlated with unobservable characteristics, such as academic motivation, and this implies that estimates in Columns (3) and (4) cannot be interpreted as revealing the purely causal impact of copper shocks on test scores that is independent of school choice.

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

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

1.8 Structural Estimation

The model presented in Section 1.2 suggested that, under certain conditions, the residual demand curves faced by private school administrators are sufficiently inelastic that their profit-maximizing response to increased demand for private schooling is to raise prices so substantially that equilibrium private school enrollment falls. The empirical results presented in Sections 1.6 and 1.7 provide evidence that these market conditions are satisfied in the Chilean setting and that students' academic performance and educational opportunities are significantly affected by private schools' pricing behavior. In this section, I estimate an extended version of the model in order to characterize average willingness to pay for particular school attributes and heterogeneity in willingness to pay for school attributes based on individual characteristics. Simulated school price and enrollment responses to rising aggregate income based on these parameter estimates then allow me to quantify the relative importance of the market power and classroom size preferences channels outlined in Section 1.2.

The estimation approach uses both micro and aggregate data in order to produce parameter values that match schools' predicted market shares to observed ones. This approach was introduced in Berry et al.

[47] Appendix Table 1.10 confirms that evaluations of teacher performance do not change in response to aggregate income shocks. This suggests that changes in teacher effort also do not appear to play a central role in explaining observed test score changes. To provide additional evidence of differential effects on academic performance, Appendix Table 1.11 shows impacts on students' grade point averages. Changes in students' grade point averages mirror estimated test score impacts.

[48] This figure is based on the following calculation: relative to low parental education students, those in parental education quartiles two and three attend a school that is 0.014 baseline price quintile points lower than the school they would have attended absent the shock (based on Column (1) of Table 1.5). Relative to low parental education students, these quartile two and three students experience a corresponding test score increase in response to a positive copper shock that is 0.0034 standard deviations smaller. Dividing the second figure by the first, I estimate that a one-point downgrade in the baseline school price quintile distribution is associated with a 0.24 standard deviation decrease in test scores. Since this estimated decrease is relative to an average increase of 0.02 standard deviations for low parental education students, the implied aggregate test score reduction for downgraded students is 0.22 standard deviations.

[2004]. The algorithm most closely follows Gallego and Hernando [2009], which first applied the Berry et al. [2004] approach to the Chilean school choice decision. The equation for household utility that is used to generate parameter estimates is the following:

$$u_{ijt} = \frac{q_{jt}}{n_{jt}}\beta_{i1} - p_{jt}\beta_{i2} + \sum_{k=3}^{4} X_{jkt}\beta_{ik} + \lambda_{ijt}\sigma_i + \gamma_j + \xi_{jt} + \varepsilon_{ijt} \tag{1.19}$$

Here, u_{ijt} is the utility of student i who chooses school j at time t. Average classroom size and price are represented by $\frac{q_{jt}}{n_{jt}}$ and p_{jt}, respectively. X_{jkt} are additional observable characteristics k of school j at time t, and λ_{ijt} is the distance from the center of individual i's home municipality to school j. γ_j is a time-invariant school fixed effect and ξ_{jt} represents unobservable time-varying school-specific characteristics.[49] ε_{ijt} is a random preference shock (error term) that is assumed to have a Type I extreme value distribution. This expression nests Equation 1.1 where school quality is expressed as a scalar s and individuals differ in their willingness to pay for school quality. The utility function here, however, exploits the availability of panel data in which case time-invariant school fixed effects are identified and baseline school quality is subsumed into the school fixed effect.

To analyze the conditions under which Equation 1.19 is identified, it is informative to characterize heterogeneity in preferences for school attributes based on observable individual characteristics. In doing so, I follow the notation from Berry et al. [2004] and rewrite β_{ik} as $\bar{\beta}_k + \sum_r d_{ir}\beta^o_{kr}$.[50] Here, r indexes individual characteristics and β^o_{kr} captures interactions between observed individual and school characteristics. I can then decompose Equation 1.19 into (1) a component characterizing the mean utility associated with a given school in a particular year (δ_{jt}), and (2) a component that captures preference heterogeneity:

$$u_{ijt} = \delta_{jt} + \sum_{kr} X_{jkt}d_{ir}\beta^o_{kr} + \lambda_{ijt}\bar{\sigma} + \sum_{r} \lambda_{ijt}d_{ir}\sigma^o_r + \varepsilon_{ijt} \tag{1.20}$$

where

$$\delta_{jt} = \sum_k X_{jkt}\bar{\beta}_k + \gamma_j + \xi_{jt} \tag{1.21}$$

Equation 1.21 highlights the fact that $\bar{\beta}_k$ and ξ_{jt} cannot be separately identified without further parametric assumptions. Consequently, I follow the two-step procedure employed in Berry et al. [2004]. The first step estimates the δ, β^o, $\bar{\sigma}$, and σ^o parameters without imposing any additional structure on the ξ_{jt} coefficients. This ensures that δ, β^o, $\bar{\sigma}$, and σ^o are consistently estimated. The second step of the algorithm requires a set of instruments Z_{jt} that satisfy the assumption $E[\xi_{jt}|Z_{jt}] = 0$ and so allow the $\bar{\beta}_k$ and ξ_{jt} coefficients to be consistently estimated.

In estimating Equation 1.19, I include fourth grade students in the Santiago metropolitan region in

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

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

the years 2006-2009. I choose fourth graders because geographic coordinates are more widely available for primary schools than for secondary schools. The sample is restricted to the Santiago metropolitan region in the years 2006-2009 given the computational intensity of the estimation algorithm.[51] In addition to average classroom size and price, the included school characteristics (i.e., the set of X_{jkt}) are indicators for whether school j is public or a low-priced private school (in baseline price quintile two or three). The included individual characteristics (i.e., the set of d_{ir}) are: parental education, parental income, imputed income change in response to a positive copper shock, and whether parents expect their child to graduate from college. The imputed income response to copper shock measure is constructed as 2.9% of income (the average change in income estimated in Table 1.1) multiplied by the copper shock corresponding to a student's home municipality in a given year. Heterogeneity in willingness to pay based on this imputed income shock is a key parameter of interest given the fact that cross-sectional differences in household income are not exogenous. As a result, changes in preferences based on cross-sectional income differences may be quite different from the changes induced when an individual experiences an income shock (and it is this within-individual income variation that is of interest for the simulations). In the analysis, observations are dropped for students with missing values for any of the individual variables and for schools with missing SIMCE survey responses (which are used to determine school attributes).

The maximum likelihood approach I employ in step one of the algorithm searches for the δ, β^o, $\bar{\sigma}$, and σ^o parameter values that maximize the sum of the predicted log probability that each student attends the school they attend in practice. Initially, a guess is chosen for the β^o, $\bar{\sigma}$, and σ^o parameter values. Then, the δ that maximizes the log likelihood function is calculated conditional on the chosen β^o, $\bar{\sigma}$, and σ^o values.[52] Calculating the maximum achievable log likelihood for each choice of the β^o, $\bar{\sigma}$, and σ^o parameters, the algorithm searches for the β^o, $\bar{\sigma}$, and σ^o values that globally maximize the log likelihood function. The advantage of this procedure over one that jointly searches for the likelihood-maximizing values of δ, β^o, $\bar{\sigma}$, and σ^o is that it eases the computational burden by reducing the dimensionality of the problem [Nevo, 2000]. In estimating parameters, I use a gradient-based algorithm in combination with the Nelder-Mead Simplex method to improve the speed with which likelihood-maximizing parameter values are found. This gradient-based approach requires that the gradient of the log likelihood function be calculated explicitly,

[51] This constitutes a subset of the full sample analyzed in Section 1.6-1.7. The Santiago metropolitan region serves 41% of Chilean students and includes 24% of schools and 20% of municipalities. Appendix Table 1.12 re-estimates Table 1.2 specifications for the Santiago metropolitan region subsample. Price and enrollment impacts are broadly similar to full sample estimates: private school prices increase, private school enrollments decline, and both price and enrollment impacts are largest at elite private schools. Initial estimates identify price decreases in public schools in response to copper shocks, but this is driven entirely by the pooled construction of year fixed effects. Public school price impacts are insignificant (with a coefficient of 0.001) when I estimate price changes separately for public and private schools or include separate year fixed effects for public and private schools in a pooled specification.

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

and I make use of the expressions for the gradient derived in the Computational Appendix of Gallego and Hernando [2009].

Table 1.7 presents the likelihood-maximizing β^o, $\bar{\sigma}$, and σ^o parameter values along with estimates of β_1 and β_2.[53] Coefficients in Table 1.7 reflect the change in utility associated with changes in school and individual characteristics. Estimates related to preferences for reduced classroom size suggest that wealthier and more educated parents actually value classroom size slightly less than their peers, while those experiencing a positive income shock increase their valuation of small class size. Higher income parents, better educated parents, those with high expectations for their children, and those experiencing positive income shocks all appear to care less about school prices than their peers. This reduced price sensitivity and increased demand for small class size by households experiencing a positive income shock will play an important role in determining school price responses to rising aggregate income in the simulation presented below.

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

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

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

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

the derivation of residual demand elasticity, I employ the following formula:

$$\hat{\nu}_{jp} = \frac{\sum_{i=1}^{N}[\hat{P}_{ij}(1-\hat{P}_{ij})(\bar{\beta}_2 + \sum_r d_{ir}\beta^o_{r2})]}{1 - \frac{1}{n_j}\sum_{i=1}^{N}[\hat{P}_{ij}(1-\hat{P}_{ij})(\bar{\beta}_1 + \sum_r d_{ir}\beta^o_{r1})]} * \frac{p_j}{q_j} \tag{1.22}$$

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

Given estimates of the elasticity of residual demand faced by schools, I apply the standard monopolist markup formula to back out schools' marginal costs at baseline enrollment levels.[55] With a full set of estimated demand parameters and school marginal costs, I am prepared to simulate school price and enrollment responses to an aggregate income shock. To do so, I impose a uniform (2.9%) copper shock-induced income change. This imputed income change is set equal to the change in mean income estimated in Table 1.1 in response to a copper shock of one standard deviation. Consequently, if I have correctly modeled the structure of the market, I should simulate price and enrollment changes that mirror the reduced-form estimates presented in Table 1.2.

To estimate equilibrium changes in private school prices and enrollment, I follow an iterative procedure. In each iteration, I construct for each school a grid of five prices (centered around the price chosen in the previous iteration) with corresponding changes in enrollment and marginal cost. Enrollment changes are calculated by linearizing the logit function. To compute marginal cost, I estimate cross-sectionally the relationship between marginal cost and class size within each baseline school price quintile bin. Then, I predict marginal cost at a given enrollment level as a function of a school's baseline marginal cost and this slope. Given expressions for school price, enrollment, and average variable cost at each point on the grid, each school is assigned the grid point corresponding to the profit-maximizing price.[56] After each iteration in which all schools have chosen profit-maximizing prices, students are re-sorted across schools based on these new prices and endogenously determined average classroom sizes.[57] This process then repeats until I arrive at a Nash equilibrium in which all schools are choosing the local best response to all other schools' prices.

Table 1.8 presents the results from this simulation algorithm for the 783 schools in the Santiago metropolitan region in 2006. The price and enrollment results provide a reasonable approximation of reduced-form

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

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

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

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

1.9 Conclusions

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

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

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

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

Table 1.1. Copper Shocks and Income

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

Notes

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

Table 1.2. School Price and Enrollment Responses to Copper Shocks

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

Notes

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

Table 1.3. School Price and Enrollment Robustness Specifications I

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

Notes

Average Copper Shock is defined as the mean of current, lagged, and twice lagged copper shocks (constructed as described in Table 1.1). Regressions are clustered at the municipality level (there are 270 municipalities). * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table 1.4. School Price and Enrollment Robustness Specifications II

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

Notes

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

Table 1.5. Copper Shocks and Student School Sorting Patterns

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

Notes

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

Table 1.6. Copper Shocks and Student Test Scores

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

Notes

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

Table 1.7. Preferences for School Attributes

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

Notes

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

Table 1.8. Simulated Enrollment Changes

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

Notes

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

Figure 1.1. Model Simulations

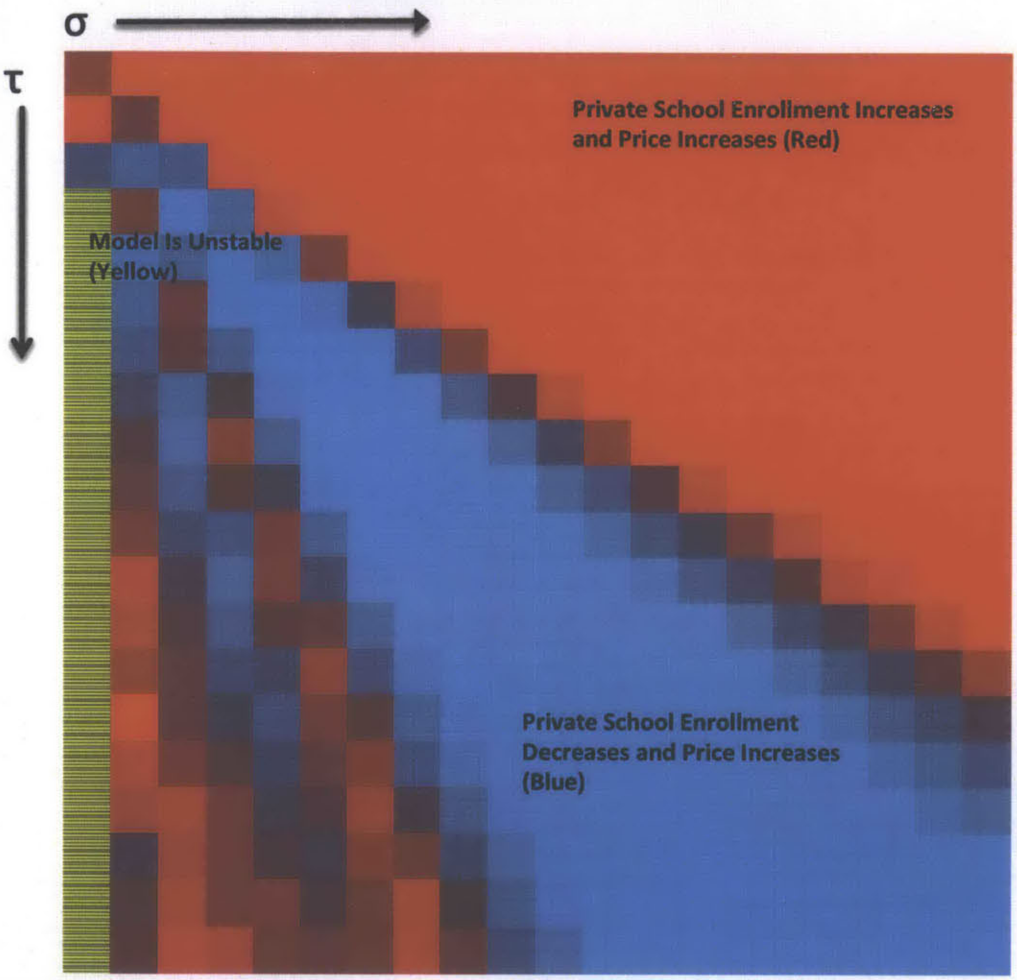

Notes: The figure displays results averaged over 50 model simulations. Y-axis values range from τ=0.25 to τ=5.0 and X-axis values range from σ=2.5 to σ=50.0. Baseline quality ratio between schools is set equal to 2. When printed in grayscale, darker cells are those in which private school enrollment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

Figure 1.2. Municipality Specific Elasticities

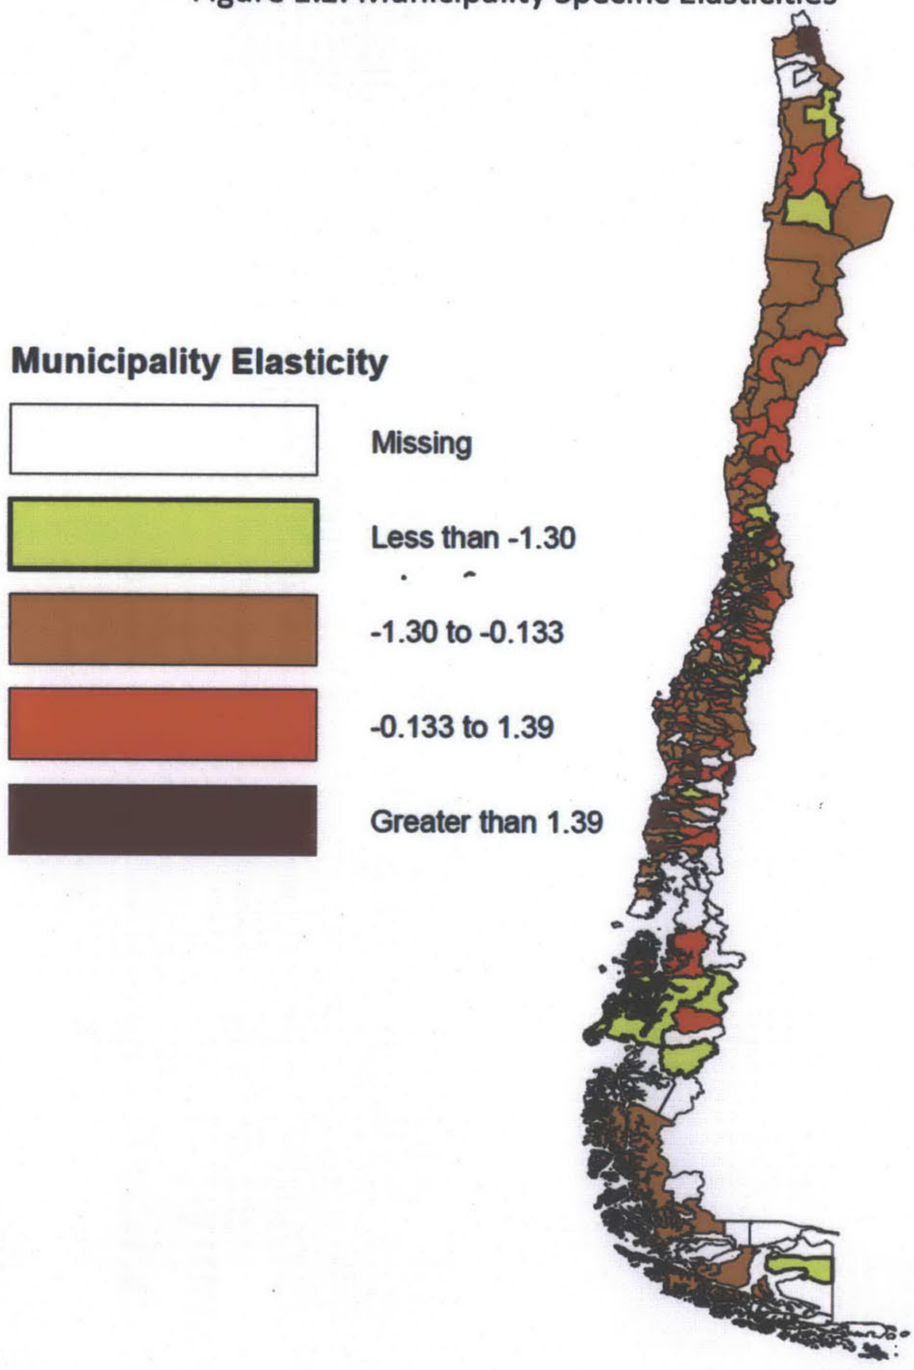

Figure 1.3. Time Series of Copper Prices and Chilean Exchange Rate

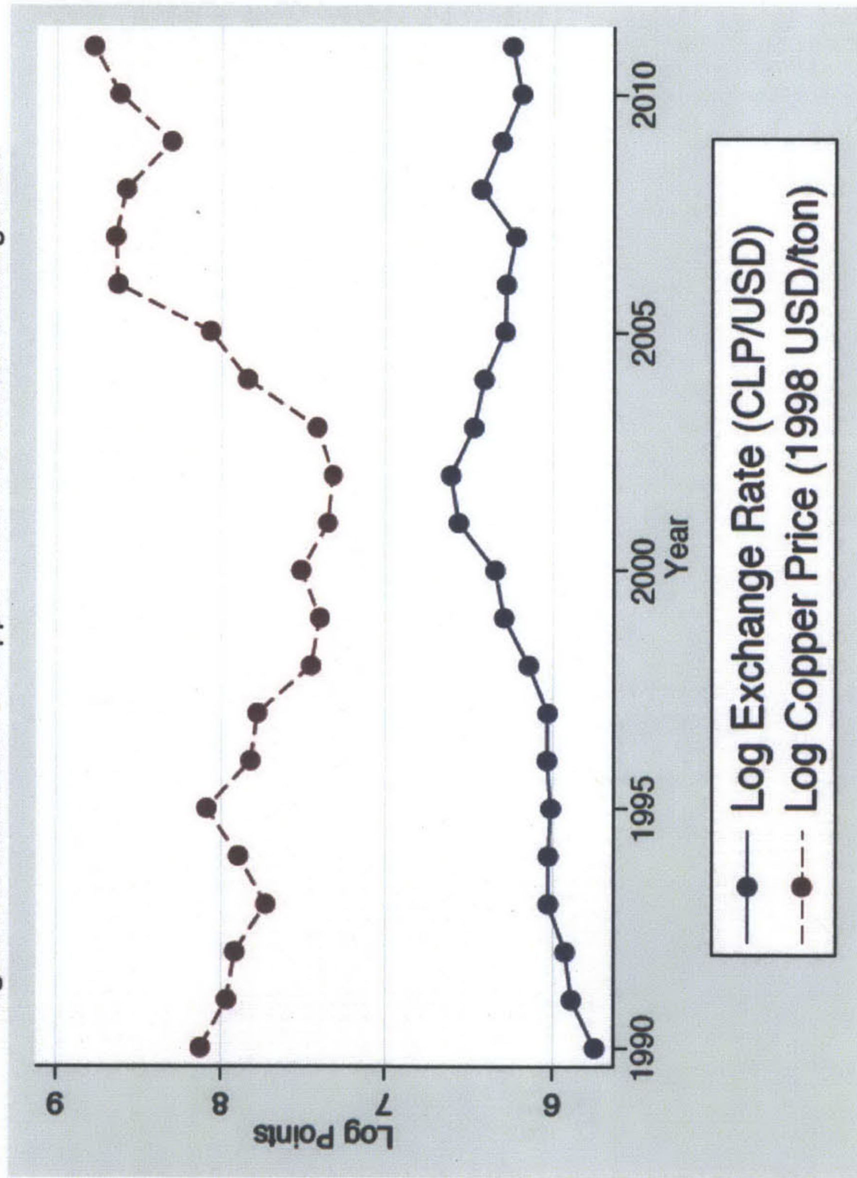

Appendix Table 1.1. Region-level Analysis

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

Notes

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

Appendix Table 1.2. School Price Determinants

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

Notes

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

Appendix Table 1.3. Copper Shocks and Public School Funding

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

Notes

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

Appendix Table 1.4. Teacher Income

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

Notes

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

Appendix Table 1.5. Copper Shocks and Teacher Characteristics

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

Notes

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

Appendix Table 1.6. Copper Shocks and Rental Income

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

Notes

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

Appendix Table 1.7. Copper Shocks and School Entry/Exit

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

Notes

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

Appendix Table 1.8. Copper Shocks and Teacher Contracts

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

Notes

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

Appendix Table 1.9. Lead Structure of Copper Shocks

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

Notes

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

Appendix Table 1.10. Copper Shocks and Teacher Performance

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

Notes

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

Appendix Table 1.11. Copper Shocks and Student Grade Point Averages

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

Notes

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

Appendix Table 1.12. School Price and Enrollment Responses to Copper Shocks (Santiago Region)

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

Notes

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

Appendix Figure 1.1. School Price Distribution

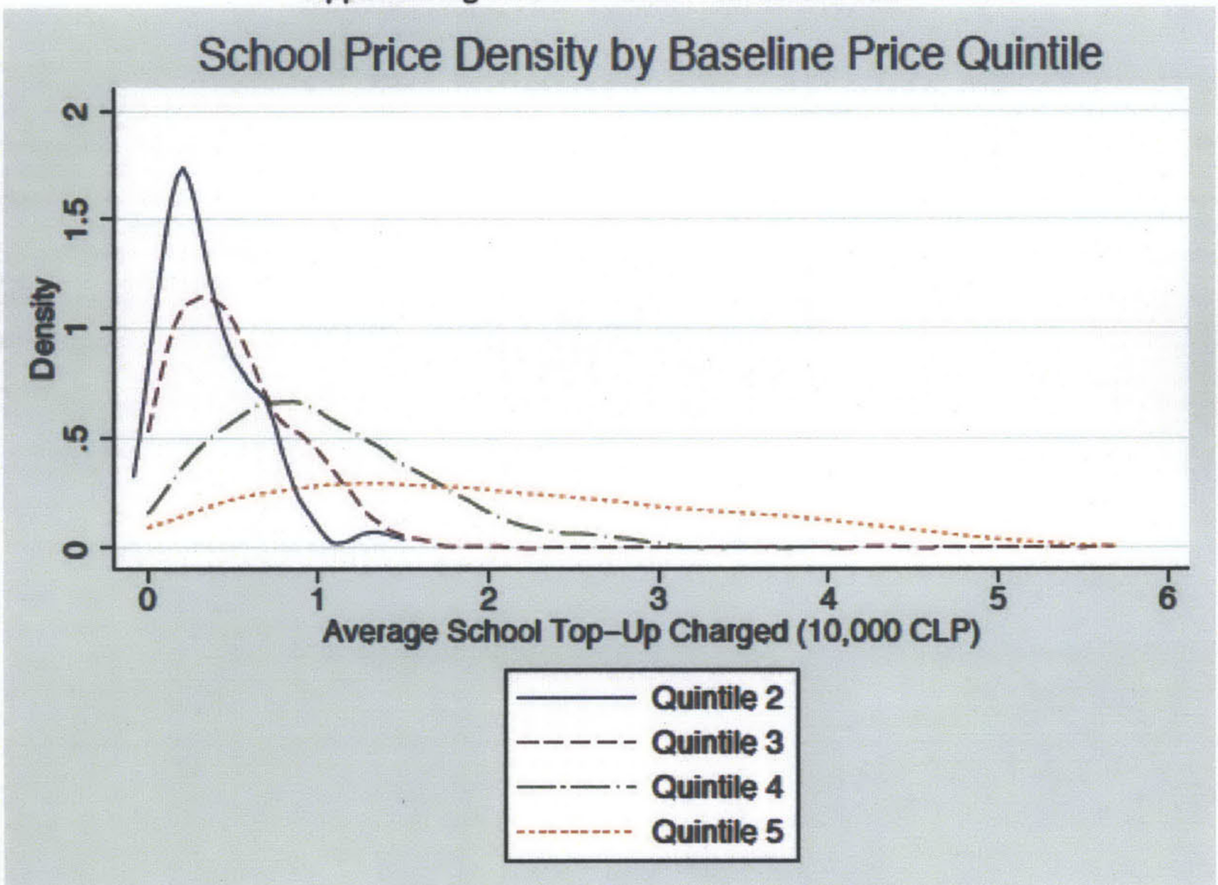

Note: Quintile 1 is excluded because 99.8% of quintile 1 schools do not charge any top-up.

1.10 Mathematical Appendix

Based on the first-order condition from the school's profit maximization and the expressions for demand and expected willingness to pay of parents that are presented in Section 1.2, I arrive at a set of three equations characterizing equilibrium prices, enrollment, and expected willingness to pay of parents whose children attend school j:

$$p^* - [\sigma + c + \alpha\tau V^* \frac{q^*}{n}] = 0 \tag{1.23}$$

$$q^* - \int_{\underline{v}}^{\bar{v}} \Gamma_j(\underline{s}|\alpha, v, p^*, s^*) N f(v) dv = 0 \tag{1.24}$$

$$V^* - \int_{\underline{v}}^{\bar{v}} v \frac{\Gamma_j(\underline{s}|\alpha, v, p^*, s^*)}{\int_{\underline{v}}^{\bar{v}} \Gamma_j(\underline{s}|\alpha, v, p^*, s^*) f(v) dv} f(v) dv = 0 \tag{1.25}$$

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

I can then apply the implicit function theorem to find expressions for $\frac{dp^*}{d\alpha}$ and $\frac{dq^*}{d\alpha}$. To do so, I construct the inverse of the Jacobian and multiply it by negative one times the vector of the partial derivatives of Equations 1.23-1.25 with respect to α. The Jacobian matrix of the partial derivatives of Equations 1.23-1.25 with respect to the equilibrium values of the three endogenous variables (p, q, and V) is as follows:

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

The inverse of the Jacobian can then be expressed as:

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

Finally, the vector of the partial derivatives of Equations 1.23-1.25 with respect to α is:

$$P = \begin{bmatrix} \frac{dEq1.23}{d\alpha} \\ \frac{dEq1.24}{d\alpha} \\ \frac{dEq1.25}{d\alpha} \end{bmatrix} = \begin{bmatrix} -\tau V^* \frac{q^*}{n} \\ -\frac{(\underline{s} - \tau\frac{q}{n}) q^* V^*}{\sigma} + \frac{N}{\sigma} \int_{\underline{v}}^{\bar{v}} v \frac{exp\frac{\alpha v s(\underline{s}) - p(\underline{s})}{\sigma} \int_{\underline{\hat{s}}} exp\frac{\alpha v s(\underline{\hat{s}}) - p(\underline{\hat{s}})}{\sigma} s(\underline{\hat{s}}) g(\underline{\hat{s}}) d\underline{\hat{s}}}{(\int_{\underline{\hat{s}}} exp\frac{\alpha v s(\underline{\hat{s}}) - p(\underline{\hat{s}})}{\sigma} g(\underline{\hat{s}}) d\underline{\hat{s}})^2} f(v) dv \\ -\frac{(\underline{s} - \tau\frac{q}{n}) E[v^2]}{\sigma} + \frac{1}{\sigma} \int_{\underline{v}}^{\bar{v}} v^2 \frac{exp\frac{\alpha v s(\underline{s}) - p(\underline{s})}{\sigma} \int_{\underline{\hat{s}}} exp\frac{\alpha v s(\underline{\hat{s}}) - p(\underline{\hat{s}})}{\sigma} s(\underline{\hat{s}}) g(\underline{\hat{s}}) d\underline{\hat{s}}}{(\int_{\underline{\hat{s}}} exp\frac{\alpha v s(\underline{\hat{s}}) - p(\underline{\hat{s}})}{\sigma} g(\underline{\hat{s}}) d\underline{\hat{s}})^2 \int_{\underline{v}}^{\bar{v}} \Gamma_j(\underline{s}|\alpha, v, p, s) f(v) dv} f(v) dv \\ + \int_{\underline{v}}^{\bar{v}} v \frac{\Gamma_j(\underline{s}|\alpha, v, p, s) \int_{\underline{v}}^{\bar{v}} \frac{v(\underline{s} - \tau\frac{q}{n})}{\sigma} \frac{exp\frac{\alpha v s(\underline{s}) - p(\underline{s})}{\sigma}}{\int_{\underline{\hat{s}}} \frac{\alpha v s(\underline{\hat{s}}) - p(\underline{\hat{s}})}{\sigma} g(\underline{\hat{s}}) d\underline{\hat{s}}} f(v) dv}{(\int_{\underline{v}}^{\bar{v}} \Gamma_j(\underline{s}|\alpha, v, p, s) f(v) dv)^2} f(v) dv \\ - \int_{\underline{v}}^{\bar{v}} v \frac{\Gamma_j(\underline{s}|\alpha, v, p, s) \int_{\underline{v}}^{\bar{v}} \frac{exp\frac{\alpha v s(\underline{s}) - p(\underline{s})}{\sigma}}{(\int_{\underline{\hat{s}}} \frac{\alpha v s(\underline{\hat{s}}) - p(\underline{\hat{s}})}{\sigma} g(\underline{\hat{s}}) d\underline{\hat{s}})^2} \int_{\underline{\hat{s}}} \frac{v s(\underline{\hat{s}})}{\sigma} exp\frac{\alpha v s(\underline{\hat{s}}) - p(\underline{\hat{s}})}{\sigma} g(\underline{\hat{s}}) d\underline{\hat{s}} f(v) dv}{(\int_{\underline{v}}^{\bar{v}} \Gamma_j(\underline{s}|\alpha, v, p, s) f(v) dv)^2} f(v) dv \end{bmatrix}$$

The intractability of the vector of partial derivatives with respect to α in turn implies that the expressions for price and enrollment comparative statics, which are equal to the first and second entries of the matrix $-J^{-1} * P$, are uninterpretable. To gain insight into the nature of price and enrollment responses to changing aggregate demand, I conduct simulations of the model (as detailed in Section 1.2).

Chapter 2

Fenced Out: Why Rising Migration Costs Matter

2.1 Introduction

In the last two decades, the total number of global migrants has increased by approximately 37%, to 214 million [UN Population Division, 2009]. In response to this rise in aggregate migration, the world's destination countries have paid increasing attention to border security and the regulation of migratory inflows. In the United States, tough talk on curbing illegal immigration has become commonplace and has been accompanied by a 240% increase in Border Patrol spending over the past decade, to $3.6 billion in Fiscal Year 2010 [Haddal, 2010]. U.S. efforts to deter illegal immigration have focused almost entirely on increasing the associated costs (by increasing policing at the border) rather than reducing the labor market returns to illegal immigration (for instance, by reducing employment opportunities for undocumented immigrants once they have entered U.S. territory). Despite this uni-dimensional approach, we know surprisingly little about the cost effectiveness of increased border enforcement, and we know even less about the mechanisms and subpopulations driving any changes in migration patterns. Studying the effectiveness of border enforcement activities is particularly policy-relevant at present given that recently-proposed immigration legislation requires that the U.S.-Mexico border be classified as "secured" before undocumented immigrants already in the United States are allowed to apply for permanent residence [Preston, 2013].

Research on the impact of changing migration costs has been limited by the difficulty of identifying exogenous shocks to costs. In the context of border enforcement, macroeconomic factors typically co-determine both the allocation of governmental resources and the labor market opportunities available to prospective migrants. While government representatives tasked with securing our borders frequently attribute reductions in the number of apprehended migrants to the improved efficacy of border enforcement activities [Napolitano, 2011], we generally cannot rule out the possibility that cited reductions in cross-border migration are, for instance, driven by business cycle fluctuations that have little to do with border enforcement resource allocation.

In this paper, I use a novel dataset to accurately estimate the effect of changing costs on migration from Mexico to the United States. In addition, I estimate how the impact of changing costs varies based on migrants' access to alternative crossing locations and to smugglers, as well as their demographic characteristics. To conduct this analysis, I exploit plausibly exogenous variation in the geographic distribution and timing of Department of Homeland Security (DHS) and Customs and Border Patrol (CBP) tactical infrastructure investment. In particular, I study the 2006 Secure Fence Act, which authorized the construction of 698 miles of double-layered pedestrian fence along the United States-Mexico border. Documentation related to border fence construction reveals that the process of land acquisition, environmental evaluation, and construction design and contracting that preceded each construction project generated variation in the timing of construction across border municipalities. Empirically, I find no evidence of municipality-level trends in migration rates or household income that can predict when a particular municipality was fenced (or whether it was fenced at all).

Based on this identification strategy, I find that fence construction reduces the migration rate to the United States of residents from the municipality that is fenced by 39% (relative to a baseline rate of 2.5 migration episodes per 1,000 residents per quarter). Using the same household survey data, I find no significant impact of fence construction on return migration by those living in the United States who migrated from a municipality in which fence construction subsequently began.[1] However, estimates are imprecise and U.S. household survey data provides suggestive evidence that fence construction reduces return migration to Mexico from U.S. border counties.

Exploiting heterogeneity in the geographic concentration of fence construction and using migration data on Mexican households living away from the border, I investigate *why* the border fence deters migration. Interestingly, I find that, for residents of a particular municipality, fence construction in an adjacent municipality has a large, independent, and statistically significant deterrent effect. I also find that fence construction reduces migration rates for the subset of prospective migrants from non-border states with historically low access to smugglers (who are typically less knowledgeable about alternative crossing routes). Taken together, these findings suggest that the impact of fence construction on migration can potentially be explained by prospective migrants having knowledge of only a small set of low-cost crossing locations, in which case reducing access to these locations would subsequently result in a significant increase in the expected cost of migration. I use evidence on heterogeneous fence impacts based on the earnings history of prospective migrants to show that a left-censoring of the migration cost distribution, rather than a mean shift, best fits the data. This finding reinforces the thesis that a model in which border crossing costs are simply a

[1] In a model in which migrants have full information regarding the impact of fence construction on the cost of migration, we should expect a negative impact of construction on return migration that at least partly cancels out the observed fall in out-migration (assuming a sufficient share of those who've already migrated to the United States remain undocumented).

function of distance to the nearest unfenced crossing location cannot explain the empirical evidence, and instead provides additional support for the hypothesis that fence construction significantly deters migration precisely because the migration costs faced by prospective migrants are sensitive to the particular set of available crossing locations. Finally, I use the high elasticity of migration with respect to fence construction ($\beta = -0.39$) along with observed changes in the risk of death associated with crossing to argue that the crossing cost increases induced by fence construction are significant.

Estimating the elasticity of migration with respect to associated costs builds on an extensive literature, dating back to Sjaastad [1962], that documents the forces driving migration and that has placed great importance on understanding so-called Push-Pull factors [Borjas, 1989]. While there is a large body of theoretical work related to heterogeneous and changing migration costs (e.g., Rosenzweig, 2007), most empirical studies have relied on cross-sectional correlations between network characteristics (that proxy for migration costs) and migration rates.[2] Previous work that relies on time-varying migration costs is limited and data issues have generally prevented the accurate estimation of migration impacts. For instance, Hanson and Spilimbergo [1999] finds that potentially exogenous increases in hours spent by U.S. Border Patrol agents patrolling the border increase apprehensions. However, a rise in apprehensions may mask a fall in migration rates given that a single individual may be apprehended multiple times during the same migratory episode, and given that the likelihood of an individual being apprehended is presumably increasing in enforcement. Consequently, the authors cannot interpret their findings in relation to migration flows. Angelucci [2012] uses household survey data to identify a negative impact of agent "linewatch" hours on both migration from Mexico to the United States and return migration to Mexico, but relies on a sample of high-migration communities.[3] [4]

Borjas [1987] provides the theoretical foundation for the cross-sectional migration selection literature in his application of the classic Roy model [1951] to the expected earnings maximization problem faced by prospective migrants. A number of papers have subsequently tested for migration selection using data on Mexico-U.S. migration patterns.[5] However, the only paper to address dynamic migration selection is Angelucci [2012], in which the author compares migration selection in the 1970s to selection in the 1990s when border control was more intensive.

My contributions to these existing literatures are two-fold. First, I identify a new source of exogenous variation in migration costs that provides municipality-level variation and so allows me to examine cross-municipality spillovers. Second, I make novel use of the Mexican ENE/ENOE, a quarterly household survey

2 See, for instance, Orrenius and Zavodny [1993], which finds that Mexicans are more likely to migrate to the United States if they have family members who have done so.

3 Angelucci [2012] does employ a choice-based sampling method in response to sample selection concerns.

4 Relatedly, Gathmann [2008] employs a similar identification strategy and finds that increased linewatch hours are associated with only a moderate rise in smuggler prices, no change in demand for smugglers, and a shift in migration towards more remote areas.

5 Findings range from neutral-negative selection (e.g., Moraga, 2011, Martinez and Woodruff, 2007) to positive selection (e.g Chiquiar and Hanson, 2005) with differences driven primarily by the datasets used for estimation.

conducted throughout the country. This rich dataset mitigates sample selection concerns, provides an accurate measure of migration to and from the United States by members of sampled households, and allows me to study migration selection dynamics by exploiting the rotating panel structure of the data.

The remainder of the paper is presented as follows. The history of border enforcement along the U.S.-Mexico border is documented in Section 2.2. Section 2.3 discusses the data used for my analysis. Section 2.4 estimates the impact of fence construction on migration flows between Mexico and the United States and examines the spatial structure of estimated impacts. Section 2.5 presents a simple model of migration behavior and uses dynamic migration selection estimates to distinguish between potential mechanisms linking fence construction to changing migration costs. Section 2.6 concludes.

2.2 A History of U.S.-Mexico Border Control Policy

The United States-Mexico border is the most frequently crossed international border in the world and has a total length of 1,920 miles (3,090 km). There are four U.S. states adjacent to the border (California, Arizona, New Mexico, and Texas) and six Mexican states (Baja California, Sonora, Chihuahua, Coahuila, Nuevo León, and Tamaulipas), including 23 adjacent U.S. counties and 37 adjacent Mexican municipalities.

Efforts to restrict migration from Mexico to the United States are a fairly recent phenomenon. Between 1942 and 1964, over four million Mexicans immigrated to the United States as contract laborers through the Bracero Program [PBS, 1999]. Although the Immigration and Nationality Act of 1965 officially restricted the number of Mexicans that could migrate to the United States, it was not until 1986 that the federal government put in place legislation aimed at curbing undocumented immigration. The 1986 Immigration Reform and Control Act (IRCA) made knowingly hiring undocumented immigrants illegal for the first time and authorized a significant expansion of Border Patrol activities [Andreas, 2001]. This 1986 legislation was followed by Border Patrol campaigns in the early 1990s in the San Diego and El Paso Sectors that increased the number of agents on the border and also authorized construction of border fencing for the first time [Hanson and Spilimbergo, 1999; Nuñez-Neto and Viña, 2006]. In 1996, the Illegal Immigration Reform and Immigrant Responsibility Act (IIRIRA) authorized the Attorney General to construct fencing along the border and mandated additional construction in the San Diego Sector, but only nine of the mandated fourteen miles of secondary fence were completed and construction outside of the San Diego Sector remained limited [Nuñez-Neto, B. and S. Viña, 2006].[6]

The REAL ID Act of 2005 paved the way for subsequent fence construction by authorizing the Secretary of DHS to waive any laws that impeded construction of security barriers along the border.[7] Then, on October 26, 2006, President George W. Bush signed into law the Secure Fence Act, which called for the construction

[6] A 2009 Government Accountability Office document reports that a total of 78 miles of pedestrian fence had been constructed by late 2005 [Stana et al., 2009a].

[7] This included, for instance, environmental legislation such as the National Environmental Policy Act and the Endangered Species Act.

of 698 miles of double-layered (pedestrian) fence along designated segments of the United States-Mexico border.[8] The Act initially specified that a 370-mile segment of fence from Calexico, California to Douglas, Arizona be completed by May 30, 2008 (along with a 30-mile segment adjacent to Laredo, Texas). However, the Consolidated Appropriations Act, enacted on December 26, 2007, significantly increased the Secretary's discretion in determining where to install fencing. This legislation required only that a minimum of 700 miles of fencing be constructed where it would be "most practical and effective," and that the Secretary identify either 370 miles or "other mileage" which would be completed by December 31, 2008.

It seems likely that relaxed restrictions on the location of fencing can be at least partly explained by concerns that DHS would not meet the timeline set out in the original Secure Fence Act legislation. By September 2007, the GAO reported that only 71 miles of pedestrian fence, along with 2 miles of vehicle fence, had been constructed since late 2005. However, significant fence construction did occur over the following year, and there were a total of 140 pedestrian fence miles and 75 vehicle fence miles in place by October 2008 [Stana et al., 2009a]. As of April 2010, 262 miles of pedestrian fence and 227 miles of vehicle fence had been constructed [Stana, 2010].[9] Declassified government documents indicate that final fence construction locations had been chosen based as much on expediency and cost as on security concerns. In one conversation, for instance, Border Patrol agents reported that "They [Army engineers] were looking at placing fencing in areas that would not be our operational priority" and that Army engineers had stated "We need to throw up fence in the areas that are most advantageous to meeting the timeline" [DHS, 2007].

2.3 Data

2.3.1 Fence Construction

This project required that I collect my own data on the timing of fence construction along the U.S.-Mexico border, given that such data is not made publicly-available by the government or compiled by any private organization. To determine the timing of fence construction, I have relied on the following two-step procedure. First, I identified a set of potential fence locations by reviewing a series of GAO reports related to border fence construction, by consulting documents produced by CBP that identify prospective fence segment locations, and by speaking with representatives from the Sierra Club (an environmental organization charting fence

[8] Technically, the Secure Fence Act called for "at least 2 layers of reinforced fencing" [Secure Fence Act of 2006]. While this language has subsequently been revised, and there has been some debate regarding what exactly is required based on the legislation, DHS documentation makes clear that the original interpretation was that all 698 miles would be pedestrian fence [CBP, 2007].

[9] While 78 miles of fencing were already in place by late 2005, some subset of these miles were subsequently replaced and counted as Secure Fence Act miles [Stana et al., 2009a].

construction in order to understand environmental impacts) to identify likely fence locations.[10] [11] After identifying a set of potential fence locations, I searched for news articles, legal documents, local government reports, and published contracts with the responsible construction firms in order to identify (1) whether a fence was indeed constructed, (2) the date on which construction began, and (3) the length of fence constructed.[12] [13] As an example, Appendix Figure 2.1 illustrates the progress of fence construction in the Mexican border state of Sonora.

2.3.2 Mexican Household Data

Data on Mexican migration rates and labor market outcomes covers the period from the first quarter of the year 2000 through the second quarter of 2011 and comes from the Mexican Encuesta Nacional de Ocupación y Empleo (ENOE) and its precursor, the Mexican Encuesta Nacional de Empleo (ENE). The ENE/ENOE is a quarterly survey that is representative at the state level. The survey is a rotating panel that includes selected households for five consecutive quarters before they are rotated out of the panel. In the first quarter that a household enters the panel, the surveyor records all current household members. In subsequent quarters, the surveyor records whether any household members left or returned to the household, and the Mexican state or country for which they left or from which they returned (i.e., the United States).[14] The respondent sociodemographic characteristics that are used to restrict the sample in a subset of specifications and to look for heterogeneous impacts of fence construction come from the sociodemographic section of the ENE/ENOE survey. This section records gender, age, literacy, and education for all household members. In addition, all household members aged 12 and older are asked a series of questions in the occupation and employment section of the survey about labor force participation and monthly earnings. While the ENE/ENOE has

[10] CBP documents were made available thanks to a successful lawsuit brought by Citizens for Responsibility and Ethics in Washington to challenge the Department of Homeland Security's failure to release border fence documentation in response to a Freedom of Information Act (FOIA) request.

[11] Although one might suspect that fence location data would be publicly-available, DHS and CBP restrict the availability of relevant data. Restricted access is justified on the basis of national security concerns- a 2008 newspaper article quotes a CBP official, who states that "all data regarding the placement of the fence is classified because you don't want to tell the very people you're trying to keep from coming across the methodology used to deter them" [del Bosque, 2008].

[12] I have not made use of fence length data given concerns about the accuracy of the data. In particular, while I can identify the date on which construction began and generate a noisy measure of final fence segment length, it is very difficult to impute the quarter-wise length of fence constructed.

[13] Although it is ex-ante unclear whether the start of fence construction or the date of completion is the point at which we expect to observe impacts on migration flows, a review of CBP documentation suggests that the start of fence construction is associated with a rise in the probability that those crossing the border illegally will be detected [DHS, 2010].

[14] Although household and respondent identifiers are not provided, according to the survey manual, they can be constructed based on a set of recorded variables: state, municipality, metropolitan status, survey area, building number, household number, household roster number, and survey number. In practice, a small share of households are not uniquely identified based on this procedure, in which case I use the year-quarter of survey to distinguish households. Based on this two-step procedure, approximately 1% of households do not have an associated identifier. Results are robust to including or excluding this subsample.

been used frequently in the past to analyze Mexican labor markets, its application to migration research has been limited despite the unique advantages of its scale, breadth, and frequency.[15] One limitation of the ENE/ENOE data is that, as is the case for other Mexico-based migration surveys, it does not capture migration episodes undertaken by entire households.[16]

In my analysis, I also use data from the Mexican Migration Project (MMP), which is co-directed by researchers at the University of Guadalajara and Princeton University and was created in 1982. The MMP initially surveyed households from communities in Western Mexico, a region characterized by high migration rates, although the sample of communities has since expanded geographically. Surveys are conducted in both Mexico and the United States. The MMP provides detailed migration histories for each respondent that include the dates of each undocumented crossing undertaken, the locations of the crossings, and whether coyotes (smugglers) were used. While the lack of representativeness of the data is an issue when responses are used to measure impacts of particular policies, I use the MMP data only to construct historical migration networks for residents of non-border states (I return to the ENE/ENOE data to estimate actual migration flows).

2.3.3 U.S. Household Data

To measure migratory impacts on the U.S. side of the border, I make use of the Current Population Survey (CPS) and American Community Survey (ACS). The CPS is conducted monthly and records demographic characteristics and labor force status of the surveyed population. Each March, responses to the Annual Social and Economic Supplement are also collected from approximately 100,000 households, and this supplement to the CPS includes questions on work history and earnings. The ACS is conducted on an ongoing basis, with approximately 2,000,000 households surveyed annually, and with responses aggregated by year. The ACS includes questions on demographics, employment, and earnings similar to those from the CPS. For the purposes of my analysis, I pool CPS and ACS data to increase statistical power, and include respondents from all counties adjacent to the U.S.-Mexico border. In order to pool data, I match Metropolitan Statistical Areas from the CPS data to the counties in which they are located (the ACS data is made available at the county level).

2.4 Fence Construction and Migration Flows

The responsiveness of migration to changes in cost is the most important determinant of the efficacy of border control policies designed to make migration more difficult. To determine whether border fence construction has a deterrent effect, I first look for an impact on residents of municipalities in Mexico that border the

[15] Research that has used ENE/ENOE data to study migration includes Rendall et al. [2010], Cadena and Kovak [2013], and a number of press releases produced by the Mexican Instituto Nacional de Estadística y Geografía (INEGI).

[16] Based on alternative data sources, it appears that undocumented migration by entire Mexican households to the United States is quite uncommon.

United States. Throughout the analysis, I include only those Mexican border municipalities which the Secure Fence Act of 2006 designated to be fenced. In practice, this includes 27 of the 37 Mexican border municipalities, but the 10 excluded municipalities are significantly more rural and less densely-populated than their counterparts in which fence construction was proposed.[17]

2.4.1 Migration to the United States

My identification strategy will rely on the assumption that, absent the construction of border fence, residents of fenced and unfenced municipalities would have exhibited parallel trends in migratory behavior. To provide evidence that differential out-migration from municipalities that experience fence construction is not driven by pre-trends (e.g., there is not some unobserved third factor correlated with both fence construction and migration rate changes that would lead me to calculate biased estimates of the impact of fence construction on migration), I estimate the following equation:

$$\begin{aligned} P(Y_{mqi} = 1|z_{mqi}) &= \wedge(\alpha + \sum_{t=-5}^{4} \beta_t * fence_{mqit} + \gamma_m + \lambda_q) \\ &= \frac{exp(\alpha + \sum_{t=-5}^{4} \beta_t * fence_{mqit} + \gamma_m + \lambda_q)}{1 + exp(\alpha + \sum_{t=-5}^{4} \beta_t * fence_{mqit} + \gamma_m + \lambda_q)} \end{aligned} \tag{2.1}$$

where Y_{mqi} is an indicator variable equal to '1' when individual **i** in municipality **m** migrates to the United States in year-quarter **q**, $fence_{mqit}$ is an indicator variable defined by whether border fence construction has begun in municipality **m** by **t** years after year-quarter **q** (or $|\mathbf{t}|$ years before for negative-valued **t**), and γ_m and λ_q represent municipality and year-quarter fixed effects, respectively. Standard errors are clustered at the municipality level to deal with the concern that the error covariance matrix may exhibit correlations varying in magnitude across municipality residents and/or across time periods within a particular municipality. The base specification employs a logit model which is appropriate due to the binary outcome measure and the fact that we expect the impact of fence construction to be proportional to pre-fence migration flows.[18] Figure 2.1 plots estimated regression coefficients. The omitted year is $\mathbf{t} = -1$ (relative to the start of construction, defined by $\mathbf{t} = 0$). In Panel A, the dependent variable is equal to '1' for all instances of migration to the United States. In Panel B, as a robustness check, the dependent variable is equal to '1' only for migration to the United States for work-related purposes. None of the years before the start of fence construction exhibit coefficients significantly different from zero in either panel, and there is no visual evidence of pre-trends in migration patterns. This suggests that differential pre-trends in municipalities experiencing fence construction are not a significant identification concern. As additional robustness checks, Figure 2.2 plots the

[17] The ENE/ENOE has data for 22 of the 27 municipalities for the relevant years. The five municipalities without data available have populations significantly lower than those 22 municipalities for which data is available.

[18] If the impact was a level effect as implied by an OLS specification, we should expect municipalities with high pre-fence migration rates to experience the same absolute change in migration rates as a result of fence construction as those with low pre-fence rates. This seems unlikely in practice.

regression coefficients from a linear probability specification, and Figure 2.3 replaces migration with income as the dependent variable and plots coefficients from years before and after fence construction based on a Poisson QMLE specification.[19]

Evidence from Figures 2.1-2.3 that the parallel trends assumption is likely satisfied suggests that the causal impact of fence construction on migration from Mexican border municipalities to the United States can be estimated based on the following specification:

$$\begin{aligned} P(Y_{mqi} = 1 | z_{mqi}) &= \Lambda(z_{mqi}) \\ &= \Lambda(\alpha + \beta * fence_{mqi} + \gamma_m + \lambda_q) \end{aligned} \tag{2.2}$$

where $fence_{mqi}$ is an indicator variable defined by whether border fence construction has begun in municipality **m** in year-quarter **q**, and the remaining terms are as defined in Equation 2.1.

Table 2.1 estimates Equation 2.2 to calculate the impact of fence construction on migration from Mexican border communities to the United States. Column (1) of Table 2.1 includes all observations from Mexican border municipalities which the Secure Fence Act of 2006 designated to be fenced.[20] The estimated coefficient on $fence_{mqi}$ suggests that fence construction has a negative and statistically significant impact on migration (significant at the 1% confidence level). In terms of magnitude, given binary dependent and explanatory variables, e^{β} reflects the estimated percent change in the dependent variable resulting from the explanatory variable moving from **0** to **1**. Hence, for this specification (with $\beta = -0.486$), the fence is estimated to reduce migration by 38.5% relative to a baseline level of 2.5 migration episodes per 1,000 respondents. Column (2) and all subsequent even-numbered columns in Table 2.1 add state*year-quarter fixed effects to ensure that estimated impacts are not being driven by differential trends in states in which fence construction is concentrated. For all specifications, the inclusion of state*year-quarter fixed effects increases statistical significance.

In Columns (3) and (4), I restrict attention to a panel defined by statewise date of initial construction such that I include, for each state, all quarters up to six years before and up to three years after fence construction began in that state. In Columns (5) and (6), I estimate fence impacts in a balanced panel that

[19] Poisson QMLE is appropriate given that I am interested in estimating the impact of fence construction on log income and have a significant number of '0'-valued observations. Estimates from this specification can be interpreted as percent changes.

[20] Since the data is a rotating panel that includes a maximum of five observations per household, and since migration is defined only for individuals who were present in the previous year-quarter, I exclude the first observation for each household member. In addition, for a household member who has migrated, I code all subsequent observations as missing. Lastly, given that a reason for household member exit is not provided in 27% of cases when the question is asked, I include all migration to the United States in the dependent variable. This is also preferred because the fact that the reason for household member exit is reported by other household members *after* the migrant has already left the household raises concerns regarding measurement error and response endogeneity. As a robustness check, I re-run my primary specification excluding non-employment-related migration for the years in which this data is available, and find that the magnitude of coefficients increases slightly but that estimates also become noisier.

includes the 18 (of 22) border municipalities from the Secure Fence Act sample which have data available for all years from 2003 to 2011. Finally, in Columns (7) and (8) I include a control for the number of adjacent municipalities in which fence construction has begun.[21] Estimates from Columns (3) to (8) are large and statistically significant, indicating that the fence is effective at deterring migration from Mexican border communities to the United States.

Appendix Tables 2.1 and 2.2 present robustness checks. In Appendix Table 2.1, I estimate the impact of the fence based on a linear probability model. I find coefficients of the order of .001, suggesting that fence construction leads to a 0.1 percentage point fall in migration rates (this is consistent with an approximately 40% fall in migration given the baseline migration rate in the sample is 0.25%). All specifications including state*year-quarter fixed effects are significant at the 5% confidence level, while the odd-numbered specifications are now significant at 10%.[22] In Appendix Table 2.2, I estimate the weighted least squares version of Equation 2.2 at the municipality*year-quarter level. The dependent variable is simply the mean value of Y_{mqi} in municipality **m** in year-quarter **q**, and the observations are weighted by the number of survey responses within each municipality*year-quarter cell. Appendix Table 2.2 estimates are consistent with the individual*year-quarter level OLS estimates from Appendix Table 2.1.

2.4.2 Return Migration to Mexico

If migrants are fully informed regarding changes in the cost of undocumented migration, we should expect to observe a fall in return migration to Mexico among those who moved to the United States from a border municipality that was subsequently fenced. These migrants should rationally respond to increased costs of future migration to the United States by, for instance, forgoing planned trips back to Mexico in between work episodes in the United States. Given the structure of the data, however, the parallel of the specification used to estimate the impact of the fence on out-migration is problematic since the potential set of migrants returning from the United States to surveyed households in Mexico is unknown.[23]

One strategy for dealing with this issue related to return migration estimation would be to consider only observations for individuals currently in the United States who migrated from Mexico in a previous survey quarter. However, there are too few such observations to allow for estimation. Consequently, I run

[21] To the extent that individuals are also deterred by construction in municipalities other than their own, Equation 2.2 will under-estimate the true magnitude of the impact of fence construction on migration. The nature of geographic spillovers is investigated more carefully in Section 2.4.3.

[22] Given the relatively low number of clusters in my sample, the p-values associated with OLS estimates are constructed using the wild bootstrapping procedure outlined in Cameron et al. [2008].

[23] In particular, suppose that individual **i** is listed as living in her household in year-quarter **q**. Then, this individual cannot be listed as returning to Mexico in year-quarter **q+1**. For the out-migration case, the parallel is that an individual who is already in the United States in year-quarter **q** cannot migrate to the U.S. in year-quarter **q+1**. However, in the out-migration case, this individual would typically not be included in the ENE/ENOE dataset in either year-quarter **q** or year-quarter **q+1**, so the data structure does not lead to biased estimates of migration to the United States when specifications are estimated at the individual*year-quarter level.

regressions considering first the municipality*year-quarter and then the household*year-quarter as the unit of analysis. In practice, these regressions estimate the effect of fence construction on the number of return migrants per survey respondent in a municipality in a given quarter, or on the probability that a household in a given quarter has a member return. While the magnitudes estimated using the municipality*year-quarter and then the household*year-quarter specifications will not be as readily interpretable as those derived from the out-migration specifications, they will address the first-order question of whether impacts on return migration are at all observable.[24]

Table 2.2 presents the relevant estimates. A negative coefficient would confirm the hypothesis that fence construction reduces return migration by those from newly-fenced municipalities who are currently living in the United States. To the extent that reducing the size of the undocumented population (rather than reducing cross-border flows) is the goal of border control measures such as fence construction, estimates of the extent to which fence construction reduces migration to the United States must be supplemented by estimates of how the fence affects return migration to Mexico.

Columns (1) and (2) of Table 2.2 use the municipality*year-quarter as the unit of analysis, with observations weighted by the number of survey responses within each municipality*year-quarter cell. Estimates are statistically insignificant, and the sign of the coefficient changes from positive to negative when state*year-quarter fixed effects are added in Column (2). Columns (3) and (4) look for differential in-migration when the unit of observation is at the household*year-quarter level, and the coefficients reflect the impact of fence construction on the likelihood that a household member returns from the United States in a given year-quarter. The -.00004 coefficient in Column (3) suggests the impact of fence construction on return migration is negligible, while Column (4) results imply that fence construction reduces migration to Mexico by 8%. Finally, Columns (5) and (6) use responses at the individual*year-quarter level to provide a comparison to the Table 2.1 analysis (with the caveat that this regression underestimates migration flows from the United States to Mexico given the data structure). The Column (5) coefficient is positive while the Column (6) coefficient is negative. Unfortunately, no clear pattern emerges across these specifications: none of the six coefficients is statistically significant, and even the sign of the effect is ambiguous (three estimated coefficients are positive and three are negative). In addition, given the large standard errors estimated, it is difficult to discern whether fence construction has any economically significant impact on return migration patterns.

As an alternative approach aimed at producing more consistent estimates, I make use of U.S. Current Population Survey (CPS) and American Community Survey (ACS) data to look for changes in return migration patterns. To do so, I first pool the two datasets in order to estimate the impact of fence construction on the share of the population that is potentially undocumented (defined as those aged 16-65 who were born in Mexico and are not college graduates). The logit specification employed is as follows:

[24] I do also present the in-migration version of Equation 2.2 for comparison's sake.

$$
\begin{aligned}
P(Y_{cyi} = 1|z_{cyi}) &= \wedge(z_{cyi}) \\
&= \wedge(\alpha + \beta * fence_{cyi} + \gamma_c + \lambda_y) \\
&= \frac{exp(\alpha + \beta * fence_{cyi} + \gamma_c + \lambda_y)}{1 + exp(\alpha + \beta * fence_{cyi} + \gamma_c + \lambda_y)}
\end{aligned}
\tag{2.3}
$$

where Y_{cyi} is an indicator variable equal to '1' if individual **i** in county **c** in year **y** is characterized as potentially undocumented. $fence_{cyi}$ is an indicator variable defined by whether border fence construction has begun in county **c** in year **y**, and γ_c and λ_y represent county and year fixed effects, respectively. Standard errors are clustered at the county level and the sample includes all respondents aged 16-65. Columns (1)-(2) of Table 2.3 suggest that the share of potentially undocumented migrants in U.S. border counties is unaffected by fence construction.

A second specification tests the impact of fence construction on the time that potentially undocumented migrants have resided in the U.S. The estimation framework is similar to the one above, although I now include only the CPS data since the ACS does not record date of arrival in the United States. In this specification, the dependent variable is now defined by whether the respondent is a recent migrant (arrived in the U.S. in the past three years) and the sample of respondents includes only those characterized as potentially undocumented. Combining estimates of the change in the share of potentially undocumented migrants in the population with estimates of the change in the distribution of migrants' time in the U.S. since arrival provides an alternative framework for identifying changes in migration to and from U.S. border counties. Logit and OLS estimates are presented in Columns (3)-(4) of Table 2.3 and provide evidence that fence construction does deter return migration- construction is associated with a 53% reduction in the share of potentially undocumented migrants who arrived in the past three years. Since the population share of potentially undocumented migrants is unchanged, I interpret this estimate as providing evidence that fence construction increases the likelihood that migrants forgo return trips to Mexico. Figure 2.4 plots the full distributions of the number of years that potentially undocumented respondents have lived in the U.S. for both fenced and unfenced border counties.

Making use of data on employment patterns, I find evidence that fence construction reduces employment among border county residents aged 16-65 (Columns (5)-(6)) and that this effect appears to be driven by reduced employment among the potentially undocumented population (Columns (7)-(8)). Using Quarterly Workforce Indicators published by the U.S. Census Bureau, I employ county-year level specifications in Column (9)-(10) and observe significant reductions in both separations and hires among Hispanics living in U.S. border counties in response to fence construction. This finding is consistent with the hypothesis that fence construction induces the population of potentially undocumented respondents to forgo return trips to Mexico and instead focus on maximizing earnings in the U.S. Those who are employed choose to remain on the job, and this behavioral change explains the observed fall in separations (and in the hiring of replacement workers). At the same time, workers who would have returned to Mexico between work episodes avoid doing so, and thereby reduce the employment rate among potentially undocumented residents.[25]

2.4.3 Geographic Spillovers

To investigate the mechanisms driving the large estimated impact of fence construction on migration from Mexico to the U.S., I test whether there are geographic spillovers in fence impacts (e.g., whether fence construction in one border municipality impacts migration from adjacent municipalities). The nature of cross-municipality spillovers can shed light on the relative importance of fence construction in a prospective migrant's home municipality, and also reflects how (if at all) the costs faced by prospective migrants vary with the set of nearby municipalities that are fenced. The specification employed here is as follows:

$$\begin{aligned}
P(Y_{mqi} = 1|z_{mqi}) &= \wedge(z_{mqi}) \\
&= \wedge(\alpha + \beta_1 * fence_{mqi} + \beta_2 * fence_{nqi} \\
&\quad + \beta_3 * fence_{oqi} + \beta_4 * fence_{mqi} * fence_{nqi} + \gamma_m + \lambda_q)
\end{aligned} \tag{2.4}$$

where $fence_{nqi}$ is a count variable that measures the number of municipalities adjacent to municipality **m** in which fence construction has begun by year-quarter **q** and $fence_{oqi}$ is a count variable that measures the number of municipalities two away from municipality **m** in which fence construction has begun by year-quarter **q**. The remaining variables are as defined in Equation 2.1. The $fence_{mqi} * fence_{nqi}$ interaction term can be interpreted as a test of the simplest model of the border crossing cost structure whereby migrants' crossing costs are determined by the distance to the nearest unfenced location/municipality. If this model is accurate, we anticipate a negative interaction term (and a coefficient on $fence_{nqi}$ of '0'). In contrast, if cross-municipality spillovers are driven by migrants having heterogeneous knowledge of low-cost border crossing locations (such that some migrants from a given municipality know where to cross in that municipality while others know where to cross in a nearby municipality), then we anticipate an interaction term equal to '0' and a negative coefficient on $fence_{nqi}$. I include $fence_{oqi}$ in the regression in order to examine the rate at which any observed spillovers diminish with distance.

Estimates of geographic spillovers in fence construction impacts are presented in Table 2.4. In Column (1), I estimate Equation 2.4 and find that the coefficient on the number of adjacent municipalities with fence construction is negative, large, and statistically significant while the coefficient on the interaction between construction in own and in adjacent municipalities is positive, but not statistically significant at standard confidence levels. The coefficient on the number of fenced municipalities that are neighbors of adjacent municipalities (i.e., two municipalities away) is also statistically insignificant. Column (2) presents corresponding OLS estimates. The estimates in Columns (1)-(2) represent a strong rejection of the shortest distance hypothesis, whereby migrants simply cross in the nearest unfenced municipality. Instead, the

[25] Implicitly, some search frictions must prevent these unemployed residents who are deterred from returning to Mexico from immediately being allocated those jobs that were designated for new arrivals who were deterred from crossing to the United States as a result of fence construction.

evidence is consistent with a heterogeneous knowledge model in which migrants are additively affected by construction in own and in adjacent municipalities because prospective migrants have private, and heterogeneous, information about potential crossing locations. The fact that the coefficient on adjacent municipality construction is approximately two-thirds of the coefficient on own municipality construction is in line with a model in which prospective migrants are most likely to have knowledge of crossing locations in their home municipality.

One identification concern that arises in this analysis of cross-municipality spillovers is that the timing of construction across adjacent municipalities may be correlated, in which case the estimated coefficients cannot be interpreted as proposed above. To address this concern, I test whether fence construction in a given municipality predicts (positively or negatively) construction in adjacent municipalities designated for construction in the Secure Fence Act. To do so, I estimate:

$$fence_{nq} = \alpha + \beta * fence_{mq} + \gamma_m + \lambda_q \tag{2.5}$$

where $fence_{mq}$ is an indicator variable defined at the municipality*year-quarter level and equal to '1' if border fence construction has begun in municipality **m** by year-quarter **q**. $fence_{nq}$ is defined as the number of municipalities adjacent to municipality **m** in which fence construction has begun by year-quarter **q**.

Columns (3) and (4) show corresponding OLS estimates. The point estimates are negative, small, and not statistically significant. This suggests that construction timing across adjacent municipalities is plausibly independent, which in turn implies that my estimated impacts of construction on residents of border municipalities are not conflated by the associated change in the probability that construction also takes place in other nearby municipalities. In Column (5), I run a municipality-level specification testing whether construction in one municipality predicts whether construction *ever* occurs in adjacent municipalities. Again, the estimated coefficient is statistically insignificant.

2.4.4 Non-Border Effects

There are a number of reasons why observed impacts of the border fence on migration might be specific to border municipality residents. First, these individuals have the shortest distance to travel to reach the border, and so the impact of the fence on total cost of migration may be proportionally higher for them than for those living farther from the border. Second, border residents may be more affected because information about fence construction reaches them before it reaches migrants traveling from more distant states (who may arrive to the border region before learning of the fence). Third, municipalities in the border region are strongly linked economically to their neighbors on the United States side of the border, and these relationships may drive heterogeneous effects of fence construction based on proximity to the border. Given this broad range of possible explanations for why non-border state migrants might not be deterred by fence construction, it is difficult to interpret a null result. However, evidence that fence construction does impede prospective migrants living outside of the border region is not likely consistent with a story in which migrants cross in the nearest unfenced location and so fence construction increases crossing costs by increasing the distance they must travel to find an unobstructed crossing location (since the change in costs for non-border residents relative to total crossing costs will be small).

To determine whether border fence constructions impacts migration for non-border state residents based on their historical migration networks, I employ the following specification:

$$\begin{aligned} P(Y_{sqi} = 1 | z_{sqi}) &= \wedge(z_{sqi}) \\ &= \wedge(\alpha + \beta * AvgFence_{sqi} + \gamma_s + \lambda_y) \end{aligned} \tag{2.6}$$

$$AvgFence_{sqi} = \sum_{m=1}^{13} P_{sm} * fence_{mqi} \tag{2.7}$$

where γ_s are state fixed effects, λ_y are year fixed effects, P_{sm} is the share of pre-2000 migrants born in state **s** who crossed in municipality **m**, and the remaining variables are as defined in Equation 2.2. The municipalities **m** are the 13 border municipalities with at least 50 crossings recorded (out of a total of 9,098 crossings in the data). Mexican Migration Project survey data, described in Section 2.3, is used to construct these pre-2000 migration patterns. By relying on historical migration patterns to predict likely crossing locations, this approach ensures that the set of border municipalities assigned to non-border residents as potential crossing locations is not endogenously affected by fence construction patterns.[26]

Column (1) of Table 2.5 estimates Equation 2.6 and the point estimate is consistent with fence construction reducing out-migration from non-border states by 17%. While this estimated impact is relatively large, the coefficient is not statistically significant at standard confidence levels (although it is significant at the 16% confidence level).[27] Column (2) presents the corresponding OLS specification.

Columns (3)-(4) of Table 2.5 makes use of Mexican Migration Project data on coyote (smuggler) use by re-running Equation 2.6 with a sample restricted to residents of those states with below-median coyote use as measured by the state-level share of pre-2000 crossings for which a coyote was used (the median coyote use rate is 72%). I find that the effect of the fence is now significant (and quite large). This evidence of impact heterogeneity is consistent with the fact that coyotes' livelihood requires that they have constant access to *some* crossing location that can be utilized, making it reasonable to assume that they are better-informed about available low-cost crossing locations than migrants themselves. Differential knowledge of crossing options implies that it is those individuals crossing without the help of a coyote who are likely to have knowledge of only a limited range of low-cost crossing locations and so are consequently most likely

[26] Note that results are not robust to including year-quarter fixed effects; the coefficients increase in magnitude, but the standard errors grow proportionately larger.

[27] One possible explanation for the large standard errors estimated is that the weighted average fence measure is computed based on the subset of municipalities surveyed in the MMP, and then this index value is assigned to all residents of the corresponding states. Since historical border crossing patterns are likely to vary within states, this procedure may be generating substantial statistical noise.

to be deterred by the closing off of particular border crossing paths. In sum, the finding that the deterrent impact of the fence is a function of historical access to smugglers buttresses the hypothesis that knowledge of (unfenced) low-cost crossing locations is a key determinant of border crossing costs.

2.4.5 Heterogeneous Effects

While the overall elasticity of migration with respect to fence construction estimated in Table 2.1 can be used to calculate the number of migrants deterred by the fence, the sociodemographic characteristics of those deterred have potentially important implications for both sending and receiving communities. Individual income and educational attainment are the observable characteristics usually examined to determine the nature of migrant selection (positive selection corresponds to disproportionate migration by those with high income and high educational attainment relative to the general population). To estimate the baseline type of selection observed in my context and to investigate how selection changes in response to fence construction, I employ the base specification outlined in Equation 2.2, but add a series of level terms for sociodemographic characteristics (gender, aged between 38 and 60, years of schooling, and baseline log income) as well as terms representing the interactions between these characteristics and fence construction. The level term coefficients can be roughly interpreted as characterizing the baseline nature of migrant selection.[28] The interaction terms reflect how migrant selection changes in response to an exogenous increase in the cost of migration (namely, construction of the fence).[29]

Table 2.6 presents estimates of heterogeneous impacts based on respondent characteristics. Columns (1) and (2) reveal that those with higher incomes are less likely to migrate at baseline, while educational attainment does not appear to significantly predict migratory behavior (the coefficient on years of schooling is positive but not statistically significant at conventional confidence levels). Investigating *changes* in migrant selection in response to fence construction, I find that individuals with more years of schooling are less deterred from migrating by fence construction, and this finding is quite robust (statistically significant at the 1% confidence level). In terms of the magnitude of estimates, each additional year of schooling reduces the deterrent effect of the fence by approximately 12%. In contrast, higher income conditional on education magnifies the deterrent effect of the fence: each one-point increase in log income increases the impact of the

28 Technically, these coefficients are also affected by geographic spillovers in fence impacts, but additional analysis reveals that the results are unchanged when the standard selection analysis is run using only the data from the pre-fence years (2000-2005).

29 One data issue that arises is that demographic characteristics are not recorded when an individual has exited the household. Consequently, I rely on the first recorded response for each household member when using these measures to restrict the sample or look for heterogeneous impacts. The value of most, but not all, of the sociodemographic measures used should not vary in the absence of measurement error. In practice, however, there is a fair amount of measurement error that necessitates choosing an approach for determining what variable values to assign to respondents. As a robustness check, I verify that results are robust to using the maximum value for demographic characteristics rather than the values recorded at the time that respondents are first surveyed.

fence by 24%.[30]

2.5 A Test of Changing Migration Costs

The nature of impact heterogeneity based on the attributes of potential migrants can be used to characterize the induced change in the distribution of migration costs. Consider a simple model in which an individual migrates if:

$$W_{us} - W_m > \pi \tag{2.8}$$

where W_{us} is the wage the individual will receive in United States, W_m is the wage received in Mexico, and π is the cost of migration.

If we assume that $\pi \sim N(\mu, \sigma^2)$, then the probability that an individual migrates is given by the expression $P = \Phi(\frac{W_{us} - W_m - \mu}{\sigma})$. Based on this expression, it is clear that migration is rising in W_{us} and falling in W_m. Suppose that fence construction simply shifts up the mean cost of migration, μ (for instance, by increasing the distance that each migrant must travel). The change in the probability of migration as a function of μ is as follows:

$$dP/d\mu = \frac{-1}{\sigma}\phi(\frac{W_{us} - W_m - \mu}{\sigma}) < 0 \tag{2.9}$$

Thus, the model offers the basic prediction that rising mean migration costs will reduce migration. A more discriminating test of whether a mean shift in the distribution of crossing costs is consistent with observed changes in migration patterns involves examining the relative deterrent effect for those with higher residual income (i.e. higher income conditional on educational attainment). In computing comparative statics based on residual income, I rely on the assumption that the income of individual i in Mexico (W_m) is independent of her migration costs (π) and expected earnings in the U.S. (W_{us}), conditional on her educational attainment. This conditional independence assumption for the expected wage in the United States seems reasonable given that undocumented migrants are most likely to work in occupations, such as agricultural production, with limited wage variation among the undocumented workforce. Similarly, while we might expect educational attainment to affect an individual's social network and consequently impact migration costs, it seems unlikely that residual income (determined by unobserved productivity and/or rents associated with a particular job) is likely to predict migratory costs. To provide empirical support for this conditional independence assumption, I examine Mexican Migration Project data on the subset of respondents who report non-zero earnings in both Mexico and the United States and I find that there is a significant relationship between U.S. and Mexican

[30] When education and income interactions are estimated separately, coefficients are somewhat smaller in magnitude. All interaction coefficients remain statistically significant except for the coefficient on the education interaction in the OLS specification. In analysis not presented, I also find that females are significantly less likely to migrate than males, but that they are not differentially affected by fence construction. Similarly, those aged 38 to 60 are less likely to migrate than those aged 16 to 37, but fence construction has the same negative impact on the propensity to migrate of individuals in both age groups.

wages when educational attainment is excluded, but that relationship is no longer statistically significant at conventional confidence levels when a control for educational attainment is added.

Given conditional independence of migration costs and expected income in the U.S., the expression for the change in $dP/d\mu$ with respect to residual income W_m is as follows:

$$\frac{\partial^2}{\partial\mu\partial W_m}P = \frac{-(W_{us} - W_m - \mu)}{\sigma^3}\phi(\frac{W_{us} - W_m - \mu}{\sigma}) \tag{2.10}$$

Based on the data, I can infer that $W_{us} - W_m - \mu < 0$ since migration rates are significantly below 50% for all education-level cells. This, in turn, implies that $\frac{\partial^2}{\partial\mu\partial W_m}P > 0$. The percentage change in the deterrent effect as residual income rises can be calculated based on the following expression:

$$\frac{\partial}{\partial W_m}\frac{\partial P/\partial\mu}{P(W_m)} = \frac{-((W_{us} - W_m - \mu)\phi(W_{us} - W_m - \mu)\Phi(W_{us} - W_m - \mu))}{(\sigma^2)\Phi^2(W_{us} - W_m - \mu)} + \frac{\phi^2(W_{us} - W_m - \mu)}{(\sigma^2)\Phi^2(W_{us} - W_m - \mu)} \tag{2.11}$$

This expression is potentially negative, in contrast with the level change in the deterrent effect as residual income rises. Thus, a shift in mean migration costs, when migration costs are normally distributed, can potentially generate a percentage-wise reduction in migration that is larger for high residual earnings individuals, but cannot produce a larger absolute fall in migration propensity for this subpopulation.[31] Consequently, the signs of the expressions derived in Equations 2.10 and 2.11 can be used to determine whether observed changes in migration patterns can be rationalized by mean cost shifts.

Columns (1)-(2) of Table 2.6 provide empirical estimates of the relevant cross-partials, and imply that fence construction has a larger deterrent effect on individuals with high residual income in both percentage and absolute terms (estimates from the logit specification in Column (1) can be interpreted as percentage changes while Column (2) linear probability model estimates reflect absolute changes). This differential deterrent effect implies that a mean shift in the crossing cost distribution is not consistent with the observed changes in migration patterns in response to fence construction. In contrast, the observed change in crossing patterns can be rationalized by a disproportionate contraction of the lower tail of the cost distribution. Such a contraction is consistent with a model in which potential migrants have heterogeneous and limited knowledge of low-cost crossing locations and fence construction has a large deterrent effect precisely because those locations that are fenced disproportionately represent the lowest-cost crossing locations.

Figure 2.5 presents a graphical depiction of the imputed shift in migration costs induced by fence construction. Panel A plots the cumulative distribution function of migration cost π, which is assumed for the purposes of this simulation to be normally distributed. Panel B presents a magnified view of the lower tail of the cost distribution and denotes the reservation value of π for high/low residual earners from the

[31] Although the assumption of normally-distributed costs allows for simple closed-form expressions for the relevant cross-partials, the basic result is unchanged when other common distributions are used instead.

period before the start of fence construction (pre-fence migration rates for these two subgroups are 0.26% and 0.38%, respectively). In Panel C, the darker-shaded curve depicts the change in the cumulative distribution function of migration costs that would result from a positive mean shift. Panel C plots the subgroup migration rates that would be observed after such a cost shift and demonstrates the fact that a positive mean shift can rationalize a larger percentage-wise decline in migration rates for high residual earners, but cannot explain the larger absolute decline observed in the data. In contrast, Panel D depicts an alternative shift in the cumulative distribution function of migration costs that can match the migration rates observed for high/low residual earners in the period after the start of fence construction (0.06% and 0.21%, respectively). Panel D makes the point graphically that the data can only be rationalized by a contraction of the lower tail of the cost distribution.

The fact that changing migration patterns are consistent with a sharp contraction in the lower tail of the border crossing cost distribution is not itself informative about the magnitude of the change in crossing costs. While the observed reduction in migration to the U.S. in response to fence construction is large, it can potentially be explained by either large increases in the costs faced by potential migrants, or by small cost increases combined with low perceived benefits of migration. However, two pieces of evidence suggest that the former interpretation is the correct one. First, the deterrent effect of the fence is large relative to the 11% decrease in migration associated with a one-point increase in log income. (A one-point increase in log income significantly reduces the 200% average earnings gain associated with migrating to the U.S.) Second, the number of crossing deaths per migration episode has roughly doubled from 0.0005 to 0.0010 in the years since the start of fence construction [Haddal, 2010]. If we place the value of a statistical life at $6.0 million based on a 2011 Department of Transportation review of the relevant research [Rivkin and Trottenberg, 2011], the implied increase in the cost of crossing from the increased risk of death alone is $3,000, which is larger than the average cost of hiring a smuggler (according to Mexican Migration Project data). Taken together, these findings provide convincing evidence that the rise in crossing costs associated with fence construction is likely quite large.

2.6 Conclusions

My analysis concludes that fence construction significantly reduced migration from Mexico to the United States. I find that there are geographic spillovers in construction impacts, as border municipality residents are deterred by construction in both their own and adjacent municipalities. In addition, non-border state residents with historically limited access to smugglers are significantly less likely to migrate to the United States after the start of fence construction. I also show that observed changes in migration rates based on the sociodemographic characteristics of potential migrants are not consistent with a model in which fence construction simply increases mean migration costs by increasing the expected distance that each migrant must travel to cross the border. Instead, the results are rationalized by a contraction of the lower tail of the cost distribution. This finding buttresses the theory that access to low-cost crossing locations is limited and that the fence has a large deterrent effect precisely because it targets these locations. I argue that the cost changes induced by fence construction are significant based on a comparison of fence impacts to the reduction in migration associated with increased earnings in Mexico and based on an evaluation of the costs associated with increased risk of crossing-related mortality.

Having estimated the impact of border fence construction on migration patterns and on the structure of crossing costs, a worthwhile exercise is to estimate the implicit cost paid by DHS per migrant deterred. Given a population of 5.4 million in Mexican border municipalities and a baseline migration rate of 0.25% per quarter, the estimated 39% reduction in migration resulting from fence construction translates into roughly 5,250 border municipality migrants deterred per quarter. In addition, border fence construction is estimated to reduce migration in non-border states with historically low access to smugglers by 38%. Given a population of 36.7 million in the non-border states with low smuggler access in my sample and a baseline migration rate of 0.26% per quarter, this implies that approximately 36,250 non-border migrants and 41,500 total migrants are deterred by the fence each quarter.

Given that passage of the Secure Fence Act was followed by fence construction in 15 of Mexico's 37 municipalities adjacent to the border, and given an estimated total cost of $6.5 billion and a fence lifespan of 20 years [Stana et al., 2009b], I estimate a cost of $16 billion to fence the border for 20 years. Based on this cost estimate and the estimated number of migrants deterred, I calculate a total cost of approximately $4,820 per deterred migrant.

Importantly, this figure is sensitive to assumptions along a number of dimensions. First, it is likely the case that construction in remaining unfenced areas is more costly than was construction in initially fenced areas, which at least partly explains why some areas were fenced before others. Second, while I have assumed a constant deterrent effect over the next 20 years, the fact that the fence was only recently constructed means that I cannot distinguish delayed from permanently deterred migration in my analysis. Third, the number I estimate abstracts from the noisily estimated impacts of fence construction on return migration of previous municipality residents currently living in the United States which could bias my estimate in either direction.[32]

In addition to these caveats, any cost-benefit analysis that ignores changes in the composition of migrants and the general equilibrium impacts of fence construction is incomplete. I have shown that the deterrent effect of the fence is driven by its impact on those with lower educational attainment and higher residual earnings. While U.S. employers may benefit from increased mean education of newly-arrived migrants, this form of positive migrant selection presumably weakens Mexican labor markets by draining human capital. In addition, in ongoing research, I find that fence construction significantly reduces earnings of border

[32] I also do not include spillover estimates in the calculation of total migrants deterred. While I have shown that the fence has impacts on residents of adjacent municipalities, noisily estimated interaction terms in these specifications make it difficult to accurately calculate the number of migrants deterred due to fence construction spillovers. If I do calculate the number of deterred migrants based on point estimates from the Table 2.4 spillovers specifications, this would imply that the average cost per migrant deterred is actually $4,580.

municipality residents, seemingly due to the contraction of local migration-related economic activity. This negative impact on local economies may increase instability in a region that already represents a significant security threat to communities on the U.S. side of the border. In an era when increasing international migration flows motivate destination country governments to enact policies aimed at deterring migration by raising its cost, a greater research emphasis on the mechanisms and subpopulations driving estimated impacts, and on the costs imposed on non-migrants, can help shed light on the efficacy of such efforts.

Table 2.1. Impact of Fence Construction on Border Municipality Migration

	Migrate to United States							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Fence Construction	-0.486***	-0.367***	-0.404**	-0.322***	-0.486***	-0.367***	-0.464***	-0.328***
	(0.162)	(0.054)	(0.163)	(0.077)	(0.163)	(0.054)	(0.161)	(0.072)
Observations	1,042,959	1,024,343	717,589	703,190	1,039,700	1,021,084	1,042,959	1,024,343
Municipality Fixed Effects	X	X	X	X	X	X	X	X
Year-Quarter Fixed Effects	X	X	X	X	X	X	X	X
State * Year-Quarter Fixed Effects		X		X		X		X
Mean of Non-Fenced	0.0025		0.0025		0.0025		0.0025	
	[0.050]		[0.050]		[0.050]		[0.050]	
Sample	All municipalities with proposed fence		Panel defined by statewise date of initial fence construction		Balanced panel (by municipality)		Control for adjacent fencing included	

Notes:

1 All regressions employ logit specifications and drop the first observation for each household member since the format of the survey means that a household member cannot be listed as out-migrating at the time of initial interview. Fence Construction is an indicator variable equal to "1" if border fence construction has begun in a given municipality. "Panel defined by statewise date of initial fence construction" includes, for each state, all quarters up to six years before and up to three years after fence construction began in that state. Balanced panel (by municipality) includes the 18 municipalities which have data available for all years from 2003 to 2011.

2 * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered by municipality.

Table 2.2. Impact of Fence Construction on Migration from U.S. to Mexico

	In-Migrate (from United States to Mexico)					
	(1)	(2)	(3)	(4)	(5)	(6)
Fence Construction	0.589	-0.247	-0.00004	-0.081	0.064	-0.184
	(0.797)	(0.955)	(0.258)	(0.333)	(0.230)	(0.277)
Mean of Non-Fenced	0.0014		0.005		0.0006	
	[0.004]		[0.068]		[0.025]	
Observations	742	742	337,197	307,863	1,042,959	952,129
Municipality Fixed Effects	X	X	X	X	X	X
Year-Quarter Fixed Effects	X	X	X	X	X	X
State * Year-Quarter Fixed Effects		X		X		X
Specification	WLS		Logit		Logit	
Unit of Analysis	Municipality*year-quarter		Household*year-quarter		Individual*year-quarter	

Notes:

1 Fence Construction is an indicator variable equal to "1" if border fence construction has begun in a given municipality. All regressions include only municipalities with fence proposed in the 2006 Secure Fence Act.

2 Columns (1)-(2) employ weighted least squares (WLS) specifications, and include one observation per municipality*year-quarter with observations weighted by number of responses within each municipality*year-quarter cell. The dependent variable in Columns (1)-(2) is the fraction of respondents who returned from the U.S. in a given year-quarter. Columns (3)-(4) use a household*year-quarter-level specification with an indicator variable for whether any household member in-migrated in a given year-quarter as the dependent variable, and Columns (5)-(6) use an individual*year-quarter-level specification with an indicator variable for whether each household member in-migrated in a given year-quarter as the dependent variable (regardless of where the household member resided in the previous year-quarter).

3 * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered by municipality.

Table 2.3. Impact of Fence Construction on Migration from U.S. to Mexico and on U.S. Labor Market Outcomes (U.S. Data)

	Undocumented (Probabilistic)		Recent Migrant		Employed		Employed		Separations	Hires
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Fence Construction	0.003	-0.005	-0.761**	-0.016*	-0.061***	-0.014**	-0.143**	-0.035	-3211.5**	-3914.8**
	(0.067)	(0.012)	(0.356)	(0.008)	(0.019)	(0.004)	(0.073)	(0.018)	(1262.6)	(1478.0)
P-value (Wild Bootstrap)		*0.238*		*0.099*		*0.020*		*0.119*	*0.020*	*0.020*
Mean of Non-Fenced	0.167		0.020		0.643		0.570		10653.9	11077.9
	[0.373]		[0.138]		[0.479]		[0.495]		[17576.7]	[18484.3]
Observations	238,645	238,645	17,234	17,234	238,645	238,645	42,257	42,257	1,396	1,396
County Fixed Effects	X	X	X	X	X	X	X	X	X	X
Year Fixed Effects	X	X	X	X	X	X	X	X	X	X
Specification	Logit	OLS	Logit	OLS	Logit	OLS	Logit	OLS	OLS	OLS
Sample	Aged 16-65		Aged 16-65, Mexican, not a college graduate		Aged 16-65		Aged 16-65, Mexican, not a college graduate		U.S. Border Counties (County-level Specification)	

Notes:

1 Fence Construction is an indicator variable equal to "1" if border fence construction has begun in a given county. All regressions include only U.S. counties with fence proposed in the 2006 Secure Fence Act.

2 Undocumented (Probabilistic) is a dummy variable for whether a respondent has the following characteristics: aged 16-65, born in Mexico, non-college graduate. Recent Migrant is defined as a dummy variable for whether an individual arrived in the United States within the past three years.

3 * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered by municipality.

Table 2.4. Geographic Spillovers in Fence Construction Impacts

	Migrate to United States		Number of Adjacent Municipalities with Fence Construction		Max # Adjacent Fenced
	(1)	(2)	(3)	(4)	(5)
Fence Construction	-0.613***	-1.431**	-0.058	-0.172	
	(0.213)	(0.517)	(0.304)	(0.399)	
P-value (Wild Bootstrap)		*0.040*			
Number of Adjacent Municipalities Fenced	-0.393**	-0.948**			
	(0.175)	(0.397)			
P-value (Wild Bootstrap)		*0.040*			
Fence Construction*Number of Adjacent Municipalities Fenced	0.329	0.683			
	(0.233)	(0.438)			
P-value (Wild Bootstrap)		*0.139*			
Number of Fenced Municipalities Two Away	-0.135	-0.283			
	(0.158)	(0.322)			
P-value (Wild Bootstrap)		*0.376*			
Max Fence Construction					0.133
					(0.335)
Mean of Non-Fenced	0.0025		0.138		0.545
	[0.050]		[0.346]		[0.522]
Observations	1,042,959	1,043,142	790	790	22
Municipality Fixed Effects	X	X	X	X	
Year-Quarter Fixed Effects	X	X	X	X	
State * Year-Quarter Fixed Effects				X	
Specification	Logit	OLS	OLS	OLS	OLS
Sample	All municipalities with proposed fence		One observation per municipality-quarter		One observation per municipality

Notes:

1 Columns (1)-(2) employ logit and OLS specifications and drop the first observation for each household member since the format of the survey means that a household member cannot be listed as out-migrating at the time of initial interview. Columns (3)-(5) employ OLS specifications. Fence Construction is an indicator variable equal to "1" if border fence construction has begun in a given municipality, and Max Fence Construction is the municipality-level maximum of the Fence Construction variable.

2 All OLS coefficients are scaled by 10^3.

3 * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered by municipality.

Table 2.5. Impact of Fence Construction on non-Border State Migration

	Migrate to United States			
	(1)	(2)	(3)	(4)
Weighted Avg. Fence Construction	-0.190	-0.283	-0.471***	-0.661**
	(0.134)	(0.215)	(0.139)	(0.240)
Mean of Non-Fenced	0.0024			
	[0.049]			
Observations	11,269,524	11,269,524	5,359,440	5,359,440
State Fixed Effects	X	X	X	X
Year Fixed Effects	X	X	X	X
	Logit	OLS	Logit	OLS
Sample	All non-border states with Mexican Migration Project data		50% of non-border states with below-median coyote use (<72%)	

Notes:

1 All regressions drop the first observation for each household member since the format of the survey means that a household member cannot be listed as out-migrating at the time of initial interview. Weighted Average Fence Construction is a state-level average generated by weighting the Fence Construction values of border municipalities in proportion to the share of pre-2000 migrants born in a given state who crossed in that municipality. Historical state-level coyote use is defined based on all pre-2000 crossings included in the Mexican Migration Project dataset.

2 All OLS coefficients are scaled by 10^3.

3 * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered by state.

Table 2.6. Heterogeneous Impacts of Fence Construction on Migration

	Migrate to United States	
	(1)	(2)
Fence Construction	-0.484**	-2.44**
	(0.221)	(0.583)
P-value (Wild Bootstrap)		*0.039*
Years of Schooling	0.021	0.026
	(0.013)	(0.052)
P-value (Wild Bootstrap)		*0.911*
Fence Construction*Years of Schooling	0.061***	0.160***
	(0.013)	(0.049)
P-value (Wild Bootstrap)		*0.000*
Log Income	-0.118***	-0.266**
	(0.011)	(0.041)
P-value (Wild Bootstrap)		*0.020*
Fence Construction*Log Income	-0.116***	-0.176*
	(0.014)	(0.065)
P-value (Wild Bootstrap)		*0.079*
Mean of Non-Fenced	0.0025	
	[0.050]	
Observations	624,580	632,565
Municipality Fixed Effects	X	X
Year-Quarter Fixed Effects	X	X
Specification	Logit	OLS
Sample	All municipalities with proposed fence	

Notes

1 All regressions drop the first observation for each household member since the format of the survey means that a household member cannot be listed as out-migrating at the time of initial interview. All regressions include only individuals aged 16-60. Fence Construction is an indicator variable equal to "1" if border fence construction has begun in a given municipality.

2 All OLS coefficients are scaled by 10^3.

3 * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered by municipality.

Figure 2.1, Panel A. Migration Trends (All Migration Episodes)

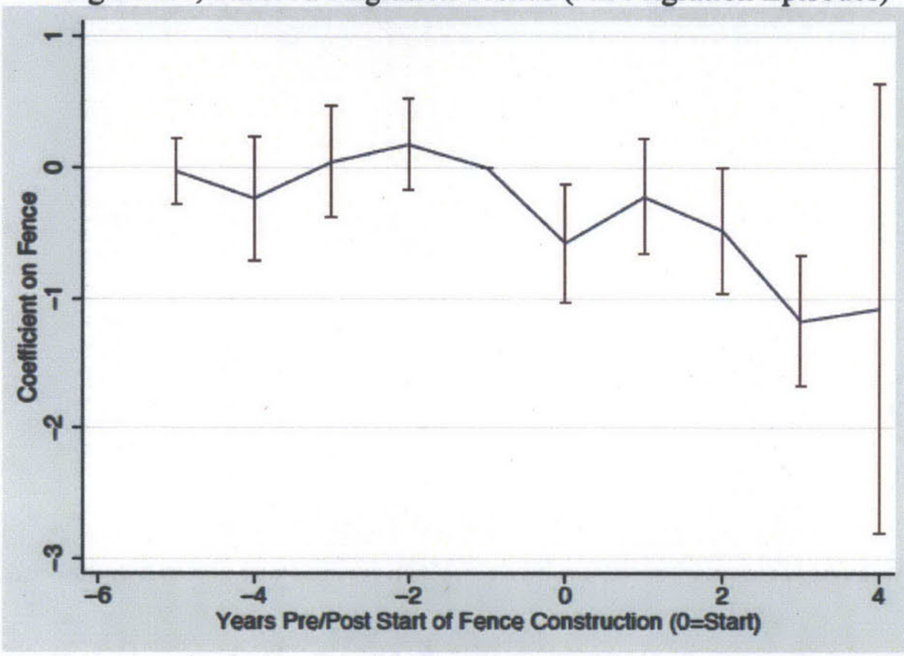

Figure 2.1, Panel B. Migration Trends (Only Employment-Related Migration)

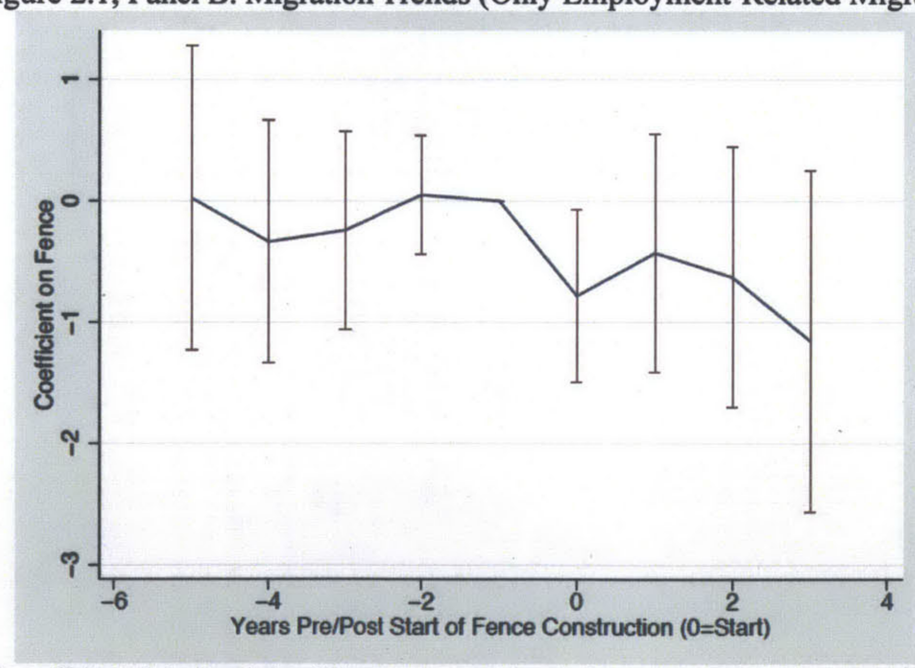

Graphs plot the coefficients of a logit regression of migration on a set of dummies for years pre- and post- fence construction. t=0 represents the start of fence construction and t=-1 is the omitted year dummy (with the coefficient plotted as equal to "0"). The regression includes municipality and year-quarter fixed effects, and clusters standard errors by municipality. The Panel A plot includes all migration episodes in the dependent variable definition, while the Panel B plot includes only migration undertaken for employment.

Figure 2.2, Panel A. Migration Trends (All Migration Episodes), OLS

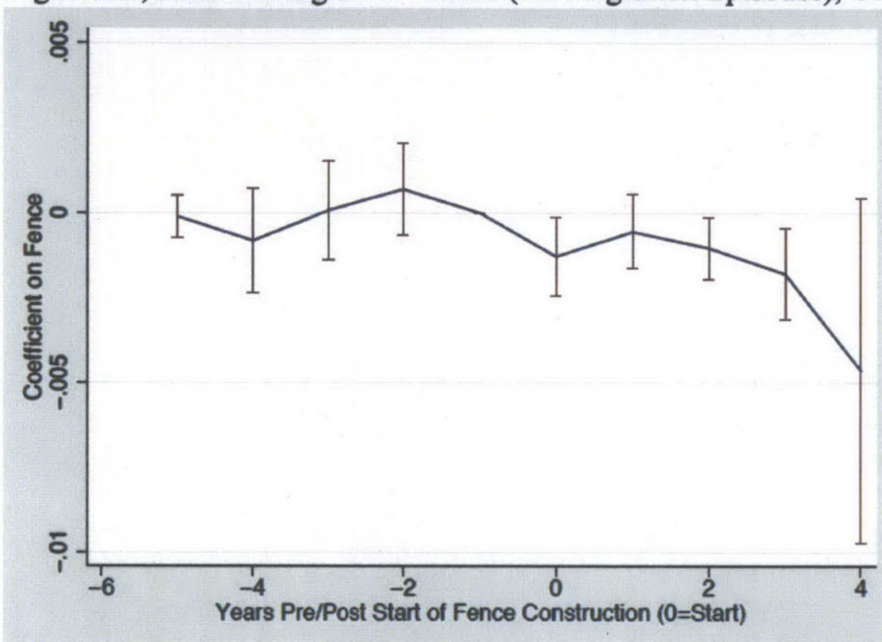

Figure 2.2, Panel B. Migration Trends (Only Employment-Related Migration), OLS

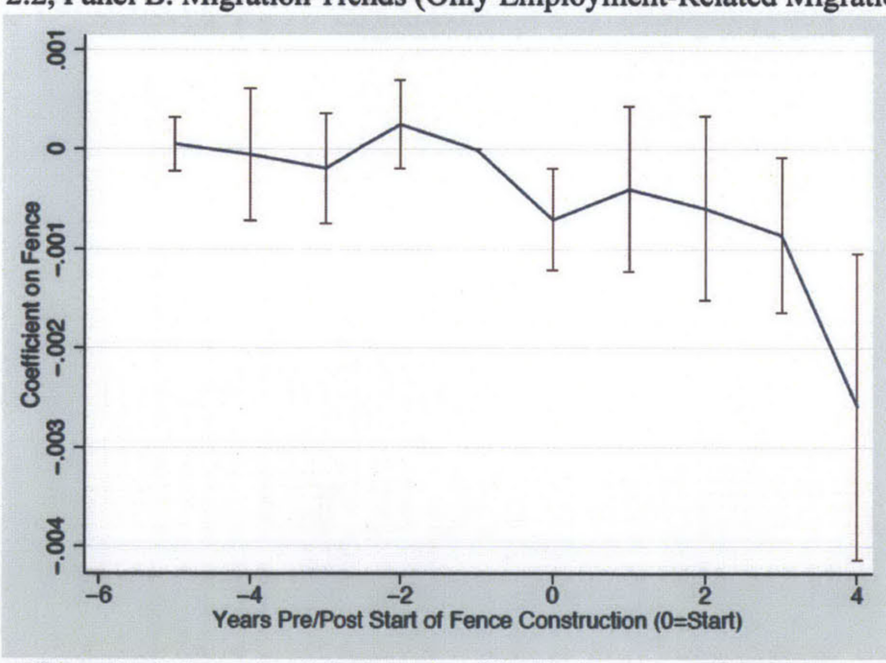

Graphs plot the coefficients of an OLS regression of migration on a set of dummies for years pre- and post- fence construction. t=0 represents the start of fence construction and t=-1 is the omitted year dummy (with the coefficient plotted as equal to "0"). The regression includes municipality and year-quarter fixed effects, and clusters standard errors by municipality. The Panel A plot includes all migration episodes in the dependent variable definition, while the Panel B plot includes only migration undertaken for employment.

Figure 2.3. Income Trends

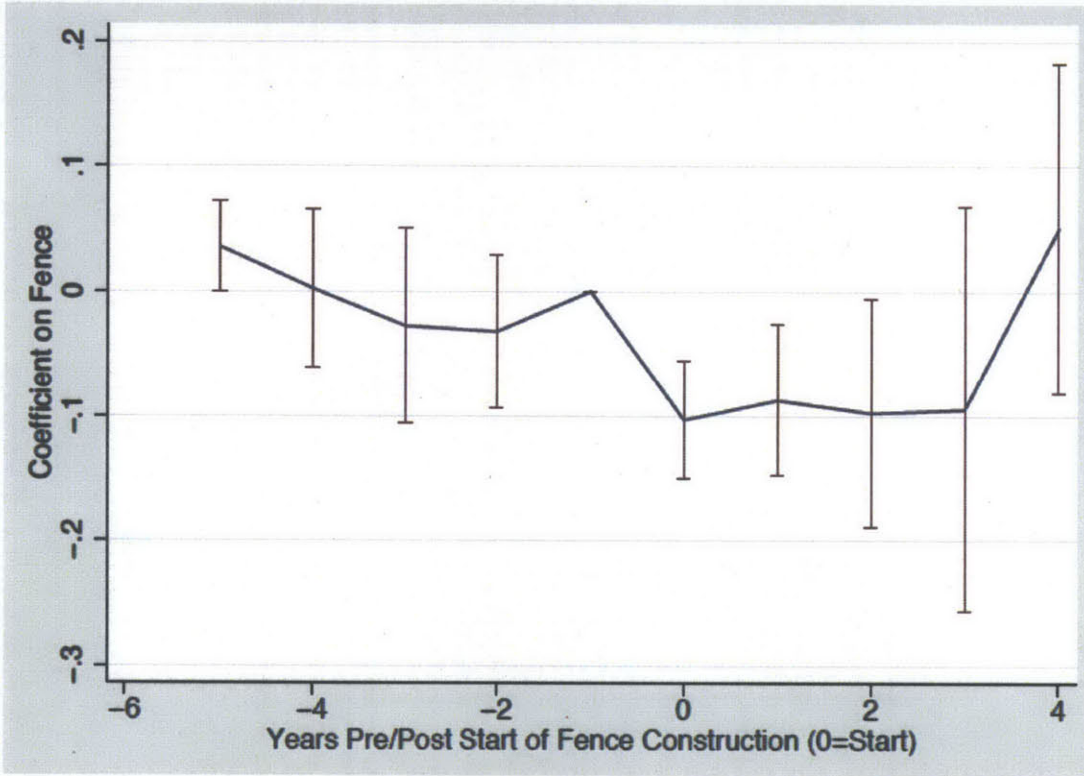

Graph plots the coefficients of a Poisson quasi-maximum likelihood estimation regression of income on a set of dummies for years pre- and post- fence construction. t=0 represents the start of fence construction and t=-1 is the omitted year dummy (with the coefficient plotted as equal to "0"). The regression includes municipality and year-quarter fixed effects, and clusters standard errors by municipality.

Figure 2.4. Years in U.S.

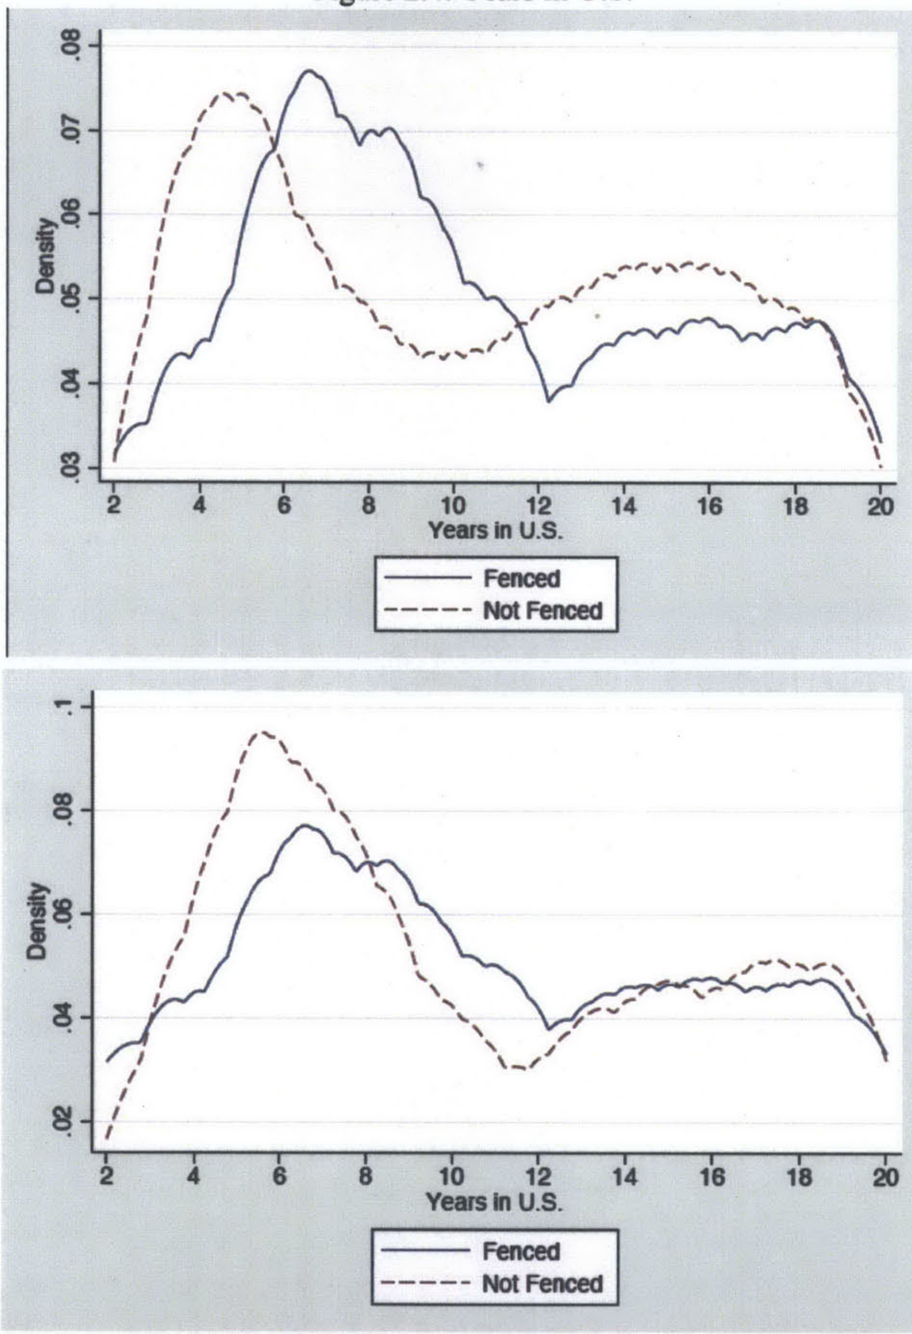

The above panel uses U.S. Current Population Survey data to plot the distribution of years lived in the U.S. by Mexican-born respondents surveyed in years 2000-2011, and the panel below plots the distribution of years lived in the U.S. by Mexican-born respondents surveyed in years 2007-2011. Fenced indicates whether border fence was ever constructed adjacent to the U.S. Metropolitan Statistical Area in which the respondent lived.

Figure 2.5. Migration Cost Curves

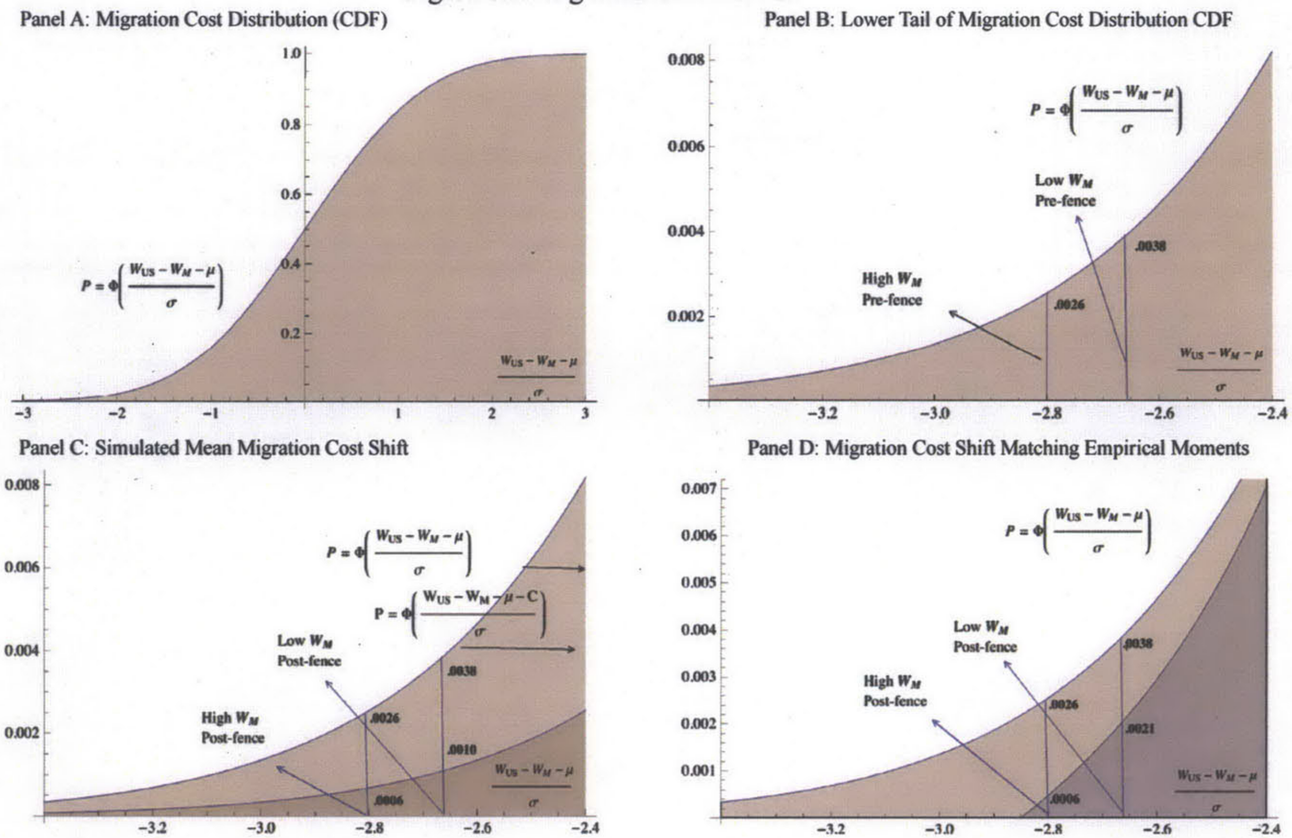

Panel A plots the CDF of migration costs, which is assumed here to be normally distributed. Panel B presents a magnified view of the lower tail of the cost distribution and denotes the pre-fence reservation value of migration costs for high/low residual earners. In Panel C, the darker-shaded curve depicts the change in the CDF of migration costs that would result from a positive mean shift in crossing costs. Panel D depicts an alternative shift in the CDF that can match the post-fence migration rates observed for high/low residual earners.

Appendix Table 2.1. Impact of Fence Construction on Border Municipality Migration (OLS)

	Migrate to United States							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Fence Construction (x 10^3)	-1.090*	-1.230**	-0.920*	-0.999**	-1.090*	-1.220***	-1.044*	-1.182**
	(0.460)	(0.304)	(0.509)	(0.273)	(0.464)	(0.305)	(0.467)	(0.305)
P-value (Wild Bootstrap)	*0.080*	*0.020*	*0.060*	*0.040*	*0.060*	*0.000*	*0.080*	*0.040*
Observations	1,043,142	1,043,142	717,772	717,772	1,039,700	1,039,700	1,043,142	1,043,142
Municipality Fixed Effects	X	X	X	X	X	X	X	X
Year-Quarter Fixed Effects	X	X	X	X	X	X	X	X
State * Year-Quarter Fixed Effects		X		X		X		X
Mean of Non-Fenced	0.0025		0.0025		0.0025		0.0025	
	[0.050]		[0.050]		[0.050]		[0.050]	
Sample	All municipalities with proposed fence		Panel defined by statewise date of initial fence construction		Balanced panel (by municipality)		Control for adjacent fencing included	

Notes:

1 All regressions employ ordinary least squares (OLS) specifications and drop the first observation for each household member since the format of the survey means that a household member cannot be listed as out-migrating at the time of initial interview. Fence Construction is an indicator variable equal to "1" if border fence construction has begun in a given municipality. "Panel defined by statewise date of initial fence construction" includes, for each state, all quarters up to six years before and up to three years after fence construction began in that state. Balanced panel (by municipality) includes the 18 municipalities which have data available for all years from 2003 to 2011.

2 * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered by municipality.

Appendix Table 2.2. Impact of Fence Construction on Border Municipality Migration (Municipality-Level WLS)

	Migrate to United States							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Fence Construction (x 10^3)	-1.063*	-1.159**	-0.875	-0.925**	-1.062*	-1.116***	-1.042**	-1.180**
	(0.488)	(0.371)	(0.544)	(0.343)	(0.493)	(0.374)	(0.490)	(0.371)
P-value (Wild Bootstrap)	*0.099*	*0.014*	*0.129*	*0.020*	*0.090*	*0.010*	*0.050*	*0.028*
Observations	742	742	606	606	713	713	742	742
Municipality Fixed Effects	X	X	X	X	X	X	X	X
Year-Quarter Fixed Effects	X	X	X	X	X	X	X	X
State * Year-Quarter Fixed Effects		X		X		X		X
Mean of Non-Fenced	0.0025		0.0025		0.0025		0.0025	
	[0.050]		[0.050]		[0.050]		[0.050]	
Sample	All municipalities with proposed fence		Panel defined by statewise date of initial fence construction		Balanced panel (by municipality)		Control for adjacent fencing included	

Notes:

1 All regressions employ weighted least squares (WLS) specifications, and include one observation per municipality*year-quarter with observations weighted by number of responses within municipality*year-quarter cell. Fence Construction is an indicator variable equal to "1" if border fence construction has begun in a given municipality. "Panel defined by statewise date of initial fence construction" includes, for each state, all quarters up to six years before and up to three years after fence construction began in that state. Balanced panel (by municipality) includes the 18 municipalities which have data available for all years from 2003 to 2011.

2 * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered by municipality.

Appendix Figure 2.1. Sample Fence Map

Panel A: Map of Mexican States

Panel B: Sonora, 2007 Quarter II

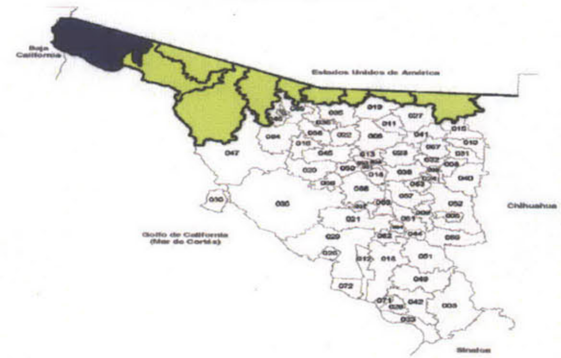

Panel C: Sonora, 2007 Quarter III

Panel D: Sonora, 2007 Quarter IV

Notes: Panel A highlights the border state of Sonora. Panels B-D depict border municipalities within Sonora in subsequent quarters of 2007. Fenced municipalities are shaded blue and unfenced municipalities are shaded yellow.

Chapter 3

The Economic Returns to Social Interaction: Experimental Evidence from Microfinance

3.1 Introduction

Social capital, famously defined by Putnam [1993] as "features of social organization, such as trust, norms and networks, that can improve the efficiency of society by facilitating coordinated actions," is considered particularly valuable in low-income countries where formal insurance is largely unavailable and institutions for contract enforcement are weak.[1] Since economic theory suggests that repeat interaction among individuals can help build and maintain social capital, encouraging interaction may be an effective tool for development policy. Indeed, numerous development assistance programs emphasize social contact among community members under the assumption of significant economic returns to regular interaction. But can simply inducing individuals to interact with one another actually facilitate economic cooperation?

Rigorous evidence on this question remains limited, largely due to the difficulty of accounting for endogenous social ties [Manski, 1993, 2000]. For instance, if more trustworthy individuals or societies are characterized by denser social networks, we cannot assign a causal interpretation to the positive association between community-level social ties and public good provision. For similar reasons, it is also not possible to assign a causal interpretation to the higher levels of cooperation observed among friends relative to strangers in laboratory public goods games.[2] In short, without randomly varying social distance, it is difficult to

[1] For instance, Guiso et al. [2004] demonstrate that residents in high social capital regions undertake more sophisticated financial transactions, and Knack and Keefer [1997] show that a country's level of trust correlates positively with its growth rate.

[2] The community ties literature includes Costa and Kahn [2003], Alesina and La Ferrara [2002], DiPasquale and Glaeser

validate the model of returns to repeat interaction and even harder to determine whether small changes in social contact can produce tangible economic returns.

The first contribution of this paper is to undertake exactly this exercise. By randomly varying how often individuals meet, we provide causal evidence on the returns to repeat social interaction. We do so in the context of a development program that emphasizes group interaction: microfinance.[3] In the typical "Grameen Bank"-style microfinance program, clients meet weekly in groups to make loan payments. Our experiment varied social interaction by randomly assigning 100 first-time borrower groups of a typical microfinance institution (MFI) in India to either meet on a weekly or a monthly basis throughout their ten-month loan cycle. Using administrative and survey data we study the effect of short-run increases in group meeting frequency on long-run social contact and an important measure of economic vulnerability: default incidence in the subsequent loan cycle.

A second contribution of this paper is to identify a key mechanism through which group lending sustains high repayment rates: risk-pooling among clients. While the theoretical literature largely emphasizes the importance of joint-liability contracts for reducing default in microfinance, recent experimental evidence suggests that joint liability per se has little impact on default [Gine and Karlan, 2009], raising anew the question of how exactly group lending achieves risk reduction without collateral. Since our clients received individual-liability debt contracts, we can isolate how a less noted feature of the classic group lending contract – encouraging social interaction via group meetings – reduces default.[4] In other words, even absent the explicit incentives for monitoring and enforcement that joint liability provides, frequent group meetings can lower lending risk by increasing social interaction among group members and, as a consequence, strengthening risk-pooling arrangements within social networks.

Our evidence consists of several striking changes in client behavior associated with experimentally in-

[1999], Miguel et al. [2005], Olken [2009], while laboratory games literature includes Glaeser et al. [2000], Carter and Castillo [2005], Do et al. [2009], Karlan [2005], Ligon and Schecter [2011].

3 Related work include Dal Bo [2005] who provide laboratory game evidence on returns to repeat *economic* interaction, where the likelihood of future rounds of exchange is randomly assigned and Humphreys et al. [2009]'s field experiment which shows that community development programs randomly assigned to villages encourage pro-social behavior (but cannot isolate the influence of social interaction from other program aspects).

4 The remarkable success of microfinance in achieving very high repayment rates on collateral-free loans to poor individuals is widely recognized, as evidenced by awarding of the Nobel Peace Prize to the Grameen Bank founder. Our findings complement theoretical research on the role of social collateral in microfinance and empirical work that identifies a significant correlation between social connections and default risk [Besley and Coate, 1995, Ghatak and Guinnane, 1999, Karlan, 2005]. For instance, MFI clients in Peru who are more trustworthy in a trust game are less likely to default, and group-level default is lower in groups where clients have stronger social connections [Karlan, 2005, 2007]. In Gine and Karlan [2009], the shift from joint to individual liability increased default among borrowers with ex-ante weak social ties. Fischer and Ghatak [2010] show that microfinance repayment schedules are attractive to present-biased borrowers and consistent with this, Bauer et al. [2012] shows that microfinance borrowers are relatively more likely to be present-biased. It is likely that microfinance induced gains in social collateral, which improve informal insurance arrangements, are particularly valued by these borrowers. In our setting, Field et al. [2013] show that impatient borrowers benefit less from added flexibility in the form of a grace period.

creasing the frequency of client contact. First, clients assigned to weekly groups during their first loan cycle increased social contact with group members outside of meetings, and sustained it in the long run. More than a year after the experiment ended, clients who had met on a weekly basis during their first loan saw each other 37% more often outside of group meetings. If client groups remained fixed for multiple loan cycles, then treatment and control groups should converge in terms of their degree of connectedness, in which case we would no longer observe long-run differences in the degree of social interaction according to first-intervention meeting frequency. Instead, the persistence of the difference makes sense in our setting given that, due to a policy change immediately after our experiment that reduced loan groups from ten to five members, the majority of pairs (68%) of first loan group members were no longer in the same loan group at follow-up. As our long-run social contact data show, clients continue to interact outside of loan groups and treatment clients do so at a significantly higher rate.

Second, greater social interaction among clients on a weekly schedule was accompanied by increased willingness to pool risk relative to monthly clients. Here, our evidence comes from a field-based lottery game conducted roughly 16 months after the first loan cycle ended. The lottery operated much like a laboratory trust or solidarity game, but in a real-world setting. Each client was entered into a (separate) promotional lottery for the MFI's new retail store. The client started with a 1 in 11 chance of winning the lottery prize, a voucher redeemable at the MFI store. She was then offered the opportunity to give out additional lottery tickets to any number of members of her first loan group.

Since ticket-giving reduces a client's individual chances but increases the probability that someone from the group would win, it captures either her unconditional altruism towards or willingness to risk-share with members of her initial group. To distinguish insurance motivations from unconditional altruism, we randomized the divisibility of the lottery prize. Assuming the more easily divisible prize reduces transaction costs of sharing and/or is perceived as more conducive to sharing, then a client should give more tickets when the prize is divisible if she is motivated at least in part by risk-sharing considerations, but should not if her sole motivation is unconditional altruism.[5]

Relative to a monthly client, a client who had been assigned to a weekly group two years prior was 32% more likely to enter a group member into the lottery when the prize was divisible, but only 16% more likely when it was not.

Finally, we show that clients on a weekly schedule were, in the long run, better able to endure financial shocks. Those who met weekly during their first loan cycle were three times (5.2 percentage points) less likely to default on their subsequent loan, despite the fact that all clients had reverted to the same repayment schedule.

To disentangle the role of meeting frequency from repayment frequency, we use a second treatment arm

[5] Similar variations of dictator or trust games have been used to parse out motives for giving [Ligon and Schecter, 2011, Do et al., 2009, Carter and Castillo, 2004]. Most similar to us, Gneezy et al. [2000] use a sequence of trust games with varying constraints on the amount that can be returned to show that individuals contribute more when large repayments are feasible.

in which clients were assigned to meet weekly but maintained a monthly repayment schedule. To address variation in actual occurrence of non-repayment meetings in this arm, we implement an Instrumental Variable (IV) strategy, which exploits the fact that loan officers were more likely to cancel a non-repayment meeting on days with heavy rainfall. Our IV estimates show that the default rate difference remains in magnitude and significance when we compare monthly clients to clients randomly assigned to meet weekly but repay monthly. Based on this evidence, we conclude that default risk falls on account of meeting more frequently rather than differences in fiscal habits that could arise from requiring clients to initially *repay* at more frequent intervals.

To summarize, a higher degree of short-run interaction is associated with increased social interaction in the long run, improved risk-sharing arrangements among clients and lower default. Our findings are consistent with several mechanisms through which social interactions reduces default, including improved monitoring, better information flows, lower transaction costs for risk-sharing and better ability to punish potential deviations from risk-sharing, which we are unable to disentangle. However, our findings substantiate theoretical claims that repeat interaction can yield economic returns by facilitating informal economic exchange, and provide an alternative explanation for the success of the group lending model. More generally, the findings demonstrate that tweaking the design of standard development programs to encourage social interaction can generate economically valuable social capital.

The paper is structured as follows: Section 3.2 describes the experimental design. Section 3.3 examines how randomized differences in meeting frequency, implemented only during the first loan cycle, influenced long-run social interaction and client willingness to share in the field lottery. Section 3.4 documents changes in long-run default rates and separates the role of meeting frequency from that of repayment frequency. Section 3.5 concludes.

3.2 Experimental Design

3.2.1 Setting

Our MFI partner, Village Financial Services (VFS), operates in the Indian state of West Bengal. In 2006 when we began our field experiment, it had $6.75 million in outstanding loans to over 56,000 female clients. VFS' gross loan portfolio to total asset ratio of 78% placed it slightly below the median Indian MFI (84%) while its portfolio at risk of 0.47% (defined as payments outstanding in excess of 30 days) was identical to the median Indian MFI [MIX Market, 2012].

Our study population consisted of first-time VFS clients living in peri-urban slums of the city of Kolkata. At the time of joining the MFI, over 70% of client households owned a business and the median client's household income placed her just below the dollar-a-day poverty line. Study population demographics, such as income, home ownership and home size, are largely comparable with similar MFIs operating in other Indian cities (Appendix Table 3.2). However, consistent with cross-city differences in MFI penetration, clients in our sample exhibit significantly lower rates of borrowing outside of the MFI.

3.2.2 Sample

Between April and September 2006 we recruited 100 first-time microfinance groups from neighborhoods in the catchment areas of three VFS branches. Following VFS protocol, the loan officer first surveyed the neighborhood and then conducted a meeting to inform female residents about the VFS loan product. Interested women were invited for a five-day training program, where clients met for an hour each day and learned about the benefits and responsibilities of the loan. At the end of the five days, the loan officer assigned women into groups of ten and identified a leader of each group.[6] Thus, clients in a single loan group lived in close proximity and were typically acquainted prior to joining. Although 63% of group members in our sample knew one another at group formation, most described their relationship with other group members as neighbors (48%) rather than friends (7%) or family (8%).

3.2.3 Experimental Design

Group Assignment At the end of the group formation process, each group member was offered an individual-liability loan of Rs. 4,000 (~\$100) and told that her repayment schedule would be assigned at the time of loan disbursal. Prior to loan disbursal, groups were randomized into either weekly or monthly schedules. In total, 38 groups were assigned to the control arm in which group meetings were held on a monthly basis, and 30 groups were assigned to the treatment arm in which group meetings occurred weekly (Treatment 1). In addition, 32 groups were assigned to an alternative treatment in which they met weekly but repaid monthly (Treatment 2), an artificial contract design for the purpose of microfinance delivery, but one that allows us to disentangle the influence of meeting frequency from the influence of repayment frequency for scientific purposes.

At loan disbursal, Treatment 1 groups were informed that they were to repay their loans in 44 weekly installments of Rs. 100 (a reasonably small amount given average weekly household earnings of Rs. 1,167), while Control and Treatment 2 groups were told that they would repay in 11 monthly installments of Rs. 400. No client dropped out after her repayment schedule was announced.

Meeting Protocol Repayment in a group setting is an integral part of MFI lending practice, and VFS followed a relatively standard "Grameen Bank" group meeting model. Each group was assigned a loan officer who conducted the meeting in the group leader's house. The average meeting lasted 18 minutes, during which clients took an oath promising regular repayment, and deposited payment with the loan officer and had their passbooks marked.[7] Thus, a client's repayment behavior was observable to other group members, although in practice most clients socialized while awaiting their turn. Anecdotally, socializing happens en

[6] Loan officers aimed to form ten-member groups. In practice, group size ranged between nine and thirteen members, with 77% ten-member groups.

[7] While the oath encourages group responsibility for loans, the loan contract is individual liability.

route to meetings, while waiting for the loan officer to arrive and begin meetings and while waiting for one's turn to pay.[8]

Overall, Control and Treatment 1 groups closely followed the assigned meeting schedule: No Control group met less than five or more than eleven times and no Treatment 1 group met less than 23 or more than 44 times, which were the minimum and maximum meetings allowed by the respective contracts.[9] While in theory clients could skip meetings and send their payment with another group member, it was rare for clients to do so, and average attendance at repayment meetings was 81%.

Treatment 2 groups did not strictly adhere to the experimental protocol: Only half of the groups met at least the minimum required number of times (23) and average attendance at meetings was only 56%. As this compliance issue necessitates a more complicated econometric strategy, we first present experimental estimates which compare Control and Treatment 1 groups only. Then, in order to identify the channels of influence, in Section 3.4.1.2 we reintroduce Treatment 2 and describe our econometric approach to isolate compliers in this arm.

3.2.4 Data

We tracked our experimental clients over two and a half loan cycles (on average 176 weeks). Figure 3.1 provides a detailed study timeline. Our analysis makes use of several data sources, which we describe in turn.

Baseline and Endline Data After group formation, we administered a baseline survey to 1016 out of 1028 clients. The short time period between group formation and loan disbursement led to a significant fraction of baseline surveys taking place after loan disbursement. We therefore exclude any potentially endogenous baseline variables from the analysis. Roughly 13 months after first loan disbursement, we conducted an endline survey with 961 clients that provides data on transfers and loan use. We observe similar attrition in both surveys across treatment and control clients (Appendix Table 3.3).

Short-run Social Contact To gauge social interaction among group members, loan officers collected data at repayment meetings during the first loan cycle. The protocol was as follows: After marking passbooks, each client was pulled aside and asked broad questions about social ties with other group members, in order to provide multiple indicators of short-run contact. The first two of these indicators measure social interaction and are constructed as the maximum values of client responses to the two questions – "Have all of your group members visited your house?" and "Have you visited the houses of all group members?" The next two indicators measure knowledge of group members: whether the client knew the names of her group members'

[8] Anthropologists have also documented that group lending increases women's opportunities for social interaction with members of their community [Larance, 2001].

[9] Variation in number of meetings *within* a repayment schedule reflects the fact that VFS allows a client to repay her outstanding balance in a single installment starting 23 weeks after loan disbursement. Once a majority of group members have repaid, remaining clients typically repay at the VFS office.

immediate family and whether she knew if group members had relatives visit over the previous month.[10] Here, we report the average effect size across these measures, defined as the short-run social contact index.[11]

Long-run Social Contact and Lottery Data collection during group meetings allowed us to gather high frequency data in an economical way. However, collecting data in a group setting could create reporting bias that confounds experimental comparisons. For instance, when responses are potentially overheard, a client may be subject to conformity bias wherein she answers questions in a similar manner to others in the group, which could potentially bias experimental estimates. To gather more reliable data on interactions, roughly 16 months after the experimental loan cycle ended, we implemented a lottery game and survey with 866 clients in their homes.[12] Surveying occurred in two phases, and client assignment to phase was random. Section 3.3.2.1 describes the lottery protocol and data. After the lottery was conducted, the client was surveyed about her current contact with every member of her *first* loan cycle group. On average we have nine observations per client. In cases where both members of a pair were surveyed, we keep the maximum value (since social contact cannot vary, in the absence of measurement error or differences in survey timing, within a pair), giving 3,026 pairwise observations.[13] The survey questions included: number of times over the last 30 days the client had visited or been visited by a group member (outside of repayment meetings), whether she talked to the group member about family and whether they celebrated the Bengali festival (Durga Puja) together. We report all three outcomes and, for comparability with the short-run index, also report a long-run social contact index defined at the pair level.

Default Data Our primary outcome of interest is default in the loan cycle *subsequent to* the experimental loan cycle (from now on, second loan cycle), during which all clients reverted to the same repayment and meeting frequency. Bank administrative records show that all clients (except one deceased) took out a loan within 176 weeks of their first loan due date. Appendix Table 3.1 shows that time between due date of first loan and disbursement of second loan does not differ by experimental arm, and we have confirmed that our default results are robust to controlling for this variable.

We define a client as having defaulted if she has not repaid her loan in full by 44 weeks after the official loan end date (i.e., one full loan cycle duration later).[14]

---

[10] To preserve anonymity (given potential observability of responses by group members) we did not ask about interactions with specific group members. We consider the maximum value for all variables, except the relative visit variable for which we take the average (only the latter was reported for an explicit recall period). To account for the delay in starting the survey and the fact that groups could choose to repay early and stop meeting after week 23 of the loan cycle, we use data collected between week 9 and week 23 of the loan cycle.

[11] The index is the equally weighted average of its components' z-scores, where each measure is oriented so that more beneficial outcomes have higher scores. The z-scores are calculated by subtracting the Control group mean and dividing by the Control group standard deviation. By construction, the index has a mean of 0 for the Control group (for further details, see Kling et al., 2007).

[12] We excluded a randomly selected 130 clients with whom we piloted the lottery game and 32 clients could not be tracked.

[13] 82% of pair member provided the same response on having spent the previous Durga Puja together. This is the only long-run social contact variable in which pair-member responses should coincide, absent measurement error (since pair members were not surveyed on the same day).

[14] Although we cannot track all second loan clients for more than 44 weeks, we have verified that second loan default rates

3.2.5 Randomization Balance Check

Panel A in Table 3.1 reports time-invariant characteristics from the baseline survey as a function of treatment assignment. Columns (1)-(3) report the randomization check for the full sample and columns (4)-(6) for clients in the lottery/long-run social interaction survey. On average, randomization created balance between treatment and control groups on observed characteristics. There is one statistically significant difference between Control and Treatment 1 clients: On average, Treatment 1 clients have lived in their neighborhood for two fewer years. With respect to the comparison between Control and Treatment 2, a higher fraction of Muslim clients fell into Treatment 2. This imbalance reflects residential segregation by religion, combined with a relatively small number of Muslim clients: 93% of our clients report living in religiously homogenous neighborhoods (90% Hindu; 3% Muslim). Our 55 Muslim clients are concentrated in eight groups, of which six were assigned to Treatment 2. Since Muslim clients tend to come from larger households, we observe a corresponding imbalance on household size. Since no variable is imbalanced in both treatment arms, the robustness of our results to alternative treatment arms provides strong evidence that imbalances are not driving our results. Nonetheless, throughout this paper we report regressions with and without the controls listed in Panel A of Table 3.1. Appendix Table 3.4 shows that our main results are robust to excluding groups with Muslim clients.[15]

Panel B reports an additional set of variables from the baseline survey that are potentially (though not likely, given the short amount of time between loan disbursement and data collection) influenced by loan receipt. We observe no systematic differences between control and treatment groups. Of the 20 comparisons, the only two (weakly) significant differences in means are that Treatment 1 clients were less likely to have a household member earning a fixed salary, and Treatment 2 clients were slightly less likely to report experiencing an illness during the last 12 months. Finally, comparing across columns we see similar patterns of mean differences in observables across the full sample and the client sample for the lottery/long-run survey.

3.3 Meeting Frequency and Client Relationships

In this section, we use data on social interactions to examine whether requiring first-time VFS clients to meet and repay weekly (Treatment 1) as opposed to monthly (Control) increased social interactions outside of group meetings, both during and beyond the experiment. To investigate whether clients also experienced long-run improvements in risk-sharing arrangements, we implemented a follow-up lottery game that measured willingness to pool risk. For ease of exposition, we restrict the sample to Control and Treatment 1 clients only, since compliance (in terms of meeting protocol) was perfect in these two arms.

are relatively constant at the 64-week mark among those clients whom we can observe for this long. This, combined with the fact that the portfolio at risk statistic officially used for MFI credit rating is defined as the share of portfolio with loan payments outstanding 30 days after due date [CGAP, 2012] makes our default definition relevant.

15 The reduction of groups makes the IV default result more noisily estimated (p-value of 0.12) but the point estimates with and without Muslim groups are of similar size and statistically indistinguishable.

In Section 3.4, we examine the economic impact of these changes by testing whether clients who met weekly in the first loan cycle exhibit lower default on their subsequent loan. Long-run financial behavior (and default) may be directly influenced by initial repayment frequency. We, therefore, complement our experimental analysis by an IV analysis in which we compare default outcomes across clients who paid monthly in the first loan cycle but differ in whether they met on a weekly or monthly basis (that is we compare Treatment 2 to Control). The IV strategy is needed to address noncompliance in the Treatment 2 arm, and exploits the fact that weekly non-repayment meetings were more likely to be canceled if they were scheduled to occur on a day of heavy rainfall. Our IV estimates verify that differences in meeting frequency *not* repayment frequency underlie changes in default.

3.3.1 Impact on Social Interaction

Data obtained during repayment meetings provide a summary measure of a client's interaction with other group members during the experimental loan cycle.

For client i in group g with short-run contact index y_{gi} we estimate:

$$y_{gi} = \beta T_{1,g} + X_{gi}\gamma + \epsilon_{gi} \tag{3.1}$$

where $T_{1,g}$ is an indicator for assignment to the Weekly-Weekly treatment arm (Treatment 1) and X_{gi} represents individual covariates (those variables included in Panel A of Table 3.1). β is interpretable as the effect of switching from a monthly to a weekly group meeting and repayment model on a client's contact with group members outside of meetings. Standard errors are clustered by group.

As reported in Table 3.2, switching a client from monthly to weekly meetings increases her social contact with group members by over 3 standard deviations (column 1). We observe similar results with and without controls (throughout the paper, Panels A and B report estimates without and with controls, respectively).[16] This impact is large but plausible. As the questions ask about a client's social contact with *all* group members, the estimated treatment effect depends on the response to treatment of the weakest pairwise tie within a group. Since 76% of clients have at least one person in their group who is a stranger at baseline and 40% have at least one member who is a *distant* (geographically) stranger at baseline, the estimates are consistent with a scenario in which it takes 5-20 meetings for two strangers to become sufficiently connected to initiate social interaction (hence the index is low for Control groups after five months, but by week 23 virtually every pair of Treatment 1 clients has connected).

However, some caveats apply. First, the presence of other clients during the survey raises the concern of aggregation and reporting biases in client responses. Second, the frequency of surveying may have influenced

[16] Component-wise regression results show large and significant effects of assignment to the Treatment 1 arm. For instance, while only 10% of Control clients report having met all group members outside of meetings, almost 100% of Treatment 1 members report having visited (or having been visited by) all other group members by the same point (see Appendix Table 3.5).

responses and generated artificial differences across treatment groups in reported interactions. A related concern is that surveying clients about social interactions may itself encourage friendship formation. Two pieces of evidence suggest that survey frequency did not directly influence real or reported interactions. First, delays in fieldwork initiation meant that group meeting surveys were implemented more than five weeks after meetings began for 26 of the 68 groups. Data on social interactions from the first group meeting survey for these groups show significant differences across experimental arms in the reported level of interaction. Second, in a later intervention we randomized groups (typically on their third loan cycle) into Weekly-Weekly and Monthly-Monthly groups and loan officers surveyed them during meetings at the same frequency (monthly). We continue to see greater increases in social contact among groups that met weekly. See Appendix Table 3.6 for results.

That said, even in the absence of data quality concerns, our interest is in lasting, not transient, changes in social networks. Therefore, we turn to long-run measures of social interaction, collected 16 months after the experimental loan cycle ended. These data have the additional advantage of being collected through careful surveying, where each client was asked in the privacy of her home about her ongoing interactions with each member of her first loan group. As before, we compare clients assigned to the Weekly-Weekly (Treatment 1) schedule to those assigned to the Monthly-Monthly (Control) schedule. For member i matched with group member m in group g we estimate:

$$y_{gi}^{m} = \beta T_{1,g} + X_{gi}\gamma + s_{gi} + \epsilon_{gi}^{m} \tag{3.2}$$

s_{gi} is a stratification indicator for whether individual i was surveyed in the first phase of surveying. The other variables are defined as in Equation 3.1 and standard errors clustered by group.[17]

Columns (2)-(5) of Table 3.2 reveal that clients engaged in a significant amount of social interaction with their first loan cycle group members at the time of the follow-up survey, and that this interaction was significantly higher among clients who met on a weekly basis during the first loan cycle. In column (2) we see that the average Control pair met 5.5 times over the last 30 days (outside of repayment meetings), and that the average Treatment 1 client pair met 37% more often than their Control counterpart. In total, 15% of Control client pairs versus 22% of Treatment 1 pairs celebrated the last Durga Puja festival together, and 23% of Control client pairs compared to 30% of Treatment 1 pairs report discussing family matters (column 4). Finally, for comparability with the short-run index we report the long-run social contact index, which aggregates outcome variables in columns (2)-(4), and see that Treatment 1 assignment increased long-run

[17] Factors common across observations involving a single member imply observations in a pairwise (dyadic) regression are not independent [Fafchamps and Gubert, 2007]. The error covariance matrix structure may also exhibit correlations varying in magnitude across group members. Group-level clustering of standard errors (which subsumes individual clustering) accounts for this potential pattern: With roughly equal sized clusters, if the covariate of interest is randomly assigned at the cluster level, then only accounting for non-zero covariances at the cluster level, and ignoring correlations between clusters, leads to valid standard errors and confidence intervals [Barrios et al., 2012].

social contact by 0.19 standard deviations.

The persistence of differences in social interaction is particularly striking given that all clients took out at least one additional loan with VFS and roughly half report having a VFS loan outstanding at the time of the follow-up survey. Thus, we might expect social interaction rates to converge as monthly members slowly get to know one another over the long run. However, an important reason *not* to anticipate convergence is churning in group membership: Due to a VFS policy change implemented immediately after our experiment that reduced roup size from ten to five members, the majority (68%) of client pairs were not in the same group for their second loan.[18] Hence, many clients lost the opportunity to get to know one another at group meetings after the experimental loan cycle ended. Put differently, the relatively low level of group membership persistence allows us to more clearly identify differences in meeting frequency during the first loan cycle as the channel for long-run differences in social interaction (which occurred *outside* of meetings).

The policy change raises the possibility that treatment assignment influenced the likelihood that group members remain together in future loan cycles, which could be an independent channel through which average levels of social interaction between treatment groups diverge over time. We are able to track group membership of clients in 51 groups. For these clients, Appendix Table 3.1 shows no difference across experimental arms in the likelihood of being paired with first group members in the second loan cycle. Thus, our experimental differences in long-run contact are likely driven by the higher propensity of Treatment 1 (Weekly-Weekly) clients to stay in touch with members of their first group who did not remain with them for a subsequent loan.

3.3.2 Impact on Risk-sharing

Clearly, the increases in social interaction documented in Table 3.2 are meaningful if they were tangibly welfare-improving, for instance by enabling information spillovers or facilitating economic exchange.[19] For poor clients who face many shocks and rigid debt contracts, informal risk-sharing arrangements are likely to be particularly valuable. Hence, we directly examine whether increasing social interaction facilitated informal risk-sharing arrangements through a series of field-based lottery games. These lotteries, a variant of laboratory dictator and trust games [Forsythe et al., 1994, Berg et al., 1995], were designed to elicit client willingness to form risk-sharing arrangements.

Our methodology contributes to a growing experimental literature on risk-sharing, which finds that increased opportunity for commitment across individuals is associated with a higher willingness to undertake profitable but riskier investments, and that close interpersonal relationships predict risk pooling [Barr and

18 On average, three out of four of a client's second loan group members were from her first loan group, so there is also some degree of change in group membership that is unrelated to the policy change. Anecdotally, the main reason for changes in group membership across cycles is that clients from the same group differed in the timing of their demand for the next loan.

19 Indeed, in and of itself, being encouraged to spend time with strangers may be utility-decreasing if one does so out of convention or social pressure.

Genicot, 2008, Attanasio et al., 2012]. Evidence from games conducted in an experimental economics laboratory also suggests that group contracts improve implicit insurance against investment losses [Gine et al., 2010]. Experimental approaches to measuring risk-sharing, inside or outside of the laboratory, depart considerably from non-experimental empirical tests which most often examine differences in networks' ability to smooth consumption in response to shocks [e.g. Townsend, 1994, Mace, 1991]. While the latter may provide a more direct test of standard hypotheses derived from models of risk-sharing, the experimental approach, in which outcomes are financially incentivized rather than merely reported, arguably enables a more reliable method of establishing risk-sharing between specific pairs of individuals.

That said, we complement our experimental measure of risk-sharing with survey data on financial transfers into and out of client households, and demonstrate similar patterns across the two types of data.[20]

Below, we describe the lottery protocol, and then key predictions of increased risk-sharing for client behavior in the lottery. Then we test these predictions using the lottery data and finally check for consistency of patterns in the financial transfers data.

3.3.2.1 Lottery Protocol and Data

Main Lottery Surveyors approached each client in her house and invited her to enter a promotional lottery for a new VFS retail store. The lottery prize consisted of gift vouchers worth Rs. 200 ($5) redeemable at the store (see Appendix for the surveyor script). The client was informed that, in addition to her, the lottery included 10 clients from different VFS branches, whom she was therefore unlikely to know. If she agreed to enter the draw (all agreed), she was given the opportunity to enter any number of members of her first VFS group into the same draw. Each chosen group member would receive a lottery ticket and be told whom it was from (typically within one day).[21] To clarify how ticket-giving influenced her odds of winning, the client was shown detailed payoff matrices (Figure 3.2), and told that the other ten lottery participants could not add individuals to the lottery. Hence, she could potentially increase the number of lottery participants from 11 to as many as 20, thereby increasing the fraction of group members in the draw from 9% to up to 50% while decreasing her individual probability of winning from 9% to as low as 5%.

We randomized divisibility of the lottery prize at the client level (randomization balance check is provided in Appendix Table 3.7). For half of the sample, the prize was one Rs. 200 voucher, while for the other half it consisted of four Rs. 50 vouchers. Appendix Figure 3.1 provides pictures of these vouchers. A voucher

[20] We lack information on consumption and, therefore, cannot directly link potential improvements in risk-sharing with consumption smoothing (for related work which links risk-sharing and social networks, see Angelucci et al., 2012). Our findings on the comparability of survey and experimental estimates is consistent with Barr and Genicot [2008] and Ligon and Schechter [2012]; both show that behavior of network members is correlated across laboratory and real-world settings.

[21] Only clients who received a ticket were told of the group members' decision. In this sense, the lottery departs from most laboratory trust games in which individuals are not given the opportunity to "opt out" of playing the game. By not giving a ticket, individuals in our sample opt out of participating in the cooperative game with the other member, which is beneficial in a non-anonymous trust game since otherwise behavior could be heavily influenced by social norms rather than pure trustingness.

could only be redeemed by one client and all vouchers expired within two weeks.

Supplementary Lottery Frequent interaction with group members could cause a client to either expand and strengthen her existing social network or to substitute microfinance group members for existing members of her network. To examine the nature of network change, we implemented a supplementary lottery with a sample drawn from five-member VFS groups formed between January and September 2008 (roughly a year and half after the experimental loan groups were formed). As before, groups were randomly assigned to either a weekly or a monthly schedule. For comparability with previous estimates, our lottery was restricted to new (first-time) borrowers, which encompasses 55 Control (Monthly-Monthly) and 51 Treatment 1 (Weekly-Weekly) clients (from 39 and 35 groups respectively). Clients were approached in the same manner as in the original lottery. The difference was that the new lottery asked each client how many tickets she wanted to give to group members (up to four), and how many tickets she wanted to give to individuals outside of the group (up to four). As in the main lottery, if an individual was given a ticket by the client then he or she was informed by the surveyor (typically on the same day). The voucher prize in this lottery was always divisible.

Lottery Data We use data on ticket-giving by a client. For each client in the main lottery, we have, on average, nine pairwise observations on whether she gave a ticket to each of her group members, and for each client in the supplemental lottery, we have eight pairwise observations.

How Artifactual Was the Lottery? Our lottery game shares many design features of the trust game. In using a lottery game in place of a trust game, our primary interest was to avoid triggering client awareness of being a participant in an experiment. Aside from banking, VFS undertakes many community interventions and conducts regular promotional activities in order to attract and retain clients. Thus, it is likely that clients perceived the invitation to participate in a VFS lottery as a natural VFS activity. The potential for the lottery to seem artifactual arises from the invitation to give tickets to other group members. However, the fact that client selection for the lottery was described as a reward for survey participation during her first loan cycle and the fact that the lottery was linked to the VFS store made it more natural that clients were offered the chance to give tickets to their very first loan cycle group members.[22]

3.3.2.2 Testable Predictions

Since group members who receive a ticket from a client are not obligated to share their winnings (as in a trust game), no ticket-giving is a Nash outcome. Risk-pooling via ticket-giving increases a client's expected payoff only if she anticipates that informal enforcement mechanisms will ensure sharing of resources (such as lottery winnings).

To see this, suppose the client gives one group member a ticket. The pair's joint chances of winning the lottery rise from 9% to 17%. There are mutual gains from risk-pooling (e.g., if the pair equally shares

[22] Furthermore, in the supplementary lottery, we expanded the set of people clients could give tickets to and, as described below, our findings are very similar across the two lotteries.

winnings then giving a ticket increases a client's expected lottery winnings from Rs. 18 to 25 and the pair member's expected winnings rise from Rs. 0 to 8.3), but costs to the client if there is no sharing (since her individual probability of winning the lottery declines from 9% to 8% as the pool of lottery entrants rises to 12; see Appendix Figure 3.2 for a graphical illustration).[23]

We use the lottery game to test the hypothesis that higher frequency of interaction can improve a client's ability to enforce risk-pooling arrangements with group members (on this mechanism, also see Karlan et al., 2009, Besley and Coate, 1995, Ambrus et al., 2014). We have already shown that higher meeting frequency in the first loan cycle strengthened long-run social ties between group members. Hence,

Prediction 1 *Higher meeting frequency in the first loan cycle will increase ticket-giving.*

However, a positive correlation between meeting frequency and ticket-giving is also consistent with a model in which more frequent interaction simply increases a client's unconditional altruism towards group members or increases her desire to signal willingness to share.

To isolate the importance of meeting frequency for risk-sharing arrangements we exploit random variation in the divisibility of the lottery prize. Divisibility reduces the transaction costs associated with sharing tickets. In addition, framing the prize as easily divisible may prime the first mover to think of the lottery in terms of potential gains from cooperation as opposed to a purely altruistic effort. However, in both cases a more divisible lottery prize will increase ticket-giving if and only if the client cares about reciprocal transfers.[24] Hence,

Prediction 2 *If ticket-giving only reflects (unconditional) altruism or signaling, then incidence of ticket-giving will be independent of receiver's perceived ability to reciprocate.*

Finally, we consider potential crowd-out of reciprocal arrangements with non-group members. The crowding out force that we consider of interest is the possibility that more time spent with individual group members reduces time spent with people outside the group, given overall time constraints on socializing. The idea is that spending more time with people either encourages or facilitates risk-sharing, so if you spend less time with non-group members, you will be less likely to pool risk with them. To examine whether higher meeting frequency caused clients to substitute social ties with group members for ties with non-group members, we use the supplementary lottery in which a client could choose to give tickets to non-group members. Hence,

Prediction 3 *If ticket-giving to group members is accompanied by substitution away from social ties with non-group members, then ticket-giving to non-group members will be lower for Treatment 1 (Weekly-Weekly)*

23 The top and bottom lines show a client's expected payoff with full and no sharing, respectively. The idea that risk-sharing can increase potential winnings is shared by a trust game, though the increase occurs with certainty in the trust game but stochastically in the lottery game. In addition, unlike a trust game, pairwise returns in the lottery depend on total ticket-giving, generating more subtle predictions on ticket-giving as a function of group composition, which we do not exploit.

24 The behavioral response to the divisibility of the lottery prize could potentially reflect the fact that framing the prize as divisible, and therefore shareable, primes a participant to think in terms of reciprocal arrangements. However, this possibility leaves our prediction unchanged: Divisibility should not matter if motivations for giving are purely altruistic or driven by signaling.

clients than for Control (Monthly-Monthly) clients.

3.3.2.3 Results

Our outcome of interest is ticket-giving: 67.2% of main lottery participants gave at least one ticket. Figure 3.3 shows the ticket distribution across Control and Treatment 1 clients (in percentage terms to account for group size differences) for the main lottery. After zero tickets, the fraction of group members that received tickets declines gradually and levels off after 60%. Control clients are more likely to not give tickets and less likely to give tickets to more than 60% of their group. Ticket-giving patterns in the supplementary lottery are qualitatively similar, with Control clients more likely to not give tickets and less likely to give multiple tickets.

In Table 3.3 we provide regression results from the specifications given by Equation 3.2. Looking across all clients, we see that Treatment 1 clients gave 23.8% more tickets than the Control group (column 1), consistent with stronger social ties among clients who meet weekly translating into higher willingness to risk-share in the lottery game.

Next, we evaluate the importance of risk-sharing relative to either unconditional altruism or a desire to signal reciprocity (independent of willingness to risk-share) in explaining the link between ticket-giving and meeting frequency.

In columns (2) and (3) we show results for clients who were randomized into either the indivisible or divisible prize lottery, respectively. Relative to the Control group, Treatment 1 clients were significantly more likely to give a ticket to a group member if and only if the lottery prize was divisible. Among clients offered the divisible voucher, Treatment 1 clients were 31.9% more likely to give tickets than Control clients (9.1 percentage points). We observe no significant difference between experimental arms when the prize was a single indivisible voucher. Furthermore, for clients in the Control group, ticket-giving behavior was similar across voucher categories.

We have posited that ticket divisibility led to actual or perceived reductions in the transaction costs associated with reciprocal behavior. A first potential explanation for the differential impact of ticket divisibility across experimental arms is non-risk-sharing motivation for ticket-giving among Control clients. Consistent with this, 76% of ticket-giving in the Control group was to either individuals that clients had not seen in the last 30 days, individuals not identified as sources of help in the case of emergency, or immediate family members. A second possibility is that only marginal risk-sharing arrangements were sensitive to the reductions in the transaction costs of reciprocal behavior which were induced by prize divisibility. If there was heterogeneity in the extent to which a client's risk-sharing network was affected by assignment to the weekly group, then the transaction cost reductions may be particularly salient for weekly clients who were less strongly affected by the treatment.

If clients are only able to sustain a fixed number of reciprocal arrangements, then one may worry that stronger ties with group members lead to crowd-out. We use the supplementary lottery to test whether greater risk-pooling among group members was accompanied by substitution away from risk-pooling arrangements with non-group members. For each client we have eight observations, four pertaining to non-group members (we capped ticket-giving to non-group members at four tickets) and four pertaining to group members. We estimate:

$$y_{gi}^m = \beta_1 T_{1,g} + \beta_2 D_{gi}^m + \beta_3 T_{1,g} \times D_{gi}^m + X_{gi}\gamma + \epsilon_{gi}^m \tag{3.3}$$

where y_{gi}^m reflects client i's ticket-giving decision, and D_{gi}^m is an indicator variable for whether individual m is i's group member. We anticipate that β_3 is positive, i.e., ticket-giving is higher among group members of Treatment 1 clients. If there is substitution then β_1 (which captures ticket-giving to non-members) will be *negative.*

Column (4) shows that, consistent with the main lottery, treatment clients are significantly more likely to give tickets to group members in the supplementary lottery ($\beta_3 > 0$). However, β_1 is close to zero and insignificant, suggesting no corresponding decline in the propensity to give tickets to non-group members. Hence, strengthening social ties among group members does not appear to cause clients to substitute away from risk-pooling arrangements with non-group members.[25]

The lack of substitution could reflect several factors: if sharing with non-group members is entirely altruistic, or if the individual time constraint is not binding (so they do not spend less time with people outside the group), or if risk-sharing arrangements with outsiders are not sensitive to small changes in time spent together because they are so well-entrenched, then we would not see any crowd-out.

While we cannot definitively identify which of the above are responsible the absence of a change, qualitative evidence suggests that traditional norms of female isolation rather than time constrains friendship formation in this setting. In interviews, study clients stated that meetings provided them with a reason to leave their home and interact with others in the community. To measure this more systematically, in December 2011 we conducted a detailed time-use survey with 50 women (randomly selected from those who entered the supplementary lottery). The survey collected hourly data over the past 24 hours on what a respondent did and with whom they spent their time. On average, a woman spent 45 minutes per day watching television by herself, 45 minutes per day resting by herself and 26 minutes engaging in other leisure time activities alone. At the end of the survey, each respondent was asked whether she would like to spend more time per week socializing with other women in her community and whether she had the spare time to do so. On average, 86% reported having time to speak with someone who wanted to talk with them, and 66% desired more friends with whom they could spend time.

Finally, we turn to financial transfers data from the endline survey conducted at the end of the first loan cycle. This both provides a consistency check on our risk-sharing interpretation of ticket-giving and tests whether behavior in the potentially artifactual field experiment correlates with behavior outside of the

[25] Since the in-group sharing option was always first (for both treatment and control), it is difficult to interpret differences in levels of in-group versus out-group sharing (β_2), although the interpretation of treatment-control differences (β_1) is still valid.

experiment. Since 43% of clients report no transfers, we focus on a binary outcome of whether the client reported transfers to or from individuals over the last year, grouped into three self-reported categories: (i) close family and friends, (ii) other relatives and neighbors and (iii) other non-relatives.[26] Unfortunately, unlike in the lottery data, we cannot identify transfers to VFS members.

Columns (5) and (7) show that transfers with close family members or friends and "other non-relatives" are equally likely among Treatment 1 and Control clients. However, Treatment 1 clients are 39% more likely to report transfers to other relatives and neighbors (column 6). Thus, consistent with the supplementary lottery ticket-giving results, we see increased risk-sharing and no displacement of risk-sharing arrangements within the immediate family or with other non-relatives.

3.4 Meeting Frequency and Loan Default

Mandating more frequent group meetings during the first loan cycle led to a persistent increase in social interactions and greater risk-pooling by group members. We now examine whether these impacts reduced household vulnerability to economic shocks.

In our setting, a carefully measured indicator of economic vulnerability that is observed for an extended period for all clients is loan default. While default reflects more than vulnerability to shocks, shocks are a strong predictor of default in our data and elsewhere, and informal insurance can be assumed to decrease the likelihood of individual default in the event of a shock [Besley and Coate, 1995, Wydick, 1999].[27]

We focus on default in the second loan cycle.[28] All clients (except one who died) took out a second loan and were placed on an identical fortnightly (every two weeks) repayment schedule for the second loan cycle. Appendix Table 3.1 Panel B reports summary statistics pertaining to clients' second loan cycle, and verifies that they do not vary systematically with treatment status in the first loan cycle. Clients took out a second loan roughly three months after the end of their first loan.[29] The typical second loan was 85% larger than the first, reflecting VFS policy that has clients start well below credit demand and graduate slowly to larger loans. Loan size and timing of disbursement is uncorrelated with first loan repayment schedule. We also have

[26] Close Family/Friend includes the following relationship types: sibling, parent, child, child-in-law, sibling-in-law, parent-in-law, uncle/aunt, cousin, grandchild and friend. Neighbor/Other Relative includes all other relatives and unrelated neighbors. Other Non-Relative includes any other type of acquaintances.

[27] In our data, illness episodes are strong predictors of default, and transfers are associated with lower default risk. Also, home ownership increases default risk, which likely reflects associated illiquidity. Higher levels of savings are negatively correlated with default risk, but corresponding point estimates are noisy (Appendix Table 3.8).

[28] Field and Pande [2008] show that loan delinquency and failure to fully repay loan 16 weeks after the first loan cycle ended do not differ by experimental arms.

[29] Given that we observe no short-run impacts of treatment on client income we do not anticipate that the variation in the timing of second loan demand should be correlated with social capital, and indeed, in Appendix Table 3.1 we do not observe significantly different time periods between first and second loan cycles across treatment and control. Since the number of first-time group members in a client's second loan group is primarily driven by variation in wait times between loans, we also do not anticipate (or observe) any treatment effect on second loan group composition.

second loan use data for a subset of clients, which reveals that most clients used the loan for business-related purposes. This also does not differ by treatment status during first loan cycle.

3.4.1 Results

3.4.1.1 Experimental Estimates: Control versus Treatment 1

Table 3.4 presents regression estimates for default outcomes. Our regression specification parallels Equation 3.1, but the outcome of interest is now an indicator variable Y_{gi} which equals one if client i who belonged to group g in her first loan cycle defaulted on her second loan. We report both Probit and OLS specifications.

As before, we first consider the sample of Control and Treatment 1 clients. In columns (1) and (2) we see that, despite the fact that all individuals faced the same loan terms for their second loan, a client who was previously assigned to a Treatment 1 schedule during her first loan cycle is nearly three times (5.2%) less likely to default on her *second* loan relative to a Control client who was previously assigned to meet on a monthly basis. The difference is strongly significant with or without controls, and is virtually unchanged across Probit and OLS specifications.

3.4.1.2 IV Estimates: Meeting versus Repayment Frequency

By considering default in the subsequent loan cycle, we avoid the possibility that *contemporaneous* differences in repayment frequency influence default outcomes.[30] However, while initial differences in repayment frequency are unlikely to influence differences in social interactions per se, they may change long-run financial habits and, thereby, default.

To isolate the long-run influence of initial differences in meeting frequency from that of repayment frequency, we now examine whether the influence of higher meeting frequency remains when we compare second loan default outcomes across clients who all repaid on a monthly basis in their first loan cycle but differed in whether they met weekly or monthly.

As described in Section 3.2, for the purpose of disentangling these influences, our experimental design included a treatment arm in which clients were required to meet weekly but repay on a monthly basis (Treatment 2). To achieve this, we interspersed the standard monthly group repayment meetings with somewhat artificial weekly "non-repayment group meetings." During non-repayment meetings, loan officers recorded attendance and collected survey data from each individual. In addition, during the first eight meetings, loan officers led a brief (ten-minute) discussion on a topic of common interest, which varied from social concerns, like street safety, to social topics such as recipe exchange.[31]

30 There was no default among Control or Treatment 1 clients during the first loan cycle. This is unsurprising given low loan repayment burden.

31 For ethical reasons, we were requested to provide information useful to clients during non-repayment meetings to justify the cost they were being asked to incur by attending the meetings. We chose topics that we did not expect to directly influence business or social outcomes. Loan officers were provided scripts for each session and required only to read information from the script. Topics covered were: awareness about street safety; geographical knowledge about India; general knowledge

Appendix Figure 3.3 documents the number of meetings held by repayment schedule. Roughly half of the Treatment 2 groups met less frequently than the minimum required by protocol, and thus can be considered non-compliers. According to interviews with loan officers (conducted after the experiment ended when noncompliance was detected), the fact that they did not need to collect and deliver money to VFS after a non-repayment meeting reduced their sense of accountability and made them more inclined to cancel non-repayment meetings (relative to repayment meetings) when inconvenient. Loan officers also acknowledged that meeting cancellations early in the loan cycle caused clients to view the institution of non-repayment meetings as dispensable, making it harder to sustain non-repayment meetings later in the loan cycle. An important reason for early cancellations was monsoon rains which caused waterlogging of neighborhoods and roads, increasing both loan officer and client commute time (50% of our loan groups were formed during monsoon months; on the impact of monsoon rains on daily life in Kolkata also see Beaman and Magruder, 2012).[32]

To address imperfect compliance in Treatment 2, we use an IV specification that makes use of this exogenous variation in monsoon rainfall shocks early in the loan cycle in order to predict Treatment 2 groups that met at least 23 times. This is the minimum number of times required by protocol, and also happens to be the median meeting rate for Treatment 2 groups. Our analysis sample for the IV estimates includes only Control and Treatment 2 clients, all of whom repaid monthly. If, among clients who repaid monthly, those who met weekly exhibit lower default incidence, then we will have identified an independent role for meeting frequency.[33] The first stage of our IV regression is:

$$M_{gi}^{23+} = \beta_1 T_{2,g} + \beta_2 Heavy_g + \beta_3 T_{2,g} \times Heavy_g + X_{gi}\gamma + \epsilon_{gi} \tag{3.4}$$

M_{gi}^{23+}, now on *group met weekly*, is an indicator variable which equals one if individual i belonged to a group g which met at least 23 times during the loan cycle.[34] M_{gi}^{23+} equals 0 for all Control groups (since there was perfect compliance in this arm). $T_{2,g}$ is an indicator variable for assignment to Treatment 2 (Weekly-Monthly). $Heavy_g$ is the number of heavy rainfall days (defined as days with rainfall above the 90th percentile of rainfall distribution for the city) during the first four weeks of meetings.[35] While it is possible that rainfall has a direct effect on social or economic outcomes, it is unlikely that rainfall shocks over such a short time period directly influence long-run social interactions and/or economic activity and,

about family ancestry; recipe exchange; questions on how they spend vacations or holidays; information on bus routes in their neighborhoods; basic physiology; basic information on state politics.

32 A VFS loan officer's average work day lasts 12 hours, and consists of conducting group meetings in the morning and then returning to the branch office by early afternoon to deposit the repayments that had been collected and complete paperwork. On an average day, a loan officer would conduct five to six group meetings and cover a distance of 20 kms on bicycle.

33 Importantly, weekly-monthly clients did not take oaths during non-repayment meetings, so we can rule out the possibility that frequency of oath-taking influenced repayment behavior for these clients.

34 We define a meeting as having occurred if at least two group members attended.

35 This corresponds to days 29–56 after group formation.

therefore, client ability to repay in the subsequent loan cycle. Hence, our exclusion restriction is likely to be satisfied. Furthermore, we have confirmed with baseline data that an additional day of heavy rain over the seven days before a client is surveyed does not influence a household's wage income or likelihood of employment (see Appendix Table 3.9).[36]

Column (3) of Table 3.4 reports this first stage regression. Treatment 2 clients at the mean value of days of heavy rain (5.1) were 47% less likely to meet the minimum required number of times than those who experienced zero days of heavy rain 29–56 days after group formation. Thus heavy rainfall very significantly influenced the sustainability of non-repayment meetings over the loan cycle.

Given this first stage, we turn to the IV estimate of the impact of increased meeting frequency, holding constant repayment frequency. Our structural equation of interest (i.e., second stage) is:

$$y_{gi} = \beta M_{gi}^{23+} + X_{gi}\gamma + \epsilon_{gi} \tag{3.5}$$

Column (4) reveals a negative and significant impact of higher meeting frequency in first loan cycle on default for the second loan.[37] The coefficient estimate is similar in magnitude (even slightly larger) to the experimental estimates in columns (1) and (2). Thus, we can rule out the possibility that lower long-run default rates among clients assigned to a weekly meeting schedule reflect improvements in their financial habits or business practices associated with having repaid their first loan on a weekly basis.

We conclude that higher meeting frequency in first loan cycle underlies the subsequent default reduction and, based on our results, posit that increased social interactions among group members is the primary channel of influence. The potential mechanisms through which social interactions influenced default potentially include better ability to monitor (and punish) group members, lower transaction costs for sharing and improved information flows across members.

3.5 Conclusions

A widely held belief among social scientists across many disciplines is that social interactions encourage norms of reciprocity and trust, which deliver economic returns. In fact, participation in groups is often used to measure individuals' or communities' degree of economic cooperation (see, for instance, Narayan and Pritchett, 1999). While theoretically well-grounded, it is not clear from previous work whether the correlation between social distance and trust reflects the causal effect of interaction on economic cooperation.

We provide experimental evidence that a development program that encourages repeat interactions can increase long-run social ties and enhance social capital among members of a highly localized community in

[36] Our results are also robust to extending the definition of Heavy Rain to include the first eight weeks of the loan cycle, or to using the 80th or 85th percentile of the rainfall distribution as the cutoff (results are shown in Appendix Table 3.9).

[37] We employ a linear IV specification given the strong functional form assumptions associated with the biprobit model [Angrist and Pischke, 2009].

a strikingly short amount of time. With only the outside stimulus of MFI meetings, close neighbors from similar socioeconomic backgrounds got to know each other well enough to cooperate in an economically meaningful way, which provided a buffer against economic shocks that lead to default. While many studies have suggested a link between social capital and MFI default rates, ours is the first to provide rigorous evidence on the role of microfinance in building social capital, and thereby broaden our understanding of the channels through which MFIs achieve low default rates without the use of physical collateral. Arguably, the improvements in risk-sharing we observe are even more striking because they were obtained in the absence of joint-liability contracts, and provide a rationale for the current trend among MFIs of maintaining repayment in group meetings despite the transition from joint- to individual-liability contracts [Gine and Karlan, 2009]. While it is difficult to account for *all* of the increased transaction costs of weekly meetings with higher loan recovery rates alone, direct cost savings from lowering default go a long way towards explaining why weekly meetings persist as the standard MFI practice.[38] Furthermore, there are many reasons to believe that the typical MFI is sufficiently delinquency- and/or default-averse to make weekly meetings cost effective.[39]

Using meetings to improve risk-sharing in a setting characterized by weak formal institutions for contract enforcement is a potentially important source of welfare gains, at least for first-time clients. Although encouraging social interaction entails higher participation costs for clients, the benefits from social network expansion are likely to outweigh the cost. We estimate that weekly compared to monthly meetings entail approximately 15 additional hours of client time over the course of an average loan cycle.[40] The benefits are likely to include, in addition to lower default risk, utility gains from consumption smoothing and other positive externalities from social interaction such as information sharing. In addition, we anticipate that lower propensity to default improves clients' long-term financial access (both ability to take out future loans and loan amount).

By broadening and strengthening social networks, the group-based lending model used by MFIs may provide a valuable vehicle for the economic development of poor communities and the empowerment of women. While we cannot expect all communities to respond equally to such stimuli, our findings are likely to be most readily applicable to the fast-growing urban and peri-urban areas of cities in developing countries (such as Kolkata) where microfinance is spreading most rapidly. An important goal of future research would

38 We estimate an additional average cost per client of Rs. 85 for a weekly relative to a monthly meeting schedule. Loan officers spend additional three hours per month per group (one hour in meeting time and two hours in commute time), which amounts to 1.9% of their monthly wage for the average group (of ten clients), or Rs. 85. Meanwhile, our data indicate that the average client who met and repaid monthly during her initial loan cycle defaulted on only Rs. 30 more than one previously on a weekly repayment schedule.

39 For instance, even delinquency reduces MFI liquidity and ability to expand lending, and MFI credit ratings are typically calculated based on share of MFI portfolio in arrears.

40 The estimate of two additional hours per month is based on meeting length of 20 minutes and an average client commute of ten minutes to and from meeting. As a client repays her loan, on average, after 7.5 months this adds up to 15 hours over the course of a loan cycle. While client cost is likely higher than just this time cost of meeting attendance, including, for instance financial and psychic burden of making regular repayments Field et al. [2012], these costs are likely less important to first time clients who receive very small loans.

be to understand how other development programs and public policies can be designed to enhance the social infrastructure of poor communities.

Table 3.1. Randomization Check

	All Clients			Lottery/Long-Run Survey Clients		
	Control Mean (Monthly-Monthly)	Treatment 1 (Weekly-Weekly)	Treatment 2 (Weekly-Monthly)	Control Mean (Monthly-Monthly)	Treatment 1 (Weekly-Weekly)	Treatment 2 (Weekly-Monthly)
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A						
Age	33.969	-0.593	-1.110	33.832	-0.806	-0.920
	(8.553)	(0.813)	(0.724)	(8.418)	(0.810)	(0.764)
Literate	0.865	-0.012	-0.059	0.880	-0.012	-0.059
	(0.342)	(0.035)	(0.039)	(0.325)	(0.036)	(0.040)
Married	0.862	0.013	0.005	0.871	0.025	-0.009
	(0.345)	(0.031)	(0.030)	(0.336)	(0.030)	(0.029)
Household Size	3.821	0.153	0.207*	3.903	0.068	0.106
	(1.335)	(0.106)	(0.114)	(1.357)	(0.119)	(0.124)
Muslim	0.023	-0.023	0.118**	0.026	-0.026	0.122*
	(0.151)	(0.021)	(0.060)	(0.159)	(0.023)	(0.062)
Years Living in	17.423	-2.010**	-0.931	17.136	-2.175**	-0.456
Neighborhood	(10.473)	(0.889)	(0.919)	(10.407)	(0.903)	(0.976)
Number of Clients	10.364	-0.086	-0.037	10.385	-0.073	-0.054
in Group	(0.727)	(0.185)	(0.192)	(0.741)	(0.199)	(0.196)
Group Formed	0.595	-0.147	-0.109	0.654	-0.154	-0.159
in Rainy Season	(0.492)	(0.122)	(0.120)	(0.477)	(0.124)	(0.119)
Heavy Rain Days	5.265	-0.128	-0.477	5.453	-0.205	-0.614
	(2.070)	(0.545)	(0.519)	(2.060)	(0.576)	(0.534)
Panel B						
Client Worked for	0.525	0.060	0.011	0.524	0.056	0.018
Pay in Last 7 Days	(0.500)	(0.053)	(0.053)	(0.500)	(0.053)	(0.053)
Household Earns	0.442	-0.079*	0.023	0.437	-0.065	0.048
Fixed Salary	(0.497)	(0.044)	(0.049)	(0.497)	(0.046)	(0.050)
Household Owns	0.717	0.038	-0.080	0.718	0.034	-0.085
Business	(0.451)	(0.049)	(0.061)	(0.450)	(0.053)	(0.061)
Household Savings	1636.2	325.7	1238.9	1828.7	103.3	1125.2
	(5793.7)	(564.8)	(762.9)	(6405.5)	(653.7)	(840.5)
Household Owns Home	0.808	-0.033	-0.035	0.828	-0.048	-0.047
	(0.395)	(0.044)	(0.047)	(0.378)	(0.046)	(0.048)
Education Expenditures	4183.9	559.5	-278.2	4490.2	112.0	-598.2
	(4868.2)	(407.8)	(356.3)	(4919.3)	(456.7)	(392.9)
Health Expenditures	3311.4	-35.0	-399.4	3241.4	-87.7	-226.9
	(5262.1)	(522.2)	(432.4)	(5154.4)	(562.9)	(432.1)
Illness in Past 12 Months	0.314	0.029	-0.080*	0.307	0.016	-0.062
	(0.465)	(0.048)	(0.046)	(0.462)	(0.053)	(0.049)
Number of Transfers into	1.388	0.172	-0.503	1.085	0.205	-0.185
Households	(6.796)	(0.542)	(0.449)	(4.659)	(0.362)	(0.335)
Number of Transfers out of	2.613	0.282	-0.253	2.563	0.311	-0.147
Households	(4.693)	(0.604)	(0.558)	(4.728)	(0.658)	(0.592)
Days between Loan				788.312	-0.211	13.977
Disbursement and Lottery				(46.182)	(11.360)	(10.968)
N	385	306	325	309	250	297

Notes

1 Group Formed in Rainy Season is an indicator variable for whether the group was formed in June, July, August, or September. Heavy Rain Days is a count variable representing the number of days within 29-56 days (5-8 weeks) after group formation in which rain was above the 90th decile for daily rainfall (14.3 mm). Illness in Past 12 Months is an indicator variable for whether any household member has been ill in past 12 months.

2 Columns (2)-(3) are the regression results of the characteristics in the title column on the two treatments for the full sample. The omitted group is clients in Control groups. In columns (5)-(6) we report the same coefficients for the sample that received the lottery. All lottery sample regressions control for survey phase. * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered by group.

Table 3.2. Meeting Frequency and Social Interactions in the Short Run and Long Run

	Short Run	Long Run			
	Social Contact Index	Total Times Met	Attend Durga Puja	Talk Family	Social Contact Index
	(1)	(2)	(3)	(4)	(5)
Panel A: No Controls					
Treatment 1	3.005***	2.045**	0.069*	0.070*	0.186**
(Weekly-Weekly)	(0.107)	(1.001)	(0.038)	(0.039)	(0.080)
Panel B: Controls Included					
Treatment 1	3.052***	2.054**	0.081**	0.071**	0.199***
(Weekly-Weekly)	(0.092)	(0.891)	(0.039)	(0.035)	(0.073)
Control Mean		5.475	0.153	0.229	
(Monthly-Monthly)		[10.386]	[0.360]	[0.421]	
Specification	OLS	OLS	Probit	Probit	OLS
N	684	3026	3023	3026	3026

Notes

1 Short-Run Social Contact Index generates average effect size from four client questions: (1) "Have you ever visited houses of all group members?" (2) "Have all of your group members visited your house?" (3) "Do you know the names of the family members of your group members?" and (4) "Do you know if any of your group members had relatives come over in the last 30 days?" The first three variables equal one if client responds yes at least once between week 9 and week 23 of her loan cycle, and the fourth is the mean value of client responses over this period. Long-Run Social Contact Index generates average effect size from three questions asked to each client during the lottery survey: (1) Total Times Met, (2) "Do you still talk to X about her family?" and (3) "During the most recent Durga Puja, did you attend any part of the festival with X?"

2 The sample is clients assigned to Treatment 1 (Weekly-Weekly) and Control (Monthly-Monthly) groups.

3 Regressions with controls include the variables in Table 3.1, Panel A. All long-run regressions also control for survey phase. * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered by group.

Table 3.3. Meeting Frequency and Risk-Sharing: Ticket-Giving and Transfers

	Main Lottery			Supplementary Lottery	Transfers		
	Gave Ticket						
	All	1-Rs. 200 Voucher	4-Rs. 50 Vouchers	All	Close Family/ Friend	Neighbor/ Other Relative	Other Non-Relative
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: No Controls							
Treatment 1	0.067**	0.043	0.091*	-0.005	0.016	0.122**	-0.019
(Weekly-Weekly)	(0.034)	(0.041)	(0.048)	(0.069)	(0.065)	(0.061)	(0.028)
Group Member				0.068**			
				(0.034)			
Treatment 1*Group Member				0.157**			
				(0.079)			
Panel B: Controls Included							
Treatment 1	0.072**	0.044	0.105**	0.0001	0.019	0.126**	-0.011
(Weekly-Weekly)	(0.033)	(0.039)	(0.048)	(0.071)	(0.066)	(0.058)	(0.024)
Group Member				0.073**			
				(0.036)			
Treatment 1*Group Member				0.158*			
				(0.081)			
Control Mean	0.281	0.277	0.285	0.223	0.426	0.309	0.067
(Monthly-Monthly)	[0.450]	[0.448]	[0.452]	[0.417]	[0.495]	[0.463]	[0.250]
Specification	Probit	Probit	Probit	Probit	Probit	Probit	Probit
N	5282	2695	2587	847	651	651	651

Notes

1 For the lottery, the dependent variable equals one for a group member if the client gave her a ticket. For each client in the sample we have (on average) nine observations for columns (1)-(3). In column (4), we include only clients borrowing for the first time during the Third Loan Cycle (see Figure 1 for details). For this column, we have eight observations for each client (four for group member ticket-giving and four for non-group member ticket-giving). In columns (5)-(7), Transfers are indicator variables for whether client's household gave or received any transfers to or from the relevant groups in the 12 months before the first loan endline survey. We divide transfers into three categories based on client's stated relationship with transfer recipient/sender at time of survey. Close Family/Friend includes the following relationship types: sibling, parent, child, child-in-law, sibling-in-law, parent-in-law, uncle/aunt, cousin, grandchild, and friend. Neighbor/Other Relative includes all other relatives and unrelated neighbors. Other Non-Relative includes any other type of acquaintances.

2 The sample is clients assigned to Treatment 1 (Weekly-Weekly) and Control (Monthly-Monthly) groups.

3 Regressions with controls include the variables in Table 3.1, Panel A. All lottery regressions also control for survey phase. * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered by group.

Table 3.4. Meeting Frequency and Default: Evidence from the Second Loan Cycle

	Default		Group Met Weekly	Default
	(1)	(2)	(3)	(4)
Panel A: No Controls				
Treatment 1	-0.052**	-0.052**		
(Weekly-Weekly)	(0.021)	(0.021)		
Treatment 2 (Weekly-Monthly)*Heavy Rain Days			-0.118***	
			(0.020)	
Treatment 2			1.086***	
(Weekly-Monthly)			(0.152)	
Heavy Rain Days			0.025	
			(0.016)	
Group Met Weekly				-0.077**
				(0.038)
Panel B: Controls Included				
Treatment 1	-0.036**	-0.045**		
(Weekly-Weekly)	(0.016)	(0.021)		
Treatment 2 (Weekly-Monthly)*Heavy Rain Days			-0.125***	
			(0.020)	
Treatment 2			1.087***	
(Weekly-Monthly)			(0.146)	
Heavy Rain Days			0.024	
			(0.018)	
Group Met Weekly				-0.092**
				(0.042)
F Statistic			*20.16*	
p-value			*[0.000]*	
Control Mean (Monthly-Monthly)	0.072 [0.258]			
Specification	Probit	OLS	OLS	Linear IV
N	698	698	720	720

Notes

1 A client is defined as having defaulted if she has not repaid the total loan amount within 44 weeks after due date. Group Met Weekly is an indicator variable for whether a group met at least 23 times during First Loan Cycle. Heavy Rain Days is as defined in Table 3.1.

2 Column (3) provides the first stage regression for the IV regression in column (4).

3 Columns (1)-(2) include clients assigned to Treatment 1 (Weekly-Weekly) and Control (Monthly-Monthly) groups, and columns (3)-(4) include clients assigned to Treatment 2 (Weekly-Monthly) and Control (Monthly-Monthly) groups.

4 Panel A regressions in columns (3)-(4) include a control for Group Formed in Rainy Season, and regressions with controls (Panel B) include the variables in Table 3.1, Panel A and also control for second loan size. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered by group.

Figure 3.1. Timeline

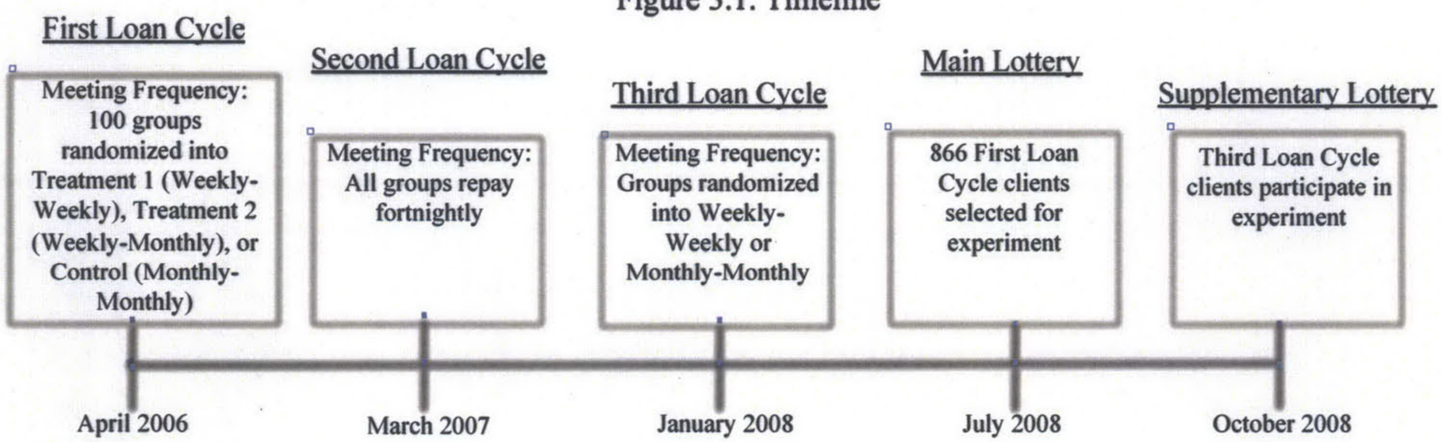

Notes: Dates reflect the start of each loan cycle and of lottery surveying. Our sample population consisted of 1028 clients who joined VFS in 2006. For their first loan cycle 392 of these clients were randomly assigned to monthly meeting and monthly repayment (38 Control groups), 307 were assigned to weekly meeting and weekly repayment (30 Treatment 1 groups), and 329 were assigned to weekly meeting and monthly repayment (32 Treatment 2 groups). All but one client continued to a second loan cycle during which all clients met for repayment on a fortnightly basis. We use this sample to evaluate second loan cycle default outcomes. Finally, clients in the third loan cycle were randomized into Weekly-Weekly or Monthly-Monthly groups. To examine the effects of meeting frequency on giving to non-group members, we restrict our sample to clients who were borrowing for the first time in the third loan cycle and who were in groups with at least one returning borrower. There are 106 such clients.

Figure 3.2. Winning Probabilities

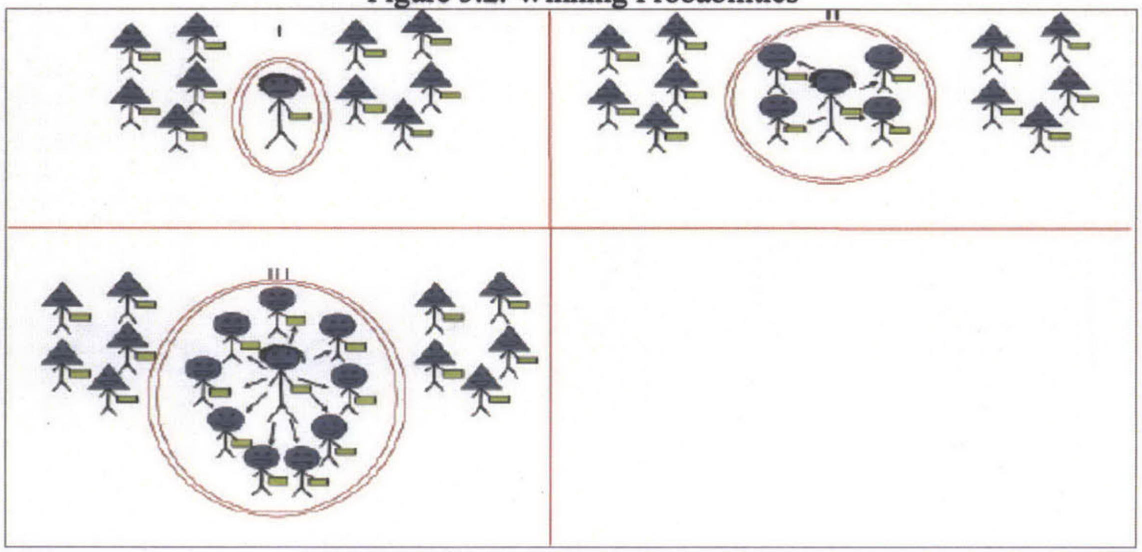

Notes:

This picture was used to explain how ticket-giving affected lottery probabilities. The explanation provided was "In Picture 1 in which you don't give out any tickets to members of your VFS group, you have a 1 in 11 chance of winning.

In Picture 2, you choose to have us give a ticket to four other members of your VFS group and there are 15 tickets total. In that case, you would have a 1 in 15 chance of winning and each of the members of your VFS group you gave a ticket to would have a 1 in 15 chance of winning.

In Picture 3, you choose to have us give a ticket to nine other members of your VFS group and there are 20 tickets total. In that case, you would have a 1 in 20 chance of winning and each of the members of your VFS group you gave a ticket to would have a 1 in 20 chance of winning." In each picture, those outside of the red circle are non-group members.

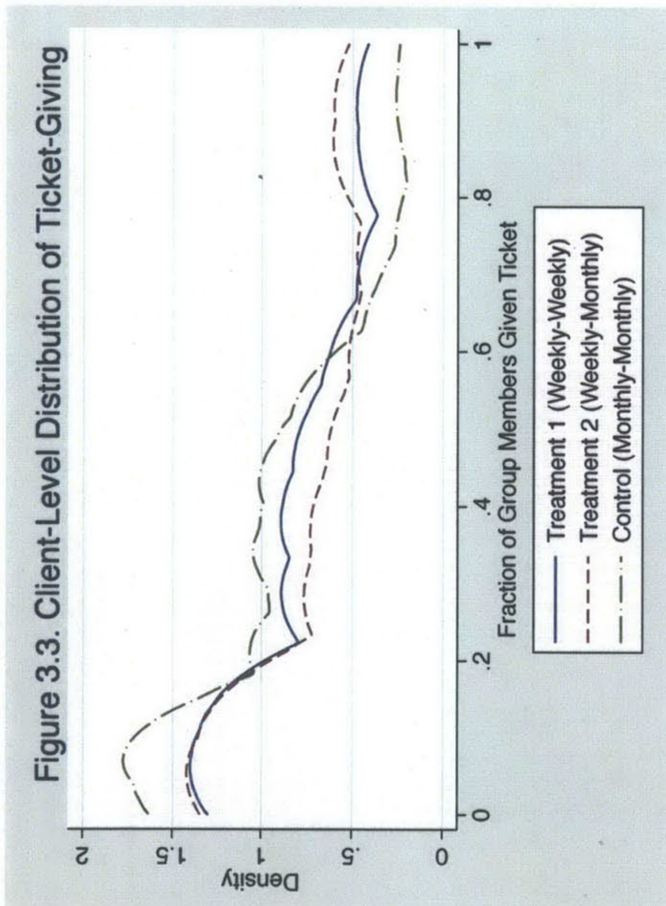

Figure 3.3. Client-Level Distribution of Ticket-Giving

Appendix Table 3.1. Robustness Checks: Impact of Meeting Frequency on Additional Outcomes

	Treatment 1 (Weekly-Weekly) **Explanatory Variable**					
	No Controls	Controls Included	Control Mean (Monthly-Monthly)	N	Specification	Data Source
Dependent Variable	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: First Loan Cycle						
Late Repayment	-0.0004	-0.0005	0.0008	698	OLS	Group Meeting Survey
	(0.0006)	(0.0006)	[0.0092]			
Loan Officer Rank	-0.057	-0.052	2.674	675	OLS	Group Meeting Survey
	(0.054)	(0.051)	[0.310]			
Present	0.005	0.004	0.754	698	OLS	Group Meeting Survey
	(0.019)	(0.014)	[0.230]			
Late	0.228***	0.210***	0.178	697	OLS	Group Meeting Survey
	(0.062)	(0.056)	[0.258]			
Meeting Duration	-0.021	-0.022*	0.172	675	OLS	Group Meeting Survey
	(0.013)	(0.013)	[0.081]			
Household Member Attending School	0.002	-0.025	0.595	460	Probit	First Loan Cycle Endline Survey
	(0.046)	(0.048)	[0.492]			
Total New Savings	120.6	88.9	128.9	446	OLS	First Loan Cycle Endline Survey
	(158.1)	(154.2)	[1354.5]			
Expanded Business in Past 30 Days	0.0008	-0.004	0.033	651	Probit	Endline Survey + Follow-up Survey
	(0.020)	(0.019)	[0.180]			
Panel B: Second Loan Cycle						
Days to Second Loan Takeup	-28.7	-27.9	116.4	698	OLS	Default Data
	(21.8)	(20.6)	[145.6]			
Second Loan Size	-30.7	-11.4	7411.8	698	OLS	Default Data
	(144.0)	(136.3)	[901.1]			
Fraction Group Members in Second Loan Group	0.028	0.019	0.718	324	OLS	Administrative Data
	(0.064)	(0.062)	[0.319]			
Loan Used for Raw Materials	-0.033	-0.028	0.210	324	Probit	Second Loan Cycle Endline Survey
	(0.045)	(0.042)	[0.408]			
Loan Used for Business Equipment	-0.020	-0.009	0.270	324	Probit	Second Loan Cycle Endline Survey
	(0.062)	(0.058)	[0.445]			
Loan Used for Health Care Costs	0.056	0.066	0.065	324	Probit	Second Loan Cycle Endline Survey
	(0.054)	(0.052)	[0.247]			
Loan Used for Housing	0.052	0.054*	0.045	324	Probit	Second Loan Cycle Endline Survey
	(0.034)	(0.032)	[0.208]			

Notes

1 Late Repayment is the fraction of group meetings at which a client failed to make the scheduled repayment. Loan Officer Rank is measured on a four-point scale, with higher rankings reflecting a higher perceived ability to repay. Present and Late are averages taken for group meetings between week 9 and week 23 of the loan cycle. Meeting Duration is measured in hours and is averaged across all group meetings. Days to Second Loan Takeup is defined as the number of days between scheduled First Loan repayment and Second Loan takeup. Fraction Group Members in Second Loan Group is defined as the fraction of second loan group members also in first loan group. Loan Used for _______ are indicator variables and multiple loan uses may be listed for each loan in Second Loan Cycle.

2 For First Loan Cycle questions on savings and school attendance, the sample excludes First Loan Cycle clients who received the follow-up survey (which did not ask about these topics and was administered to clients who repaid their initial loans faster than anticipated). For Second Loan Cycle loan use questions, the sample includes only First Loan Cycle clients who remained research clients during the Second Loan Cycle (and so continued to be surveyed regarding loan use).

3 The sample is clients assigned to Treatment 1 (Weekly-Weekly) and Control (Monthly-Monthly) groups.

4 Regressions with controls include the variables in Table 3.1, Panel A (except for Days to Second Loan Takeup specification which includes all controls except for Group Formed in Rainy Season). * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered by group.

Appendix Table 3.2: Representativeness of VFS Borrowers

	VFS	SEWA	Spandana
	(1)	(2)	(3)
Income in Last Month	4845.4	6164.7	3323.0
	(2705.4)	(4140.9)	(3838.4)
Household Owns Business	0.703	0.369	0.494
	(0.457)	(0.483)	(0.500)
Number of Paid Employees	0.232	0.699	0.287
	(1.206)	(3.120)	(1.094)
Profit Last Month (Rs.)	2846.8	2668.7	-
	(2339.5)	(2420.4)	-
Number of Loans in Past Year	1.030	1.443	4.410
	(0.175)	(0.832)	(2.704)
Largest Loan (Rs.)	4178.9	23309.2	36792.6
	(2165.0)	(28941.5)	(44452.4)
Fraction Households with Savings	0.263	-	0.722
	(0.440)	-	(0.448)
Household Owns Home	0.786	0.710	0.788
	(0.410)	(0.454)	(0.409)
Number of Rooms in Home	1.757	1.822	2.221
	(1.087)	(0.919)	(1.143)
Household Owns TV	0.783	0.882	0.859
	(0.412)	(0.323)	(0.348)
Household Owns Two-wheeler	0.046	0.293	0.275
	(0.210)	(0.456)	(0.447)
Number of Household Members	3.933	5.693	5.902
	(1.336)	(2.262)	(2.248)
Has Insurance	0.352	-	0.500
	(0.478)	-	(0.500)
N	1016	846	1599

Notes

1 Number of paid employees is defined only for business owners. Profit last month is defined as the average of minimum and maximum monthly profits for VFS borrowers.

2 Each column presents the mean and standard deviation for the relevant sample and the given outcome variable.

3 VFS data comes from the 2006 First Loan Cycle baseline survey. SEWA data comes from a 2009-2010 survey of SEWA clients conducted (and made available) by Field and Pande. Spandana data comes from a 2007-2008 endline survey conducted (and made available) by Banerjee and co-authors, and is restricted to respondents who have an outstanding MFI loan.

Appendix Table 3.3. Survey Attrition

	Received Baseline Survey	Received Endline Survey
	(1)	(2)
Panel A: No Controls		
Treatment 1	0.015*	0.035
(Weekly-Weekly)	(0.008)	(0.023)
Treatment 2	0.006	0.026
(Weekly-Monthly)	(0.010)	(0.021)
Control Mean	0.982	0.916
(Monthly-Monthly)	[0.133]	[0.278]
Specification	OLS	OLS
N	1028	1028

Notes

1 * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered by group.

Appendix Table 3.4. Main Results, excluding Muslim clients

	Short Run	Long Run	Lottery		Default	
	Social Contact Index	Social Contact Index	All	4-Rs. 50 Vouchers	Default	Default
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: No Controls						
Treatment 1	3.373***	0.159**	0.063*	0.082*	-0.038*	
(Weekly-Weekly)	(0.104)	(0.079)	(0.035)	(0.049)	(0.020)	
Group Met Weekly						-0.072
						(0.045)
Panel B: Controls Included						
Treatment 1	3.385***	0.182**	0.071**	0.102**	-0.029**	
(Weekly-Weekly)	(0.098)	(0.073)	(0.033)	(0.049)	(0.014)	
Group Met Weekly						-0.078
						(0.047)
Specification	OLS	OLS	Probit	Probit	Probit	Linear IV
N	667	2905	5061	2452	675	635

Notes

1 Column (1) replicates column (1) of Table 3.2, column (2) replicates column (5) of Table 3.2, column (3) replicates column (1) of Table 3.3, column (4) replicates column (3) of Table 3.3, column (5) replicates column (1) of Table 3.4, and column (6) replicates column (4) of Table 3.4.

Appendix Table 3.5. Short Run Social Contact Index Components

	Short Run			
	I Visited All Members in Their Homes	All Members Visited Me in My Home	Know Names of Family Members	Know if Relatives Visited
	(1)	(2)	(3)	(4)
Panel A: No Controls				
Treatment 1	0.920***	0.928***	0.931***	0.127***
(Weekly-Weekly)	(0.030)	(0.030)	(0.030)	(0.025)
Panel B: Controls Included				
Treatment 1	0.941***	0.948***	0.951***	0.123***
(Weekly-Weekly)	(0.024)	(0.023)	(0.024)	(0.024)
Control Mean	0.080	0.072	0.069	0.015
(Monthly-Monthly)	[0.271]	[0.258]	[0.254]	[0.094]
Specification	OLS	OLS	OLS	OLS
N	684	684	684	684

Notes

1 Dependent variables in columns (1)-(4) are constructed respectively from client indicator variables which equal one if the client responded "Yes" to the questions, (1) "Have you ever visited houses of all group members?" (2) "Have all of your group members visited your house?" (3) "Do you know the names of the family members of your group members?" and (4) "Do you know if any of your group members had relatives come over in the last 30 days?" The first three variables equal one if client responds yes at least once between week 9 and week 23 of her loan cycle, and the fourth is the mean value of client responses over this period.

2 The sample is clients assigned to Treatment 1 (Weekly-Weekly) and Control (Monthly-Monthly) groups.

3 Regressions with controls include the variables in Table 3.1, Panel A. * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered by group.

Appendix Table 3.6. Short Run Social Contact Robustness Checks

	Short Run (Initial Meeting for Delayed Groups)		Third Intervention Clients			
	Know Names of Family Members	Know if Relatives Visited	Number of Members I Visited	Number of Members Visited Me	Number of Group Members Discuss Business	Number of Group Members Discuss Personal Matters
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: No Controls						
Treatment 1	0.913***	0.043	0.555***	0.506***	0.726***	0.534***
(Weekly-Weekly)	(0.056)	(0.039)	(0.106)	(0.183)	(0.128)	(0.154)
Panel B: Controls Included						
Treatment 1	0.918***	0.041	0.555***	0.525***	0.719***	0.531***
(Weekly-Weekly)	(0.050)	(0.036)	(0.107)	(0.170)	(0.124)	(0.147)
Control Mean	0	0	0.472	1.282	0.740	1.125
(Monthly-Monthly)	.	.	[0.562]	[1.179]	[0.758]	[1.035]
Specification	OLS	OLS	OLS	OLS	OLS	OLS
N	221	221	707	707	707	707

Notes

1 Dependent variables in columns (1)-(2) are constructed from client indicator variables which equal one if the client responded "Yes" to the questions, (1) "Do you know the names of the family members of your group members?" and (2) "Do you know if any of your group members had relatives come over in the last 30 days?" Dependent variables in columns (3)-(6) are constructed from client responses to the questions, (3) "How many group members have you visited in their houses in the last 2 weeks?" (4) "How many group members have visited you in your house in the last 2 weeks?" (5) "How many people in the group did you talk to about business matters in the last 2 weeks?" and (6) "How many people in the group did you talk to about personal matters in the last 2 weeks?" All four variables in columns (3)-(6) are averaged over the first five months of the loan cycle.

2 We observe perfect difference in means (all Weekly-Weekly clients responded "Yes" and all Monthly-Monthly clients responded "No") to the following two short-run questions: (1) "Have you ever visited houses of all group members?", and (2) "Have all of your group members visited your house?"

3 The sample in columns (1)-(2) is clients assigned to Treatment 1 (Weekly-Weekly) and Control (Monthly-Monthly) groups that did not receive group meeting surveys until more than five weeks after meetings began. The sample in columns (3)-(6) is clients in the Third Loan Cycle.

4 Regressions with controls include the variables in Table 3.1, Panel A. * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered by group.

Appendix Table 3.7. Lottery Randomization Check

	4-Rs. 50 Vouchers			1-Rs. 200 Voucher		
	Control Mean (Monthly-Monthly)	Treatment 1 (Weekly-Weekly)	Treatment 2 (Weekly-Monthly)	Control Mean (Monthly-Monthly)	Treatment 1 (Weekly-Weekly)	Treatment 2 (Weekly-Monthly)
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A						
Age	33.919	-1.353	-1.037	33.750	-0.214	-0.709
	(8.275)	(1.023)	(0.866)	(8.574)	(1.191)	(1.090)
Literate	0.866	-0.010	-0.050	0.894	-0.015	-0.070
	(0.342)	(0.046)	(0.053)	(0.309)	(0.040)	(0.045)
Married	0.866	0.015	-0.010	0.875	0.038	-0.004
	(0.342)	(0.045)	(0.043)	(0.332)	(0.036)	(0.034)
Household Size	3.899	0.257*	0.054	3.906	-0.128	0.151
	(1.314)	(0.139)	(0.165)	(1.400)	(0.168)	(0.170)
Muslim	0.034	-0.034	0.140*	0.019	-0.017	0.104**
	(0.181)	(0.032)	(0.074)	(0.136)	(0.014)	(0.052)
Years Living in	17.336	-3.687***	-0.478	16.950	-0.624	-0.388
Neighborhood	(10.285)	(1.190)	(1.247)	(10.548)	(1.278)	(1.175)
Number of Clients	10.309	0.071	0.010	10.456	-0.207	-0.102
in Group	(0.706)	(0.210)	(0.183)	(0.768)	(0.197)	(0.217)
Group Formed	0.644	-0.142	-0.108	0.663	-0.162	-0.207
in Rainy Season	(0.480)	(0.132)	(0.125)	(0.474)	(0.132)	(0.126)
Heavy Rain Days	5.517	-0.235	-0.482	5.394	-0.176	-0.758
	(2.075)	(0.607)	(0.598)	(2.050)	(0.619)	(0.526)
Panel B						
Client Worked for	0.564	0.003	0.008	0.488	0.113*	0.038
Pay in Last 7 Days	(0.498)	(0.068)	(0.064)	(0.501)	(0.065)	(0.062)
Household Earns	0.443	-0.069	0.085	0.431	-0.061	0.008
Fixed Salary	(0.498)	(0.066)	(0.064)	(0.497)	(0.058)	(0.061)
Household Owns	0.718	0.029	-0.074	0.719	0.039	-0.096
Business	(0.451)	(0.066)	(0.069)	(0.451)	(0.067)	(0.067)
Household Savings	1631.2	-117.2	1272.0	2016.1	298.1	887.4
	(4209.7)	(596.5)	(787.2)	(7958.3)	(1171.5)	(1205.7)
Household Owns Home	0.879	-0.160***	-0.118**	0.781	0.055	0.011
	(0.327)	(0.054)	(0.053)	(0.415)	(0.058)	(0.057)
Education Expenditures	4809.9	57.8	-940.2	4186.9	152.6	-248.4
	(5110.7)	(661.8)	(636.7)	(4727.1)	(544.9)	(422.8)
Health Expenditures	3264.5	-590.6	-295.4	3220.0	441.1	-125.2
	(5460.2)	(662.5)	(629.1)	(4870.2)	(727.5)	(559.6)
Illness in Past 12 Months	0.302	-0.042	-0.026	0.313	0.076	-0.010*
	(0.461)	(0.066)	(0.061)	(0.465)	(0.060)	(0.055)
Number of Transfers into	1.544	-0.838	-0.469	0.663	1.213**	0.038
Households	(6.263)	(0.545)	(0.628)	(2.320)	(0.538)	(0.271)
Number of Transfers out of	2.658	-0.578	-0.490	2.475	1.195	0.187
Households	(5.564)	(0.696)	(0.747)	(3.803)	(0.858)	(0.667)
Days between Loan	788.7	-1.1	14.1	788.0	0.6	13.6
Disbursement and Lottery	(46.1)	(12.8)	(11.7)	(46.4)	(11.3)	(11.9)
N	149	126	154	160	124	143

Notes

1 Columns (2)-(3) are the regression results of the characteristics in the title column on the two treatments for the 4-Rs. 50 Vouchers lottery sample. The omitted group is clients in Control groups. In columns (5)-(6) we report the same coefficients for the sample that received the 1-Rs. 200 Voucher lottery. All lottery sample regressions control for survey phase. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered by group.

Appendix Table 3.8. Default Determinants

	Second Loan Cycle Default
	(1)
Household Savings (x 10,000)	-0.010
	(0.006)
Illness in Past 12 Months	0.027*
	(0.016)
Number of Transfers out of Households	-0.0043***
	(0.0016)
Household Owns Home	0.032**
	(0.015)
Specification	OLS
N	1015

Notes

1 Variables are as defined in Tables 3.1 and 3.4. Regression also controls for First Loan Cycle repayment/meeting schedule.

2 * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered by group.

Appendix Table 3.9. Rainfall Robustness Checks

	Any Household Wage Income in Past 7 Days	Total Household Wage Income in Past 7 Days	Default	Default	Default
	(1)	(2)	(3)	(4)	(5)
Panel A: No Controls	0.020	18.416	-0.068*	-0.067*	-0.063
Group Met Weekly	(0.018)	(24.451)	(0.034)	(0.035)	(0.049)
Panel B: Controls Included					
Group Met Weekly	0.020	18.692	-0.079**	-0.077*	-0.076
	(0.017)	(22.934)	(0.039)	(0.039)	(0.051)
Mean of Dependent Variable	0.669	608.9	0.058		
	[0.471]	[620.0]	[0.235]		
Specification	Probit	OLS	Linear IV (29-56 days, 85th percentile rain)	Linear IV (29-56 days, 80th percentile rain)	Linear IV (29-84 days, 90th percentile rain)
N	996	996	720	720	720

Notes

1 Columns (3)-(5) replicate column (4) of Table 3.4 using different rainfall measures as a robustness check.

2 For column (3)-(5), Panel A regressions include a control for Group Formed in Rainy Season. For all columns, regressions with controls (Panel B) include the variables in Table 3.1, Panel A. * , **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered by group.

Appendix Figure 3.1. Lottery Vouchers

VILLAGE WELFARE SOCIETY **VILLAGE MICRO CREDIT SERVICES**

TOGETHER TOWARDS PROSPERITY

Rs. 50 Voucher

Single Use Only

Whoever reclaims this voucher must bring their VWS passbook with them to the VWS village bazaar when making their purchase. If the claimant is no longer a VWS client, they should bring their voter identification card.

Date of Lottery: ____________________ **Deadline to Claim:** ____________________

Group Name: ____________________ **Name of Claimant:** ____________________

Name of Winner: ____________________ **Signature of Claimant:** ____________________

Signature of Winner: ____________________

VILLAGE WELFARE SOCIETY **VILLAGE MICRO CREDIT SERVICES**

TOGETHER TOWARDS PROSPERITY

Rs. 200 Voucher

Single Use Only

Whoever reclaims this voucher must bring their VWS passbook with them to the VWS village bazaar when making their purchase. If the claimant is no longer a VWS client, they should bring their voter identification card.

Date of Lottery: ____________________ **Deadline to Claim:** ____________________

Group Name: ____________________ **Name of Claimant:** ____________________

Name of Winner: ____________________ **Signature of Claimant:** ____________________

Signature of Winner: ____________________

Note:

Clients were randomly offered entry into the lottery for a Rs. 200 Voucher or four Rs. 50 Vouchers. This figure shows the final vouchers which were given to the winner of the two lotteries.

Appendix Figure 3.2. Expected Returns to Lottery by Ticket-Giving Decision

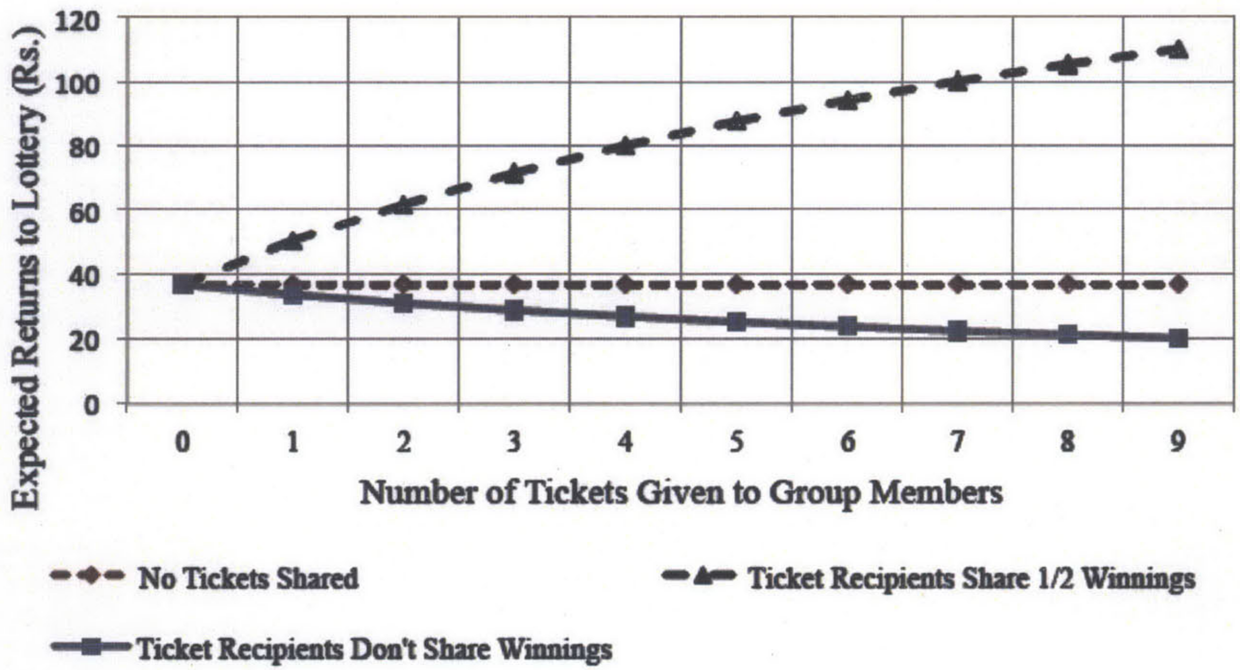

Notes:
Appendix Figure 3.2 shows the expected returns to the lottery based on ticket-giving decision, and extent of reciprocal behavior by ticket recipient.

Appendix Figure 3.3. Distribution of Group Meetings Held

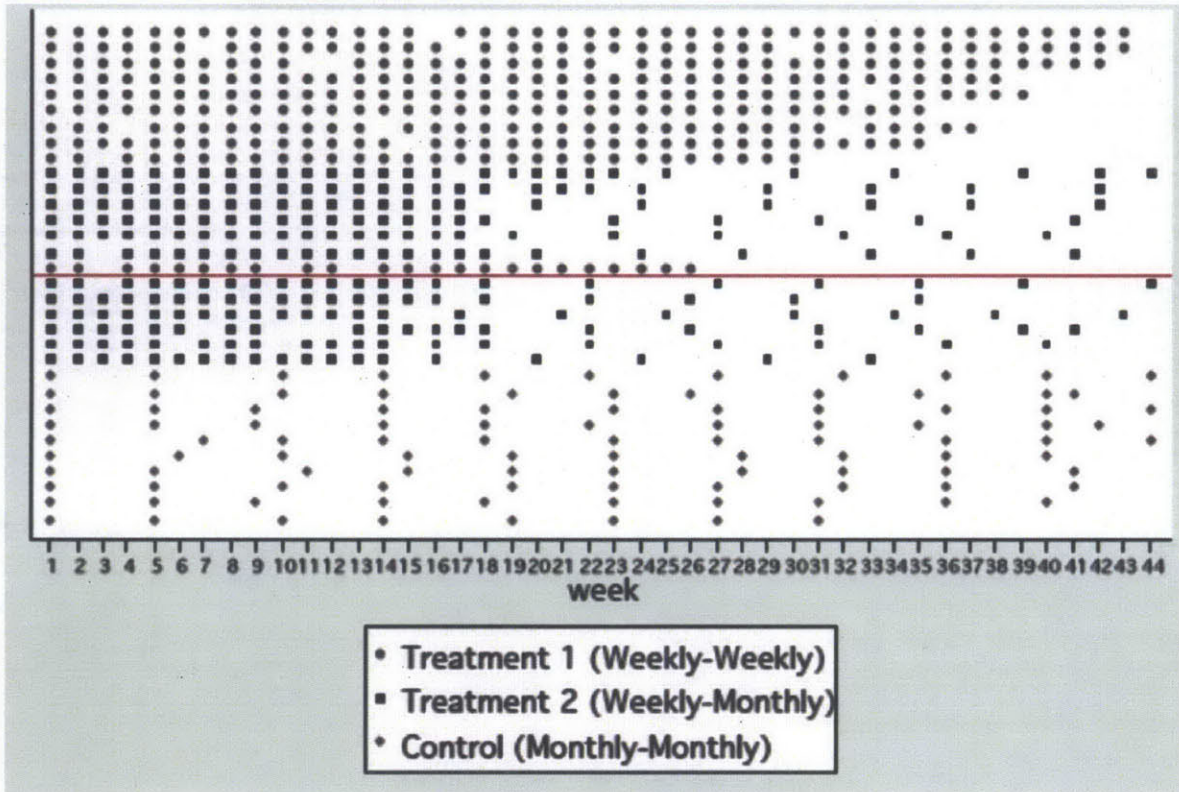

Notes:
We sample one-third of groups from each experimental branch, and stratify by quartile of number of group meetings held within each branch to ensure representativeness. Groups above the horizontal line met 23 or more times over First Loan Cycle.

Bibliography

Agency of Educational Quality. What is the SIMCE? `http://www.agenciaeducacion.cl/simce/ques-es-el-simce/`, 2013.

A. Alesina and E. La Ferrara. Who Trusts Others? *Journal of Public Economics*, 85(2):207–234, 2002.

A. Ambrus, M. Möbius, and A. Szeidl. Consumption Risk-sharing in Social Networks. *American Economic Review*, 104:149–182, 2014.

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

P. Andreas. The Transformation of Migrant Smuggling across the U.S.-Mexico Border. In D. Kyle and R. Koslowshi, editors, *Global Human Smuggling: Comparative Perspectives.* The Johns Hopkins University Press, 2001.

M. Angelucci. U.S. Border Enforcement and the Net Flow of Mexican Illegal Migration. *Economic Development and Cultural Change*, 60(2):311–357, 2012.

M. Angelucci, G. de Giorgi, and I. Rasul. Resource Pooling within Family Networks: Insurance and Investment. Working Paper, 2012.

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

J. D. Angrist and J.-S. Pischke. *Mostly Harmless Econometrics: an Empiricist's Companion.* Princeton University, Princeton, NJ, 2009.

O. Attanasio, A. Barr, J. C. Cardenas, G. Genicot, and C. Meghir. Group Formation and Risk Pooling in a Field Experiment. *American Economic Journal: Applied Economics*, 4(2):134–167, 2012.

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

A. Barr and G. Genicot. Risk Sharing, Commitment, and Information: An Experimental Analysis. *Journal of the European Economic Association*, 6(6):1151–1185, 2008.

T. Barrios, R. Diamond, G. Imbens, and M. Kolesar. Clustering, Spatial Correlation, and Randomization Inference. *Journal of the American Statistical Association*, 107(498):578–591, 2012.

M. Bauer, J. Chytilov, and J. Morduch. Behavioral Foundations of Microcredit: Experimental and Survey Evidence from Rural India. *American Economic Review*, 102(2):1118–1139, 2012.

L. Beaman and J. Magruder. Who Gets the Job Referral? Evidence from a Social Networks Experiment. *American Economic Review*, 102:3574–3593, 2012.

J. Berg, J. Dickhaut, and K. McCabe. Trust, Reciprocity and Social History. *Games and Economic Behavior*, 10(1):122–142, 1995.

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

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

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

T. Besley and S. Coate. Group Lending, Repayment Incentives and Social Collateral. *Journal of Development Economics*, 46(1):1–18, 1995.

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

G. J. Borjas. Self-Selection and the Earnings of Immigrants. *American Economic Review*, 77(4):531–553, 1987.

G. J. Borjas. Economic Theory and International Migration. *International Migration Review*, 23(3):457–485, 1989.

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

B. Cadena and B. Kovak. Immigrants Equilibrate Local Labor Markets: Evidence from the Great Recession. 2013. Working Paper.

C. Cameron, J. Gelbach, and D. Miller. Bootstrap-Based Improvements for Inference with Clustered Errors. *Review of Economics and Statistics*, 90:414–427, 2008.

M. Carter and M. Castillo. Does It Take More than Markets to Get Ahead? An Experimental Approach to Social Capital in South Africa. Agricultural and Applied Economics Staff Paper Series, Paper No. 448, University of Wisconsin, 2004.

M. Carter and M. Castillo. Coping with Disaster: Morals, Markets and Mutual Insurance - Using Economic Experiments to Study Recovery from Hurricane Mitch. In C. Barrett, editor, *The Social Economics of Poverty: On Identities, Communities, Groups, and Networks*, New York, 2005. Routledge.

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

D. Chiquiar and G. Hanson. International Migration, Self-Selection and the Distribution of Wages: Evidence from Mexico and the United States. *Journal of Political Economy*, 113(2):239–281, 2005.

Consultative Group to Assist the Poor (CGAP). `http://www.cgap.org/p/site/c/`, 2012.

D. Costa and M. Kahn. Understanding the Decline in American Social Capital, 1953-1998. *Kyklos*, 56(1): 17–46, 2003.

P. Dal Bo. Cooperation Under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games. *American Economic Review*, 95(5):1591–1604, 2005.

M. del Bosque. Holes in the Wall: Homeland Security won't Say Why the Border Wall is Bypassing the Wealthy and Politically Connected. `http://www.texasobserver.org/archives/item/15288-2688-holes-in-the-wall`, February 22 2008. The Washington Times.

D. DiPasquale and E. Glaeser. Incentives and Social Capital: Are Homeowners Better Citizens? *Journal of Urban Economics*, 45(2):354–384, 1999.

Q.-A. Do, S. Leider, M. Möbius, and T. Rosenblat. Directed Altruism and Enforced Reciprocity in Social Networks. *Quarterly Journal of Economics*, 124(9):1815–1851, 2009.

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

G. Elacqua. For-Profit Schooling and the Politics of Education Reform in Chile: When Ideology Trumps Evidence. 2009. CPCE Working Paper No. 5.

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

M. Fafchamps and F. Gubert. The Formation of Risk-Sharing Networks. *Journal of Development Economics*, 83(2):326–50, 2007.

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

E. Field and R. Pande. Repayment Frequency and Default in Microfinance: Evidence from India. *Journal of European Economic Association*, 6(2-3):501–509, 2008.

E. Field, R. Pande, J. Papp, and Y. J. Park. Repayment Flexibility Can Reduce Financial Stress: Experimental Evidence from Microfinance in India. *PLoS One*, 7(9):e45679, 2012.

E. Field, R. Pande, J. Papp, and N. Rigol. Does the Classic Microfinance Model Discourage Entrepreneurship Among the Poor? Experimental Evidence from India. *American Economic Review*, 103(6):2196–2226, 2013.

G. Fischer and M. Ghatak. Repayment Frequency and Lending-Contracts with Present-Biased Borrowers. London School of Economics Working Paper, 2010.

R. Forsythe, J. Horowitz, N.Savin, and M. Sefton. Fairness in Simple Bargaining Games. *Games and Economic Behavior*, 6:347–69, 1994.

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

F. Gallego and A. Hernando. School Choice in Chile: Looking at the Demand Side. Working Paper No. 356, Catholic University of Chile, 2009.

C. Gathmann. Effects of Enforcement on Illegal Markets: Evidence from Migrant Smuggling along the Southwestern Border. *Journal of Public Economics*, 92:1926–1941, 2008.

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

M. Ghatak and T. W. Guinnane. The Economics of Lending with Joint Liability: Theory and Practice. *Journal of Development Economics*, 60(1):195–228, 1999.

X. Gine and D. Karlan. Group versus Individual Liability: Long Term Evidence from Philippine Microcredit Lending Groups. 2009. Working Paper.

X. Gine, P. Jakiela, D. Karlan, and J. Morduch. Microfinance Games. *American Economic Journal: Applied Economics*, 2(3):60–95, 2010.

E. Glaeser, D. Laibson, J. Scheinkman, and C. Soutter. Measuring Trust. *Quarterly Journal of Economics*, 115(3):811–846, 2000.

U. Gneezy, W. Guth, and F. Verboven. Presents or Investments? An Experimental Analysis. *Journal of Economic Psychology*, 21(5):481–493, 2000.

L. Guiso, P. Sapienza, and L. Zingales. The Role of Social Capital in Financial Development. *American Economic Review*, 94(3):526–556, 2004.

C. C. Haddal. Border Security: The Role of the U.S. Border Patrol, August 2010. Congressional Research Service (CRS).

G. H. Hanson and A. Spilimbergo. Illegal Immigration, Border Enforcement, and Relative Wages: Evidence from Apprehensions at the U.S.-Mexico Border. *American Economic Review*, 89(5):1337–1357, 1999.

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

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

N. Humphreys, J. Fearon, and J. Weinstein. Can Development Aid Contribute to Social Cohesion After Civil War? *American Economic Review*, 99(2):287–291, 2009.

D. Karlan. Using Experimental Economics to Measure Social Capital and Predict Real Financial Decisions. *American Economic Review*, 95(5):1688–1699, 2005.

D. Karlan. Social Connections and Group Banking. *Economic Journal*, 117(517):F52–F84, 02, 2007.

D. Karlan, M. Mobius, T. Rosenblat, and A. Szeidl. Trust and Social Collateral. *Quarterly Journal of Economics*, 124(3):1307–1361, 2009.

E. King, L. Rawlings, M. Gutierrez, C. Pardo, and C. Torres. Colombia's Targeted Education Voucher Program: Features, Coverage, and Participation. 1997. Working Paper No. 3, Development Economics Research Group, World Bank.

J. Kling, J. Liebman, and L. Katz. Experimental Analysis of Neighborhood Effects. *Econometrica*, 75(1): 83–119, 2007.

S. Knack and P. Keefer. Does Social Capital Have an Economic Payoff? A Cross-Country Investigation. *Quarterly Journal of Economics*, 112(4):1251–1288, 1997.

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

L. Larance. Fostering Social Capital Through NGO Design: Grameen Bank Membership in Bangladesh. *International Social Work*, 44, 2001.

E. Ligon and L. Schechter. Motives for Sharing in Social Networks. *Journal of Development Economics*, 99 (1):13–26, 2012.

E. Ligon and L. Schecter. The Value of Social Networks in Rural Paraguay. CUDARE Working Paper No. 1116, 2011.

B. J. Mace. Full Insurance in the Presence of Aggregate Uncertainty. *Journal of Political Economy*, 99(5): 928–956, 1991.

C. Manski. Identification of Endogenous Social Effects: The Reflection Problem. *Review of Economic Studies*, 60(3):531–542, 1993.

C. Manski. Economic Analysis of Social Interactions. *Journal of Economic Perspectives*, 14(3):115–136, 2000.

J. Martinez and C. Woodruff. Selectivity of Migrants from Mexico: What Does Net Migration Tell Us? 2007. Working Paper.

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

E. Miguel, P. Gertler, and D. Levine. Does Social Capital Promote Industrialization? Evidence from a Rapid Industrializer. *Review of Economics and Statistics*, 87(4):754–762, 2005.

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

MIX Market, 2012. URL `http://www.themix.org/`.

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

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

J. F.-H. Moraga. New Evidence on Emigrant Selection. *Review of Economics and Statistics*, 93(1):72–96, 2011.

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

J. Napolitano. Remarks on Border Security at the University of Texas at El Paso. `http://www.dhs.gov/ynews/releases/pr_1296496548700.shtm`, January 2011.

D. Narayan and L. Pritchett. Cents and Sociability: Household Income and Social Capital in Rural Tanzania. *Economic Development and Cultural Change*, 47(4):871–97, 1999.

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

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

Nuñez-Neto, B. and S. Viña. Border Security: Barriers Along the U.S. International Border, September 2006. Congressional Research Service (CRS).

B. Olken. Do Television and Radio Destroy Social Capital? Evidence from Indonesian Villages. *American Economic Journal: Applied Economics*, 1(4):1–33, 2009.

P. M. Orrenius and M. Zavodny. Self-Selection among Undocumented Immigrants from Mexico. *Journal of Development Economics*, 40(1):1–24, 1993.

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

J. Preston. Beside a Path to Citizenship, a New Path on Immigration. `http://www.nytimes.com/2013/04/17/us/senators-set-to-unveil-immigration-bill.html`, April 16 2013. The New York Times.

Public Broadcasting Service (PBS). The Border. `http://www.pbs.org/kpbs/theborder/`, 1999.

R. Putnam. *Making Democracy Work: Civic Traditions in Modern Italy.* Princeton University Press, Princeton, NJ, 1993.

M. Rendall, P. Brownell, and S. Kups. Declining Return Migration from the United States to Mexico in the late-2000s Recession. 2010. RAND Working Paper.

R. Rivkin and P. Trottenberg. Treatment of the Economic Value of a Statistical Life in Departmental Analyses- 2011 Interim Adjustment, July 2011. Department of Transportation.

M. Rosenzweig. Education and Migration: A Global Perspective. 2007. Mimeo, Yale University.

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

A. Roy. Some Thoughts on the Distribution of Earnings. *Oxford Economic Papers*, 3:135–146, 1951.

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

L. A. Sjaastad. Individual-level Evidence for the Causes and Consequences of Social Capital. *Journal of Political Economy*, 70(5):80–93, 1962.

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

R. M. Stana. DHS Has Faced Challenges Deploying Technology and Fencing Along the Southwest Border, May 2010. Government Accountability Office (GAO).

R. M. Stana, S. Quinlan, and J. Espinola. Secure Border Initiative Fence Construction Costs, January 2009a. Government Accountability Office (GAO).

R. M. Stana, S. Quinlan, and J. Espinola. Secure Border Initiative: Technology Deployment Delays Persist and the Impact of Border Fencing Has Not Been Assessed, September 2009b. Government Accountability Office (GAO).

The Secure Fence Act of 2006. H.R. 6061 (109th).

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

R. Townsend. Risk and Insurance in Village India. *Econometrica*, 62(3):539–591, 1994.

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

United Nations, Department of Economic and Social Affairs, Population Division (UN Population Division). Trends in International Migrant Stock: The 2008 Revision. `http://esa.un.org/migration/p2k0data.asp`, 2009.

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

U.S. Customs and Border Protection (CBP). Environmental Assessment: Pedestrian Fence Near Sasabe, Arizona, July 2007.

U.S. Department of Homeland Security (DHS). U.S. Customs and Border Patrol: Regarding Border Fence, April-July 2007.

U.S. Department of Homeland Security (DHS). U.S. Customs and Border Patrol: Regarding Border Fence: 4/28/10- RE: for Future Consideration Redacted 2- RE: Requests Redacted 2, April 2010.

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

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

B. Wydick. Can Social Cohesion be Harnessed to Repair Market Failures? Evidence from Group Lending in Guatemala. *Economic Journal*, 109(July):463–475, 1999.