

Introduction

The development process is in large part about addressing market failures. Whether it is the inability of productive firms to obtain financing or children to receive quality education, private and public responses can be significantly constrained by market deficiencies. A central theme in my work has been to identify and address such limitations faced by developing countries in the process of accumulating physical, human, and institutional capital – three prominent factors that promote an economy’s development.

My primary training, at both the doctoral and undergraduate levels at Harvard and MIT, has been in applied microeconomic theory. I have built on this foundation through extensive fieldwork, comprehensive data collection, and rigorous empirical analysis in order to answer questions that are motivated by and engage with policy. This policy engagement has not only provided conceptual insights and facilitated data access but has also given me unique opportunities to design, implement, and evaluate interventions. It is this combination of theoretical and empirical analysis with fieldwork and policy engagement that allows my research to uniquely contribute to both academia and policy in three broad areas: finance, education, and institutions.

My work in finance uncovers informational and agency problems in emerging financial markets. This includes examining market manipulation, firm networks, and the impact of aggregate liquidity shocks. The empirical analysis exploits detailed secondary, though unexplored, data sources ranging from individual trade data in stock markets to credit registry databases that cover the universe of bank and institutional lending. My research has been published in leading economic journals, such as the *American Economic Review* and the *Quarterly Journal of Economics*, and is at the forefront of a new and rapidly growing field in emerging economy financial markets. Funded by a large seed grant from the Google.org Foundation to establish the *Entrepreneurial Finance Lab* (EFL) at Harvard, I have also begun to examine how alternative financial instruments may unleash the entrepreneurial potential in emerging economies, valued at several trillion dollars. Given the current financial crisis, the informational and regulatory issues I examine in my research are increasingly salient in developed markets, and I plan to contribute to research and policy in this area, as well.

My research in education addresses another challenge central to an economy’s growth – the development of human capital. With a mushrooming of affordable, private, “mom-and-pop” schools in many developing countries, there is a growing recognition that education is both publicly and privately provided. My work considers the nature, evolution, and implications of having such an educational market. In order to do so, I have lead an ambitious and novel effort that surveys the entire educational universe – schools, teachers, parents, and children – in over one hundred villages in Pakistan, capturing its evolution over a four-year period. In the process, I have directly engaged with local policy makers on a series of feasible, market-level experimental interventions that address informational and accountability failures. With well-received initial publications, including a paper winning the *George Bereday Award* in 2006, this work demonstrates the value of such longer-term, comprehensive projects. I use both observational and experimental evidence in order to shed light on the dynamics of the educational marketplace, to examine supply and demand-side constraints, and to design and analyze the market-wide impact of interventions that alleviate them.

In addition to financial and human capital, institutional factors – as determined by the culture, norms, and history of a region – are also central to the development process. My research also examines such factors, from a group’s ability for collective action to its beliefs and values. My work on understanding how religious institutions like the Hajj, the annual Islamic pilgrimage to Mecca, shape beliefs and behavior is particularly salient given the current concerns about Islamic extremism. Such work has drawn both academic and policy attention, being published in the *Quarterly Journal of Economics* and informing discussions in Washington and Islamabad, as well as receiving coverage in media outlets such as the *NY Times*, *Washington Post*, *International Herald Tribune*, *Economist*, *Al-Jazeera*, *BBC*, and *CNN*. More recently, I was awarded the 2009 *Carnegie Scholars Fellowship* in order to continue this research and identify mechanisms that may moderate extremism and promote tolerance in a world at risk of increasing divisions.

My empirical work so far has mostly focused on Pakistan due to my familiarity with the context, as well as to reduce the large fixed costs of gathering detailed data. While this region is of particular global interest, it is also fairly representative of emerging economies in the themes I examine. Therefore, the research insights generated have broader implications for developing economies and, in several cases, for the developed world. I am also expanding my research to other countries in ongoing work and I anticipate this will increase further.

In addition to the research I have also directly contributed to policy. This is best demonstrated in my co-founding the web-portal www.risepak.com set up to coordinate disaster relief following the 2005 Pakistan earthquake. The portal was extensively used and widely recognized, winning the 2006 *Stockholm Challenge Award* in Public Administration. In the coming years, I expect that my scholarship and policy engagement will remain deeply intertwined and that the positive synergies between the two will allow me to make unique contributions to both. Below I present my research in each area in more detail.

A. Finance and Firms

Financial markets intermediate the matching of money to ideas. To the extent that there are failures in such markets, the effectiveness of this role will be diminished, with potentially large consequences for the economy. My research in this area, mostly developed jointly with Atif Mian (Chicago, Booth GSB), examines this question from a series of related perspectives that focus on informational and agency problems prevalent in emerging and, increasingly, in developed financial markets.¹ I do so primarily by taking advantage of detailed data sources from Pakistan – ranging from individual trade data in stock markets to credit registry databases that cover the universe of formal bank and institutional lending – and developing methodologies that permit better causal inference.

There are four broad directions of enquiry that I am pursuing: uncovering market manipulation; tracing the impact of aggregate liquidity shocks; examining the nature and role of firm networks; and developing alternative instruments to help facilitate financial access.

Market manipulation

It is widely believed that manipulation of financial markets is prevalent in emerging economies and imposes a substantial cost. Recent financial scandals in the US suggest that manipulation may also still be at play, though

¹ While most of my work is co-authored, both in this section and subsequent ones, for the purpose of this document after explicitly identifying my coauthor(s), I shall subsequently use “I” rather than “we.”

potentially less visible, in the developed world and with just as severe consequences. However, finding direct evidence for such manipulation has been challenging.

I directly identify the presence of market manipulation in both equity and debt markets in two different papers that exploit novel micro-data from Pakistan. “Unchecked Intermediaries”² (*Journal of Financial Economics*, 05, with A. Mian) uncovers manipulation by market intermediaries in equity markets. I utilize a unique dataset of all daily trades for every broker trading during a two and a half year period in the main stock exchange in the country to identify a “pump and dump” price manipulation scheme believed to be prevalent in emerging economies and also reminiscent of trading behavior on the NYSE in the 1900s. When prices are low, colluding brokers trade among themselves to artificially raise prices and attract positive-feedback traders. Once prices have risen the brokers exit, leaving the traders to suffer the ensuing price fall. Estimating the profits generated by such trading manipulation shows that manipulating brokers earn up to 90 percentage point higher returns. Conservative estimates suggest a \$100 million annual transfer of wealth from outside investors to manipulating brokers – 10 percent of market capitalization. These substantial profits may explain the resistance of brokers to governance reforms and also suggest why emerging equity markets often remain marginal, with few outsiders investing and little capital raised.

Banks offer the dominant form of firm financing in most economies and manipulation in these markets is likely to have serious consequences on an economy’s growth and productivity. Moreover, rents obtained by the politically connected are singled out as a source of corruption in developing countries. In “Do Lenders Favor Politically Connected Firms?” (*Quarterly Journal of Economics*, 05, with A. Mian), I identify politically-based rent provision by exploiting an incredibly rich dataset that covers every loan by each bank to all firms borrowing in Pakistan over a seven-year period.

I find that politically connected firms receive substantial preferential treatment from public banks. A firm is classified as “politically connected” if any of its directors participated in a recent national or state election. The results show that politically connected firms borrow one and a half times as much and have 50 percent higher default rates. Suggestive of rent provision, such preferential treatment is only observed for loans given by government banks, with private banks displaying no such favorable treatment. A contribution of this work is to empirically identify that this preferential treatment is indeed a likely result of rent-provision to the politically influential. First, using firm fixed effects (and thereby exploiting variation for the same firm across lenders and over time) shows that unobserved time-invariant firm characteristics are unlikely to be a confounding factor. Second, I am also able to provide direct evidence against alternative explanations such as socially motivated lending by government banks to politicians. This work shows political corruption in financial markets can have a tremendously adverse impact – tentative estimates using the results in this paper suggest a cost of 1.9 percent of GDP per year, quite similar to the numbers obtained in cross-country macro-studies.

The results also show rents are affected by the nature of politics. They increase with the strength of the firm’s politician and political party and fall with electoral participation in his constituency. This raises several questions on how such rents can be checked, for example, by imposing constraints on political candidacy, governance reforms, and privatization, or by their electorate not re-electing rent-seeking politicians. On a

² For the sake of expositional clarity, I have shortened the paper titles. Complete titles are provided in my CV.

broader note, both these papers highlight the governance challenges faced in curbing manipulation. While the overall economic costs of such activities are substantial, there are large individual rents accrued – especially by those with influence. Therefore, unless one can design mechanisms to circumvent or compensate those who accrue rents, governance reforms, even if straightforward, may be hard to implement.

Liquidity Crises and Booms

An important and much debated question in the literature and one that is of particular concern these days is to what extent liquidity shocks can have large and real effects on the economy. While it has been argued that severe recessions like the Great Depression and the Asian Financial Crisis arose due to liquidity shocks, identifying the full impact of a bank liquidity shock has remained a challenge. This is primarily because of empirical difficulties in separating out the demand/productivity shocks that typically accompany the supply shocks.

In a series of papers, I am able to isolate the supply channel and highlight the agency problems that may underlie why liquidity shocks can have real consequences. I can do so by using methodologies that take advantage of both comprehensive data and two recent and large unanticipated liquidity shocks to the Pakistani financial markets: a liquidity crisis after the multilateral aid withdrawals due to the 1998 nuclear tests carried out by India and Pakistan; and a liquidity boom due to the reverse capital flight after 9/11. These shocks provide novel “natural experiment” settings that permit cleaner empirical analysis.

In “Tracing the Impact of Bank Liquidity Shocks” (*American Economic Review*, 08, with A. Mian), I find evidence of a substantial “bank lending” channel (shocks to the supply of a bank’s liquidity affect its lending) and trace its distributional consequences. Differential (across banks) liquidity shocks arising from the unanticipated nuclear tests and a dataset linking firms to all banks allow me to isolate the supply shock. Contemporaneous demand shocks are “netted out” by comparing changes in lending to the *same* firm borrowing from two different banks, thereby isolating the supply-side channel. The results reveal a large bank-lending channel, with a bank decreasing its lending to firms by 0.37 percent for every percent drop in its own liquidity (deposit base).

In addition to providing improved empirical identification of the bank-lending channel, I also contribute to the literature by tracing the impact of this channel on a client firm’s overall borrowing and subsequent financial health to reveal the distributional consequences of such shocks. I find that larger firms fully compensate the initial shock by increasing their borrowing from more liquid banks and establishing new banking relationships. In sharp contrast, smaller firms bear the entire brunt of the initial shock and are unable to compensate because they are unable to access new sources of bank financing. Such firms not only see large drops in their overall borrowing from the financial sector but are also more likely to enter financial distress in the years following the shock.

These results have important policy implications: They highlight the substantial imperfections faced in inter-bank markets. Moreover, they show that in addition to aggregate effects, financial crises may have even more costly distributional consequences if smaller firms face the primary brunt of such shocks. In the development literature, there is suggestive evidence of far fewer middle-sized firms compared to developed economies.

These results suggest that this may be partly due to constraints prevalent in financial markets that prevent smaller firms from growing.

One may hope that the negative impact of liquidity crises can ultimately be undone by liquidity booms. While this is challenging, particularly for smaller firms that may not have the resources to “ride out” the crises, what is even more worrying is that the impact of liquidity shocks may be asymmetric – while adverse shocks may lead to lending cuts, positive shocks may not raise lending as much. This asymmetry is plausible since the former involves frictions at the bank level (shocked banks are unable to raise liquidity) whereas the latter depends on banks being able to overcome frictions at the client firm level (liquidity flush banks still need to identify secure lending options).

In “Dollars Dollars Everywhere, Nor Any Dime to Lend” (*Review of Financial Studies*, R&R, with A. Mian and B. Zia), I find that banks may indeed be unable to exploit positive liquidity shocks even in fairly open emerging financial markets. The events of 9/11 resulted in large unanticipated capital inflows into Pakistan as citizens pulled back their foreign investments and increased remittances due to fear of account freezes in the West. As a result, interest rates in Pakistan fell by over 50%. Given the magnitude of the shock and the fact that it was unexpected, this provides an ideal setting in which to study the impact of positive shocks.

Basic economic theory would suggest that given such a large drop in the cost of capital, lending to firms should increase dramatically as one moves down a firm’s marginal product of capital curve. However, the results are quite stark and show that overall lending to firms remained relatively flat in the following year and a half, with most of the new liquidity channeled into the stock market and real estate speculation. While this is reconciled if 9/11 also led to large negative demand shocks, the macro evidence suggests the opposite. Given Pakistan’s almost immediate cooperation in the efforts against terrorism, 9/11 resulted in an improvement in aggregate demand as sanctions imposed after the 1998 nuclear tests were lifted and reconstruction efforts in Afghanistan started.

The evidence instead reveals that firm-level frictions were likely behind this inability of banks to lend despite having liquidity. I identify the presence of firm-level frictions by taking advantage of a novel feature in our data: a measure of a firm’s borrowing capacity (credit limit) as determined by its lender. To the extent that such limits are backward looking (based on a firm’s collateral, reputation, etc.), they won’t adjust rapidly enough to take advantage of the liquidity boom. Consistent with this, I find that firms with “financial slack” (unused borrowing capacity) at the time of the 9/11 shock indeed increased their borrowing afterwards while firms that had no slack were unable to do so. This suggests that the inability of banks to lend is not driven by a lack of firm demand but rather by lending frictions. Since firm collateral did not show commensurate increases, banks were unwilling to extend further credit to those borrowing up to their credit limits. Tentative estimates put the cost of such foregone lending to the economy at 2.3 percent of GDP.

While the finance literature typically checks for firm credit constraints using investment-cash flow sensitivity analysis (i.e., credit constrained firms are likely to have their investments respond to cash flows), there is a recognition that such analysis likely suffers from omitted firm-level factors. The paper therefore also contributes by proposing a methodology that is less susceptible to such concerns since it takes advantage of

an unanticipated shock and a measure (financial slack) that directly captures firm borrowing capacity, and because it accounts for unobserved, time-invariant firm attributes through firm and management fixed-effects.

These results emphasize that the problem in financial markets may not simply be a shortage of liquidity: If informational and other lending constraints are large enough, increasing liquidity may not be sufficient to generate an investment boom and, in fact, could have potentially perverse consequences such as speculative hikes and inflation. In future work, I hope to explore the theoretical basis and consequences for the asymmetric response to liquidity shocks and examine their impact on firm production, investment, and growth. Countries such as Indonesia have detailed firm censuses that permit such an analysis and contribute to an analysis of how financial markets affect the size, distribution, and growth of firms in a developing economy.

Firm Networks

In a series of papers, I also consider how firm networks reflect and contribute to financial market performance, adding to a growing area of research in network analysis.

In “Identifying Business Networks” (with A. Mian & A. Qamar), I document the presence of a robust “super-network” of firms linked through common directors. Using an ownership database of over 140,000 firms in Pakistan, I use board-interlocks to construct a network topology of firms. While the majority of firms are stand-alone or part of small networks, there is a single “super-network” of 10,000 firms. The super-network displays a “small-world,” loosely connected pattern with few “star” or highly-connected nodes and is robust to changes in the definition of links and to the removal of heavily-linked nodes. While this exercise has not been replicated for the universe of borrowers in other economies, analysis on a smaller sample of listed firms in the US and UK shows that these “super-networks” of inter-locked boards are also present in these countries.

In a follow up paper, “The Value of Business Networks” (with A. Mian & A. Qamar), I find that network membership generates substantial value in financial markets, leading to a 16.5 percent increase in a firm’s borrowing and a 9.7 percent lower probability of it defaulting. A challenge in estimating these impacts is to deal with selection effects. Differences in outcomes between network and non-network firms may therefore reflect prior differences between firms rather than the effect of joining the network. While the paper makes substantial progress in dealing with time-invariant firm attributes by exploiting variation *within* the same firm for firms that enter and/or exit the super-network over time, it takes a step further by dealing with time-varying selection issues. I instrument for a firm’s entry/exit by exploiting the fact that some firms enter/exit a super-network not because they experience any change in their board but because one of the firms they were previously connected to experiences such a change. By focusing on such firms that are “indirectly” pulled in or out of the network, I am able to lessen time-varying selection concerns. To my knowledge, this methodology has not been applied in the network literature previously and offers an additional contribution that can be replicated in a variety of network settings in other empirical fields.

While the theoretical literature on networks suggests heterogeneity in network returns, there are few empirical results that establish this. Another contribution of this work is to provide such evidence. Moreover, the results show that this heterogeneity varies depending on the particular channel through which the network

generates value. For financial access, a firm benefits more from entry when it connects to powerful parts of the network and benefits relatively less if it was already powerful, i.e., entry into the super-network is a substitute for a firm's preexisting power. In contrast, for financial distress, network entry is a complement to a firm's pre-entry power. This hints that the mechanism for lowering the default rate may be quite different from that for improving financial access. While the latter likely reflects leveraging neighbors' connections with lenders, the former may arise by benefiting from one's neighbors through internal insurance and credit flows, and business contacts. Hence, the more powerful firms take more advantage of their network neighbors' resources. As further evidence for the importance of these internal flows, I also find that networked firms are better insured against industry and local demand shocks than non-networked firms.

Networks are also salient in manufacturing and retail settings. In "Subcontractors for Tractors" (*Journal of Development Economics* 05, with T. Andrabi & M. Ghatak), I model how a network setting allows buyers to manage aggregate demand shocks and enables both low and high quality suppliers to coexist. I then use primary data on contracts between the largest tractor assembler in Pakistan and its suppliers to document large differences in how suppliers of the same product are treated. The results show that more "tied" suppliers (that make dedicated investments to the buyer) receive smaller and more unstable orders, as well as lower prices. The model reconciles this puzzle by recognizing that relationship-specific investments are not just technologically driven, as is typically assumed in the literature, but are also a matter of choice. Given that dedication has higher costs for better quality suppliers, low-type suppliers are kept as marginal suppliers because of their greater willingness to invest in buyer-specific assets. Low quality suppliers survive in equilibrium because the buyer uses them as "capacity" buffers in the absence of requisite insurance markets. These results demonstrate that market imperfections have an impact not only on the organization of firms but also on their "type" distribution, with less productive firms surviving despite competitive pressures.

In an ongoing project, "Divided We Fall" (with J. Das and A. Mian), I examine retail clusters, an organization structure prevalent in developing countries that is defined as a network of shops selling similar products at a common location. I first develop a model that shows how retail clusters may emerge as an equilibrium outcome with the same product (optimally) priced differently across shops in the same cluster. I then analyze a unique primary survey of 700 retail shops in 15 retail categories in a large city (Lahore, Pakistan) and show that, apart from inventory sharing considerations, clusters also arise due to an inability of retailers to smooth input price shocks. This inability, in turn, produces strong "price comparison" incentives on the part of buyers and, therefore, a preference to shop at clusters (i.e., an individual shops in a cluster not only to obtain a greater variety but also to find the best price since no one seller can consistently "beat the market"). The prevalence of retail clusters in developing countries shows how powerful market context may be in shaping firm organizational structure.

Alternative Financial Instruments

Driven by the shortcoming of standard lending instruments, I have begun to examine and develop alternate financial instruments that facilitate financial access for borrowers often excluded from traditional lending markets.

In "Can Individual Lenders Infer Borrower Creditworthiness?" (with R. Iyer, E. Luttmer & K. Shue), I examine peer-to-peer (p2p) credit markets that lend to smaller borrowers (individuals and small firms) in

developed and developing economies. While these relatively new, usually web-based p2p lending markets such as Prosper, Zopa, Kiva, Myc4, Lending Club, and Fynanz, have grown both in number and in size (Prosper has funded over \$178 million in loans in the US and currently has 830,000 members), the long-run viability of this market remains to be seen.

I use proprietary data from Prosper.com to show that, despite not being financial experts, individual lenders can infer up to a third of the relevant information content regarding a borrower's credit-worthiness, as proxied by their credit score. This inference is economically meaningful and allows them to lend at a 200 basis points lower rate than had they not been able to draw inferences. The methodology developed in the paper to estimate this inference parameter is fairly novel. It takes advantage of lenders not observing a borrower's true credit score and only seeing an aggregate credit category. The results also show that while lenders infer the most from standard "hard" information traditionally used by banks for screening, they also use non-standard (subjective) information for inference. The methodology shows, without needing to code information contained in pictures or personal descriptions a borrower posts, that lenders learn even from such "softer" information. This is particularly the case when such information provides credible/costly signals regarding borrower credit worthiness. These findings suggest that p2p markets may even provide advantages over traditional lending markets by extracting information from non-traditional, "softer" sources.

I have also recently obtained a seed grant from Google.org to set up the Entrepreneurial Finance Lab (www.cid.harvard.edu/efl) at Harvard. The Lab's focus is exploring alternate ways to extend capital to the relatively untapped small and medium entrepreneurs in developing countries – a sector that offers a growth potential of as much as several trillion dollars a year. The initial idea is to develop and validate a scalable micro-venture capital (VC) model. Like VCs, this model screens entrepreneurs but proposes to substantially lower the cost of doing so by relying on automated psychometric testing tools and by making smaller investments in thousands of smaller enterprises. Moreover, rather than exit out of these enterprises like banks, the proposed model generates returns through a repeated relationship. The initial phase of the work – validating the screening tool by partnering with banks, VCs, and MFIs across several client types and countries in Africa, Latin America, and Asia – is already underway, and preliminary results are encouraging.

B. Education

There is an increasing realization that education in developing countries is no longer characterized by a simple choice environment where parents only decide to enroll a child in a (public) school or not. Instead, it is one where parents choose among multiple schools, of differing price and quality, across multiple children. While recent empirical work has made substantial contributions through experimental interventions that estimate returns to educational inputs, we have just started examining the overall educational marketplace and its evolution in developing countries. This is particularly relevant given the rising affordable private sector in developing countries, with a rapid growth of small and independent "mom-and-pop" schools. My research seeks to understand the complexities and challenges of having such a choice-rich educational market. I examine imperfections faced on both the supply and demand side and evaluate interventions intended to alleviate market-level shortcomings by drawing on both observational and experimental evidence.

A challenge of this research is the data requirements imposed. In order to understand the overall market, one needs matched information on all significant market players and must follow them over time. For the past

several years I have jointly led, with J. Das (World Bank) and T. Andrabi (Pomona College), an ambitious and novel exercise to gather a comprehensive, four-year panel dataset that captures the educational universe in 112 villages in rural Pakistan (www.leapsproject.org). The “Learning and Achievement in Pakistan Schools” (LEAPS) data starts with a baseline census of all 85,000 households and consists of four (annual) rounds of matched surveys of all 800 primary schools, 5,000 teachers, and 2,000 households, as well as data from achievement tests we designed and administered to more than 12,000 children tracked over the 4 years. To our knowledge there are few datasets in developing or developed countries that provide such a comprehensive picture of the educational landscape and its evolution. While I am still in the early stages of producing this work, the LEAPS project has already generated significant insights and begun to uniquely contribute to the literature. My research is broadly grouped in three themes that examine supply-side constraints, demand considerations, and how these two forces interact to create the educational marketplace.

The Supply of Education

The 2000 Millennium Development Goals (MDGs) reflected a powerful global consensus to improve the development of poor countries by 2015. Central to this promise is ensuring universal primary education. However, by 2005, most countries had already fallen well behind this goal. While this has led to renewed calls for greater public investment in schools, the role of private schools has been notably absent from the debate. This is partly because private schools at the primary level have traditionally been thought of as catering to the elite. In the US, only 3 percent of parents with less than a high school education send their children to private schools while the share is 19 percent for parents with graduate or professional degrees. However, the private sector in many emerging economies in South Asia, the Middle East and Gulf States, and to a lesser extent in Africa and Europe, has increasingly begun to play a salient role. In India, the 2005 National Sample Survey shows that private schools account for 30 to 40 percent of primary enrollments in several states.

In “A Dime a Day” (*Comparative Education Review* 08, with T. Andrabi and J. Das), I detail the phenomenal rise of the private sector in Pakistan – a ten-fold increase in less than two decades and comprising over a third of all enrollment, with an increasing share of children from rural areas and middle-class and poorer families. These are not NGO or religious schools but small, “mom-and-pop,” for-profit schools set up by local entrepreneurs. If replicable in other countries, this suggests that separating the financing of education (the government) from the provision of education (through the private sector) could have large gains. Private schools are able to offer low cost education – with median fees less than a day’s unskilled wage – by providing a bare-bones, “teacher only” model and by taking advantage of local labor markets. They hire mostly locally-educated women who are paid a sixth of the typical salary of public sector teachers. While part of this reflects their lower education and experience levels, in the case of Pakistan, this is driven by the fact that local/rural educated women have fewer alternative work opportunities due to occupational and mobility constraints. This suggests private schools are more likely to arise in villages where there is a relatively cheaper source of teachers, such as in villages that have a local girls’ high school.

In a companion paper, “Students Today, Teachers Tomorrow?” (with T. Andrabi and J. Das), I show that supply-side factors can indeed play an important role in the rise of private schools and provide clues to why this phenomenon may not be more ubiquitous. I argue that public investment in education facilitates future education provision since it lowers production costs for the private sector by creating an affordable local pool of potential teachers.

There are two parts to this argument. First, by utilizing school construction guidelines specified under government school expansion programs in Pakistan, I obtain exogenous variation in school provision to show that the construction of government girls' high schools has a large causal impact on the education market. Villages where such schools were constructed are three times more likely to have a private school. The identification strategy exploits non-linearities and discontinuities in the population-based girls' high school eligibility rule, which arises because two villages with equal population may or may not be eligible depending on their population rank among their neighbors.

Second, I argue that this effect works through a teacher supply channel. What helps substantiate this claim is that the context of geographically and occupationally restricted labor markets for women (but not men) generates cross-sectional variation in the impact of public school provision. In support of a "women as teachers" channel, I find that, in sharp contrast to the large impact of a girls' high school, having a government boys' primary/high school or a girls' primary school has no effect on private school existence. Although this could be consistent with a situation where only highly educated mothers are the prime drivers of demand, price movements in the education market provide compelling evidence for a net supply shifter: The wages of private school teachers are 18 percent *lower* in villages with a preexisting girls' high school.

While the recent focus both in education practice and scholarship has been on incentivizing educational demand, these findings highlight an important supply-side constraint to schooling – the supply of teachers. In doing so, this research provides compelling evidence in support of an older literature that argued for higher education as a means of facilitating primary education. Moreover, it underscores the prominent role of women, not only as mothers but as teachers in facilitating educational access, and resonates with similar historical evidence from developed economies.

The Demand for Education

The premise of many of the increasingly popular conditional cash transfer programs in developing countries is that parental demand for and investments in education may be suboptimal if parents do not internalize educational benefits or reap the full benefits of their investments. Therefore, parents need to be incentivized in making educational choices. For example, in Pakistan, programs that condition payment on female enrollment have been implemented in response to perceived gender discrimination within the household. However, an initial examination of the LEAPS data reveals that not only are parents actively engaged in the education process but the nature of discrimination may be quite different. Detailed household surveys show that parents spend considerable financial resources on their children's education, with the typical sample household spending 3 to 5 percent of their overall budget on each child's education. Even in the poorest third of households, out-of-pocket expenditures amount to 75 percent of government educational spending on these households (through free public schooling). Discrimination (in terms of educational investments) within the household is much larger across perceived child ability than across gender, that is, parents tend to invest a lot more resources in children they perceive to have higher ability.

Whether parental investments in children compensate for or reinforce existing endowment differences is a question that has received attention in the literature since the 1970s. However, empirically identifying the extent and motivations for such investments has proven challenging due to an array of omitted variable concerns. I plan to overcome such challenges by using novel data on parental perception of each child's

ability, detailed time use and investment data for every child, and exogenous variation in perceptions. Taking advantage of an experimental intervention that distributed child performance report cards (based on the LEAPS achievement tests), I will examine how parents reallocate their investments from children who they (exogenously) discover are better/worse in terms of educational performance than they had previously thought. Since the report cards were only distributed in treatment villages and only to tested (Grade 3) children, exploiting cross-village, cross-household, and cross-sibling variation will better isolate how households respond to changes in perceived child ability. In addition, by examining which types of investments are reallocated, a further contribution of this work will be to uncover the motivations - investment or preference based - for why parents differentially invest in their children.

While it is widely believed that the presence of educated mothers leads to beneficial child outcomes, it is less clear how this benefit arises. Bargaining theories suggest that more educated women have a greater say in household decision making. However, given the low levels of education for women in most developing countries, this would imply that such beneficial effects may be hard to come by. In “What Did You Do All Day?” (with T. Andrabi, & J. Das), I find that even low levels of maternal education may lead to substantial child benefits. By matching the availability of schools in the mother’s birth village to her educational levels, I present instrumental variable estimates that account for the selection into maternal education. I find that in households with educated mothers, time spent by their children on educational activities at home goes up by over an hour a day. While this is partly driven by the mother spending more time directly working with her children, most of it is generated by having the child devote more time to studying at home. The learning impact of maternal education associated with this extra child effort is large. Test scores in English, Urdu, and Mathematics go up by a fifth to a third of a standard deviation. These results show that even with low levels of maternal education, which may not be sufficient to affect household bargaining outcomes, an educated mother may significantly raise her child’s learning outcomes by creating a more conducive learning environment within the household.

In recognizing that parents have to make costly tradeoffs between price and quality, understanding the learning process becomes central. In “Here Today, Gone Tomorrow?” (with T. Andrabi, J. Das & T. Zajonc), I show that retaining learning levels may be particularly challenging since only a fifth to a half of achievement persists between grades. This complicates the parental decision-making process. Not taking into account that children with higher levels of learning face larger future falls may significantly understate the effect of educational inputs and therefore bias decision-making. Using private schools as an example, the paper demonstrates that incorrectly assuming high persistence in learning significantly understates and occasionally yields the wrong sign for private schools’ impact on achievement. The paper illustrates the potential pitfalls of using the value-added estimators commonly used in the literature and contributes methodologically by applying a wider set of econometric tools from the dynamic panel data literature than have been typically used in the education literature. Finally, towards an economic interpretation of low learning persistence, I use question-level exam responses, as well as household expenditure and time-use data, and find little support for measurement error, mechanical psychometric testing, and behavioral response explanations, suggesting that forgetting may be an important factor. While this is a promising area of future enquiry, it does underscore the challenges faced in improving learning outcomes and parental decision making, particularly in developing countries.

The Educational Marketplace

Despite the policy and academic interest in education, we have only just begun to build a picture of the educational landscape in developing countries. In a forthcoming book (Oxford University Press, with T. Andrabi, J. Das, T. Vishwanath, & T. Zajonc,) that uses the first year of the LEAPS data, I provide a basic but comprehensive view of this landscape in one setting and reveal several insights.

As noted earlier, the data shows that the schooling market is characterized by engaged, active parents and substantial school choice and competition. What is perhaps even more revealing, though, is that while low enrollment is an issue it is improving, and increasingly the challenge going forward is low learning levels – a fact that has been relatively overlooked in setting the Millennium Development Goals. In fact, learning across schools differs much more within than across villages.

This learning difference within villages is driven primarily by school type, with private schools showing a standard deviation higher learning outcomes than public schools. While this could be partly driven by selection differences, with richer and more educated parents sending their children to the (slightly more) expensive private schools, preliminary results show that this is unlikely. Controlling for parental and child background and using instrumental variable techniques suggest that this difference is not due to selection effects and that the private sector does add greater value. Given that private school teachers are less qualified and paid substantially less than their public counterparts, this suggests incentives may be the key difference. In support of this, I find that while private sector wages do respond positively to student performance, those in the public sector show a weak relationship and may, in fact, be downward sloping since wages depend on seniority levels. This raises important questions regarding teacher quality versus incentives. Even though the public sector selects more qualified teachers (in terms of education, experience, and training levels), its inability to provide effective incentives causes it to significantly underperform compared to a poorly endowed but highly incentivized private sector.

In an effort to better understand what factors may hinder the development of an effective educational marketplace, I have begun to examine experimental interventions that address the potential informational and organizational market failures suggested by the descriptive analysis.

In “Report Cards” (with T. Andrabi and J. Das), I examine the impact of a market-level experimental intervention that provides school and child-level report cards on learning. Evidence on the impact of information provision on service delivery has been mixed, with overall outcomes potentially worsening if providers cream-skim and there is excessive sorting. The findings in this paper, however, paint a more promising picture. I find that providing report cards improves overall learning by a tenth of a standard deviation and lowers private school fees by 18 percent, with little switching or sorting of children across schools. The results reveal substantial heterogeneity in the impact. Initially below-median (bad) private schools show the largest learning gains – over a third of a standard deviation. While the worst of the bad private schools are more likely to close, on average, it is the under-performing children who benefit the most in terms of learning gains. While initially above median (good) private schools show no learning gains, they show large drops in fees. Parents whose children attend these schools also gain by receiving similar levels of learning quality but at a lower cost. Given that there is little school switching, this suggests that providing report cards generates credible competitive pressures on all schools to increase price-adjusted quality, with the

specific tool used – decreasing prices or raising quality – determined by the nature of production costs and market demand. Consistent with this, I find that schools increase investments, while there is little evidence for greater (direct) parental investments.

Information, therefore, enhances school competition and leads to both improved learning and lowering educational costs. These results provide a sense of the underlying market structure, suggesting that parents and schools are able to agree on relatively forward-looking arrangements (since parents do not switch schools immediately), that school quality is not only quite sensitive to school effort (rather than just school “type”), but that there may be limits to raising quality (since initially better quality schools respond to competition by lowering prices instead of quality). While still a working paper, this research has already received significant academic attention. The combination of a market-level experimental intervention that directly addresses a (informational) market failure and a comprehensive panel dataset that allows examining the impact on all market players – schools, teachers, parents, and children – enables a unique contribution in understanding how educational markets work and how their shortcomings may be addressed.

While public schools also respond to the informational intervention – showing a tenth of a standard deviation gain – the public-private quality gap remains large, likely due to the steeper performance incentives in private schools. Public schools in developing countries often lack such explicit incentives as those in the US “No Child Left Behind” policy. In another ongoing experimental intervention in collaboration with the Pakistani education department, I consider the impact of making public schools responsive to market forces by reactivating parent-teacher school councils. Preliminary results indicate that the intervention leads to more representative (of parents) councils. I am now examining how public schools and parents respond to this change, and how this interacts with the private sector and affects the market equilibrium.

C. Institutions

While economic growth hinges on how an economy leverages its human and financial capital, we now increasingly realize that this process is intermediated and directly affected by the culture, norms, and history of a region. In a more recent line of research, I examine how such “institutional capital” – in the form of beliefs, religious and cultural norms, and specific institutions – develops and interacts with the development process. While I expect these research interests to refine and gain further depth in the coming years, I anticipate two broad themes: one that examines these institutions and their impact; and the second that seeks to design better institutions and mechanisms.

Examining Institutions

Religion, particularly Islam, is increasingly forwarded post-9/11 as a set of beliefs that has profound impact on how people engage in the development process. Within the context of education in Pakistan, this was brought to the forefront due to a policy concern highlighted by the international media regarding the high and increasing enrollment in Pakistani religious schools (madrassas). In a series of papers (with T. Andrabi, J. Das, & T. Zajonc), I argue that quantitative social science tools can contribute to a better understanding of such issues. In “Madrassa Metrics” (Routledge, 2008) and “Religious School Enrollment in Pakistan” (*Comparative Education Review*, 06), I show that, contrary to popular views, madrassas in Pakistan account for less than 1 percent of enrollment with no evidence of a dramatic increase in recent years. I also demonstrate that existing and widely promoted theories are unable to adequately explain the fact that even among the

(few) households with a child in a madrassa, three-fourths send their other child(ren) to a public or private school. The paper won the George Bereday Award in 2006 for outstanding article in the *Comparative Education Review* and has received substantial policy and media attention with coverage in the *Economist*, *Foreign Policy*, *New York Times*, *LA Times*, *Reuters*, *UPI*, *Washington Times*, *Al-Jazeera*, and several other major publications.

While religious schools may factor as an important determinant of belief formation, an influential institution in Islam is the Hajj – the annual pilgrimage to Mecca that draws over two million from more than 70 countries each year. In “The Impact of the Hajj” (*Quarterly Journal of Economics*, Forthcoming, with D. Clingingsmith, & M. Kremer), I take advantage of a lottery used to allocate Hajj visas in Pakistan and find that, while participation in the Hajj increases religious orthodoxy, inducing a shift towards more mainstream practices, it is accompanied by greater tolerance. It increases belief in equality and harmony among ethnic groups and Islamic sects, leads to more favorable attitudes toward women (including greater acceptance of female education and employment), and indicates increased beliefs among Hajjis in peace as well as in equality and harmony among adherents of different religions. The evidence suggests that these changes are likely due to exposure to and interaction with Hajjis from around the world, highlighting how institutions such as the Hajj may play an important role in the evolution of beliefs. This work has also generated media attention, with coverage in the *Economist*, *CNN*, *International Herald Tribune*, and *Slate*, and formed the basis of my recently being selected as a 2009 Carnegie Scholar. In future work, I intend to use large sample surveys and experimental methods to examine what mechanisms shape and may be used to moderate the intolerant and extremist preferences of individuals and groups.

In order to understand how political and economic institutions are shaped, I am examining the consequences of the 1947 partition of British India – one of the largest and most rapid migrations in history, with over 14 million people migrating within four years. In “The Big March” (*Economic and Political Weekly*, 08, with P. Bharadwaj & A. Mian) and “The Partition of India” (with P. Bharadwaj & A. Mian), I document the population flows and its consequences on the gender, educational, and occupational composition of affected regions. In ongoing work, I am investigating how these changes, in turn, affected future political institutions and economic productivity.

Institutional Design

Better institutional design is likely to improve service delivery. In a series of papers, I examine the role of community participation and “social capital” in public good upkeep. In “Measuring Empowerment” (World Bank, 05) and “Is Increasing Community Participation always a Good Thing?” (*Journal of the European Economic Association*, 04), I use the property rights framework to model empowerment as an agent’s ownership over decision rights. The model is empirically substantiated by showing that while community participation in non-technical decisions improves public project outcomes, increasing it in technical decisions can worsen outcomes. In “Can Good Projects Succeed in Bad Communities?” (forthcoming, *Journal of Public Economics*), I use primary data on community-maintained infrastructure projects in Northern Pakistan to show that better design indeed matters. I find that while community-specific factors such as inequality, social fragmentation, and lack of leadership have adverse effects on collective success, they can be compensated by design factors such as task-complexity, inequality in task returns, and beneficiary participation.

Certain services that are typically considered the domain of the public sector – such as highways, airwave frequencies, and education – display properties of excludable public goods. Studying how a (private) monopolist can provide these goods and how they respond to underlying consumer inequality sheds light on designing alternate provision mechanisms. In “Bunching and Heterogeneity” (with J. Das), I take a theoretical approach and derive conditions under which a monopolist provides an excludable public good under a simple pricing scheme: All individuals with willingness to pay above a threshold receive the same quantity, and those below are fully excluded. This simpler contract then allows me to examine when changes in consumer inequality will adversely impact the provision of the good, particularly to the poor.

Decentralization reforms are a frequently forwarded as a means of improving public good delivery in developing countries. However, there is limited evidence on their success, partly because it is hard to isolate the impact from secular time trends and other interventions that often accompany such reforms. In an effort to do so, I first examine the specifics and political economy considerations for decentralization reforms in “Local Government Reforms in Pakistan” (MIT Press, with A. Cheema & A. Qadir). In ongoing work, I then exploit several features in the reform to better isolate their impact. Specifically, the reforms did not devolve all expenditure types – wages in a given sector were not decentralized, while other expenditure heads were. This provides a natural “control” expenditure category, allowing for difference-in-difference estimates. Initial results show that the reforms had a large impact in increasing local spending. Furthermore, an unanticipated infusion of external multilateral funds (after 9/11) in specific social services (education) reduced local politicians’ electoral return for spending in these services. If reelection incentives drive improved delivery, one would expect to see little increased spending allocated to such services, even though they are the ones typically considered to appeal to voters.

Finally, I am exploring alternate mechanisms that address coordination problems in service delivery. A particularly salient context is disaster relief. Oxfam reports that the number of natural disasters has quadrupled in the past two decades. Yet as Katrina demonstrated, even the most advanced economies may be unable to act in time. [RISEPAK](#) (Relief and Information System for Earthquakes – Pakistan), the disaster relief tool I helped design, offers an alternate solution to coordinate disaster relief. It provided searchable GIS-mapped information on the damage, demographics, and real-time relief access and receipts for the more than 4,000 villages affected. The key insight of RISEPAK was to recognize that the internet era successes show that instead of relying on a centralized coordinator and a few experts, a network of actors – individuals and organizations – can effectively, rapidly, and reliably exchange information and coordinate efforts. Widely used by relief providers during the weeks that followed the earthquake, RISEPAK was well-received and won the 2006 Stockholm Challenge Award for the most innovative Information and Communication Technology (ICT) project in the Public Administration category. Having such a system that covers the entire world would mean that within minutes of any disaster, its location, how many are affected, and how to get to the victims would be easily accessible online. Within hours, we could know what is needed, who is helping, who is being helped, and who is not. Using technology to combine the global knowledge network with local information could help inexpensively save thousands of lives. Inspired by the potential of such a tool, I hope to explore how web-based mechanisms that create a central informational marketplace based on decentralized provision can work as an effective means of sharing, coordinating, and verifying information in real time.