



## Copenhagen Consensus 2012: Education Perspective Paper

### Learning Quality as the Primary Objective

Lant Pritchett

April 2012

Professor Orazem’s Challenge Paper on Education has the big picture exactly right. At this juncture, the most productive investments will be those that increase the learning of children per year of school. Currently nearly all children do attend at least some school, but often are getting very little out of it and, at least in part because of this, do not attend and/or drop out early. Therefore investments in getting marginal children into school, while perhaps justified on moral grounds of a drive to universal schooling, are not particularly high.

However, this conclusion, that the focus on learning is the right approach, leads into the difficulty with the question “how should resources be spent?” From a purely economic perspective many of the “investment” possibilities are in fact “win-win”—*cheaper* and *produce better learning outcomes*. Activities that are “win-win” have essentially infinite economic returns. This leads directly to the puzzle of why they are not adopted—as it must mean some relevant cost to some relevant decision maker is not being taken into account. This makes the question of how to spend resources even more difficult as perhaps rather than thinking of logistical tasks by existing actors at existing institutions and incentives the highest returns are to changing ideas and institutions.

#### **I) Infinite returns available in schooling and what they imply about “opportunities”**

Let me give two examples of “win-win” policies that have ample evidence that they work at scale. This leads to some reflection on what an “opportunity” is and suggest that the current exercise is, while headed in the right direction, not bold enough.

##### **I.A) Contract Teachers**

An extreme illustration of the infinite economic returns comes from recent research (Kingdon & Atherton, 2012 (forthcoming)) in two states of India, Uttar Pradesh and Bihar with large samples of students and with repeated measurements of student learning across grades. Their estimates

of learning gains, which are plausibly identified even if not experimental, show that contract teachers, hired on one year contracts and hence subject to at least some sorts of “high powered” incentives in Uttar Pradesh produced roughly twice the learning gain per year of the regular, civil service, teachers. That is, students learned in one year with a contract teacher the equivalent of two years worth of gain of students with civil servant teachers. In 2009 the contract teachers were making roughly 3000 rupees a month whereas the regular teachers were making around 11,000 rupees a month.

What is the cost-benefit of replacing civil service teachers with contract teachers in this case? Crude calculations (which suffice for this purpose) suggest that the savings would be roughly **billion dollars** (around 600,000 primary and upper primary teachers times 8000 rupees monthly wage differential (11000 less 3000) times 12 months divided by 45 rupees/\$). And for this billion dollars in cost savings learning per year would double. At the margin the benefit-cost ratio of hiring a teacher on contract instead of civil service is infinitely large as benefits go up and costs go down (even over the very short run as this is a feasible option even with no incremental training or recruitment costs).

What did the government of UP actually do? Of course they *raised* the wages of civil service teachers (as part of an overall pay hike to government workers in the Sixth Pay Commission) and hired *more* of them. Kingdon’s current calculations, updated from 2009, suggest that the replacement of civil service with contract teachers would save UP eight billion dollars, not just one billion.

This finding of infinite returns with contract teachers in India is not a fluke of this particular study or its methodology. (Muralidharan & Subrararaman, 2010) evaluate experimentally the impact on learning of an extra contract teacher in Andhra Pradesh (in assessing the relevance of evidence from states of India that Andhra Pradesh is larger in population than Germany) and find that an extra contract teacher raises learning by .1 to .13 effect sizes (student standard deviations). Moreover, when they evaluate the gains to contract teachers compared to regular teachers they find the gain is roughly the same (usually slightly higher for contract teachers, but not statistically significantly so). In their experiment the contract teachers made *one fifth* the pay of regular teachers.

Again, on a cost benefit ratio there is an infinitely high return at the margin to hiring a choosing a contract teachers over a regular teacher as the output is the same or higher and the cost is much less (and not just on salaries, as the contract teachers have less education and less training, both of which are costly, and still do as well).

These findings are perhaps a bit specific to India (and similar South Asian countries) since (a) the wage gaps between regular teachers and the market wage for equally capable teachers have reached such astronomic proportions and (b) the dysfunction of the civil service cadre has reached such dispiriting levels (evidence by such high absenteeism and lack of effort).

But recent evidence on contract teachers in Kenya finds similar results Duflo, Dupas and Kremer (2009). Here the contract teachers were of the same qualifications and paid the same as regular teachers and hence just contractual status varied. In this case the learning gains were much larger for students exposed to contract teachers than regular teachers. In fact, in assessments a year after the intervention (to account for learning “depreciation”) the students who were randomly assigned an extra contract teacher had learning .19 student standard deviations higher than the control group while those assigned a civil service teacher were .011 units *lower* (obviously not statistically significant). Amazingly, and not amazingly if one is cynical enough, the addition of an extra teacher caused the existing civil service teachers to attend *less* so that the net teaching from an additional teacher was much less than the net impact of adding a teacher.

Again, at the margin, the return to contract teachers (with big learning effects) over civil service teachers (with zero learning effects) at the same economic cost is *infinitely large*.

Turning back to India, not only are these returns infinitely high at the margin but they are also massive in total. Suppose, that spending on education in India is 4 percent of GDP and primary education is 40 percent of that and teacher wages are 80 percent of spending then teacher wages in primary schools are roughly 1.25 percent of GDP. If these costs could be reduced by a factor of 4 with no reduction in quality then there is a full 1 percent of GDP to be saved. Suppose this is India alone, this is still a savings of 18 *billion* dollars—with no loss in learning or educational output.

### **I.B) Private Schooling in Pakistan (and elsewhere)**

A team of researchers has done a massive exercise in studying education in Punjab, Pakistan, tracking students and schools for years. One of the main things to emerge from this study (Andrabi, Das, Khwaja, Viswanath, & Zajonc, Pakistan: Learning Achievement in Punjab Schools, 2009) is that low-cost private schools are both much *cheaper* in terms of total expenditures per child per year of schooling (the costs in private schools were about half of government schools) *and* provide *massively* more learning per child than government schools (on a reasonably, if not experimentally, identified causal estimates)—by .7 of a student standard deviation in Maths, for instance ( (Andrabi, Das, Khwaja, & Zajonc, 2009).

Already parents are voluntarily switching from public to private schools (and there is higher reported parental satisfaction with private schools over government schools). At the margin even small incentives would be sufficient to move children from public to private schools. Suppose the government took the Rps 2000 it was spending to education a child in a public school and instead paid a voucher that covered the average cost of a private school (Rps 1000 in the LEAPS sample). This would have the effect of *reducing* costs and *increasing* learning substantially.

Again, at the margin this has *infinite* rate of return as costs go down and learning goes up as the private sector spends Rs 1 per percent correct while the public sector spends Rs 3.

While the “global” literature on the causal impact of private schools is “mixed” (meaning that in well functioning and high state capability countries like the USA the gains are small) there is little doubt that in countries with poorly functioning states there has been a massively voluntary shift into private schools even with no policy encouragement, which is suggestive of much higher quality. Evidence from India is consistent with large learning gains from shifting to the private sector with plausible identification—especially in low performing states of India.

But suppose that even just in Pakistan and India (over 1.2 billion people) one could reduce costs of primary schooling by half and have equal or better outcomes. The cost savings alone exceed the wildest hypothetical of what the additional money the world would spend on addressing “global problems.” That is, back to our earlier calculation of what could be saved, between just India and Pakistan about 30 billion dollars. This is not additional expenses but *savings*. Given that one way of expressing the previous version of the Copenhagen Consensus was “how to spend an additional 50 billion” across all global problems, in the education area just in two countries we can identify *savings* of 30 billion with equal (likely much better) outcomes.

### **I.C) How to “cost” opportunities**

The difficulty of course is in deciding, what is an “opportunity”? In education it is always politically feasible to do incrementally more of the same. Especially if governments are provided more money for it they will build more schools, hire more teachers, buy more chalk and textbooks, train more teachers. That is, they will willing do BAUWMM—“business as usual with more money”—and it is an interesting conjecture to ask whether if among the clearly politically feasible opportunities in the BAUWMM set some have higher impacts (in enrollments or learning) than others.

But if one is imagining spending 50 billion dollars to improve the human condition, why take current politics and hence existing political feasibility as a “hard” constraint? Spending money on advocacy that creates new political possibilities might well be fantastically more cost effective than anything in the BAUWMM set. Imagine what I could do with a billion dollars spaced over ten years (100 million a year) just to promote the idea and hence raise the political feasibility of contract teachers or private schools or other cost-saving/learning achievement raising win-win options. I could set up institutes around the world to do more research, I could promote the research findings in popular media, I could hold conferences with policy makers, I could fund exemplars around the world, I could do lots of things. Maybe I would succeed, maybe I would not.

Suppose one does the rate of return to spending a billion dollars for policy advocacy for contract teachers in India over ten years followed by, after ten years, some probability of success which would mean 18 billion dollars in gains for each of the next 30 years (relative to the counterfactual). If this would work for sure the internal rate of return is 68 percent. But, suppose the odds are only 50 percent it would work, the IRR is still 57% (because after all we are still only

investing 100 million a year to then get 9 billion a year). Even if the probability of the success of the advocacy falls to just 1 in 10 that spending a billion on advocacy will reach the goal, the IRR on such advocacy in expected value is still 34 percent.

The problem is that there are two separate domains “what will get done” and “what the impact on outcomes of what gets done will be.” Right now there is seemingly more and more attention to the latter. There is a drive for more rigorous evidence about cause and effect relationships between programs/projects/policies and outcomes. But it is not at all obvious to me, or any careful observer, that the latter plays an important role (or much of any role at all) in the former.

Advocacy seems to play a much larger role in what gets done, particularly in the development sphere, than evidence. For instance, there was a massive expansion of US Foreign Assistance under President Bush—most of which was accounted for by a 15 billion dollar program over five years devoted to HIV/AIDS. While this may have been the optimal allocation of resources (or may not have been) as measured by some objective metric of improved human well-being per dollar, it also may not have been. But importantly, no one really cared. The decision was not made on the basis of evidence about the relative merits of investments in health versus education versus roads nor even on the evidence of the relative merits of investments in health between day, clear water or pre-natal care and HIV/AIDS nor even on the evidence of the relative merits of spending on prevention versus treatment. This is not to say that there was not some kinds of evidence on these various topics put forward, just to say in no one’s causal narrative of “what got done” was compelling evidence the operative explanation.

Not to mention the rise of many other issues in the development agenda, e.g. the environment, gender, or even how some issues come, go and then come back—like the importance of “infrastructure.” There is no question that advocacy plays a role in what gets advanced and hence “what gets done” on the international agenda, by creating persuasive cases to core constituencies and stakeholders. The question is the relationship between “persuasive” and “evidence.”

In part because I am so well trained at the production and evaluation of evidence as construed by a particular academic discipline it has taken me a long time to escape from the myth that what I am good at is also important. That there is no (reputable) “science” or “discipline” of “what gets done” does not mean that there is not a way in which things get done. Moreover, it does not mean that there are people out there engaged in getting things done and that those people are often better funded, cleverer, and overall just better at getting things done than are academic experts.

This also means that potentially the highest return activity could be the creation of a persuasive case for the adoption of good “policies” or good programs—knowing that “persuasive” and “evidence based” need not be the same. There are “successful failures” in which it is possible to get something done politically even though its cause-effect efficacy and accomplishing the

putative objectives is zero. Prohibition in the USA in the early 20<sup>th</sup> century is a great example. They were massively successful in doing what is extremely difficult: amending the US Constitution. However, they wanted Prohibition because its advocates believed the *impact* of Prohibition would be certain desirable goals—like stronger families, less domestic violence. However, they proved to be much better at advocacy than at social science and their means did not accomplish their ends, which led to the end of Prohibition by re-amending the US Constitution.

But there are also successful successes and it only takes a few successful successes to justify enormous amounts of spending. For instance, Professor Orazem suggests that reforms in the economic climate in which education was used would have enormous returns. This again falls into the space of the “political feasibility” assessment. If these reforms are “win-win”—good for the country’s economic growth and also good for raising returns to schooling—why are they not already adopted? The answer is, tautologically, because the people who have the power to adopt them have not done so, and only slightly less tautologically, have not done so either because (a) they are ignorant (e.g. are unaware of means to ends actions that would promote their own interests) or (b) it is not in their interests. Rather than glossing over this as he does as kind of a result to be pointed out, but not really an “opportunity” lets ask what it would be cost effective to do to achieve this result.

Take the case of the economic reform in India in the early 1990s. India has been growing rapidly in the 1980s but, as many other countries had, hit an impasse in 1990/1991 as a delayed devaluation combined with a controlled, anti-competitive, anti-outward oriented economy was taking India into a potentially severe macro-economic crisis. The Latin American countries which hit similar crisis after decades of growth in the early 1980s saw their growth drop to near zero for a decade or more. In India, this incipient crisis was used as an opportunity and the macro aspects were handled well (a devaluation restored external balance very quickly) and in addition a set of reforms in trade (lowered barriers in the most restrictive import regime in the world at that time), regulation, and in the financial sector were implemented. Of course one can debate cause and effect but the fact of the matter is that after these reforms the economy quickly returned to a rapid growth path and has persisted in rapid growth (accelerating even further prior to the global crisis of 2008) until today (maybe that would have happened anyway, but maybe not).

Let us do the simple calculation of the gains of some simple counter-factuals about Indian growth. What if a poor handling of the macro crisis had led to stagnation in growth? Or more modestly, if India had returned to its 1970s level of growth? Table 1 shows the resulting PPP GDP per capita from the actual trajectory up until 2007 versus the counter-factuals of zero growth or 2.2 ppa growth since 1991. The simple math is that at 2.2 ppa per annum growth GDP per capita would be roughly 1,000 PPP\$ per capita lower and India has roughly a billion people so total GDP would be a *trillion dollars* lower *per year*. If mishandling of the crisis has led to

zero growth, as in the Latin American experience, then GDP would be lower by *two trillion* dollars *per year*.

If we are thinking about global welfare and well-being then it is perfectly legitimate to adjust those gains to OECD equivalent dollars by adjusting for declining marginal utility as the average GDP per capita in India over this period was 2,600 PPP\$ compared to 26,000 PPP\$ in the OECD. If we adjust for marginal utility at log utility this means the welfare gains from the difference in the actual growth rate compared to the downside counter-factual of zero growth are larger than US GDP—20 trillion dollars.

Table 1: Gains from return to rapid growth after crisis in India in 1991 versus downside scenarios		
	GDP per capita at 0 ppa growth since 1991	GDP per capita growth are 2.2 ppa growth since 1991
GDP per capita in PPP in 2007 (actual): 3,826	1,964	2,782
Loss in total annual GDP in 2007 from slower growth since 1991 in PPP\$ <i>trillions</i>	2.10	1.18
Gains adjusted for higher marginal utility of consumption with log utility (coefficient=1)	20.64	11.57
Coefficient=2	202.48	113.49
Source: Author's calculations using PWT6.3 data for GDP per capita. Memo: US GDP in 2007: 14 trillion		

There were international organizations like the World Bank and private foundations that were supporting think-tanks that brought together economists and supported them in the research and advocacy to create the intellectual climate in which the costs and benefits of alternative growth strategies—including more market-oriented—were debated. This small group, challenging the conventional wisdom at the time of state-led growth, eventually got more and more intellectual traction. I do not want to tell any simplistic cause and effect story that this “caused” the reforms in India, but it is hard to believe it played no role. Moreover, many of the key actors responsible for the reforms spent significant time in these organizations and think-tanks, essentially preparing themselves for the time when the opportunity would come. Let us suppose that all this contributed only 1 percent to the reforms and suppose (I agree, debatably, but suppose) the reforms accounted for sustained growth versus reversion to mean growth of 2.2 percent. How much would it have been optimal to invest in this research and advocacy for economic policy change *ex ante*? Even one percent of 11 trillion (welfare gains at log utility by 2007) is 110 billion a year. In the 1980s the *total* operating budget of the World Bank was on the order of 1 billion dollars a year—obviously even if *all* of that had been spent on nothing but creating the possibility of reform in India this would have been dramatically under-investing in the creation of conditions propitious for successful handling of the crisis.



One might think I have wandered far from the economics of the opportunities in education—but I have not. The point is that if “opportunities” are defined as “incremental ways of spending money that are currently politically feasible” then one will get a certain array of mostly programmatic answers: e.g. nutrition programs, conditional cash transfers, targeted inputs (e.g. girl’s toilets). One could ask which of these is the most cost effective and hence what fraction of some incremental available resources should be devoted to that—and come up with numbers in the, say, few billions.

Total education spending by developing country governments around the world is on the order of 800 billion dollars now. There is no question that much of this money is spent in extremely inefficient, if not absolutely wasteful or counter-productive ways. Suppose, that through some combination of new information, research, and advocacy one could expand the set of the politically and administratively feasible ways to do education that improved the effectiveness of existing spending by say, a tiny fraction, say 1 percent. That is an equivalent of 8 billion in gains per year. What is the “opportunity” there?

Of course one could say, “we have science of ‘what works’ but do not have a science of ‘what gets done’ and hence we should stick to science.” That is true if one’s objective is exclusively science for its own sake. But in terms of improving the human condition this is the classic response of the drunk searching under the streetlight. So far, economists have been very stubborn drunks.

## II) Minor quibbles with the challenge paper

Three technical comments:

First, I am glad Professor Orazem was clear about “conditional cash transfers” as an *education* intervention. If the government is *already* making a cash transfer the incremental costs and benefits of making that cash transfer conditional on school attendance are almost certainly worth it as the *incremental* cost is only the cost of adding the C to the CT and the bulk of the cost of the CT is “sunk” in this calculation.

However, if one does a CCT and includes the entire cost of the CT itself then I do not think that in most circumstances this makes for an attractive education intervention as the cost per additional enrolled student are very high. For instance, take the Colombia results in his table 3. Households received 17 percent of pre-transfer income as the transfer at child time in school for ages 8-12 went up 2 percent. There is just no way this works out to be cost-effect compared to other interventions in the education sector as nearly all of the Cash Transfer went to infra-marginal households (that is, households that would have had their children in school anyway). In my mind a CT has to be mostly justified on its CT merits (e.g. transferring purchasing power to people with low incomes) and the C in the CT against the incremental cost benefit of adding the C.

Second, I just disagree with Professor Orazem’s evaluation of fee elimination in Sub-Saharan Africa, particularly with the view that it did not decrease quality. This all depends on what was done with the fees before their elimination. The only empirical evaluation I have seen of the impact of fee elimination on quality is for Kenya and there the evidence is pretty compelling that fee elimination led to massive deteriorations in, at the very least, parental perception of school quality. In Kenya the fees had been locally collected and controlled and used at the school level. Their elimination plus replacement with a central transfer was by and large a disaster. As (Bold, Kimenyi, Mwabu, & Sandefur, 2011) show Kenya eliminated fees in primary by not secondary schools and nearly all of the incremental enrollment in primary schools happened in *private* schools while most of the incremental enrollment in secondary schools was in *government* schools. This differences in differences seems pretty compelling that abolishing fees led to massive flight out of public schools even though their money price had fallen which is only consistent with a very substantial reduction in perceived quality.

Several of the countries on the list in Table 1 of having eliminated fees have also been participating in the SACMEQ testing and show pretty substantial deteriorations in learning achievement at grade 6. For instance on a SACMEQ norm of 500 Mozambique fell from 516 in 2000 to 476 in 2007, Malawi fell from 462 in 1995 to 433 in 2007. Of course we don’t know cause and effect and this could just be compositional shifts as less prepared students come in—but the case that this does not deteriorate quality has to be made case by case.

The following table is pretty sobering. In Tanzania (reported as eliminating fees in 2001) the number registering for the secondary schooling leaving test from 2007 to 2011 more than doubled but the absolute number getting division 1 to 3 marks actually *fell* by more than 10,000 students (almost a quarter of the baseline total) as the percent of exam takers getting good marks fell by a third. Hard to say what is test comparability, cause and effect, and so on, but when only 8 percent are passing the leaver’s exam it is hard to be sanguine about quality or the impact of increased enrollments on quality.

Registered for CSEE	Sat CSEE	Scores in Divisions 1-3	Share of exam takers getting divisions 1-3	Year
<b>199,283</b>	189,398	44,567	23.5%	2007
<b>241,472</b>	233,848	41,915	17.9%	2008
351,152	339,925	42,790	12.6%	2009
458,114	441,426	40,807	9.2%	2010
450,324	426,314	33,869	7.9%	2011

Third, and much more minor. The assessment of whether or not there are “positive externalities” to schooling has to be balanced against the fact that nearly everywhere and always schooling is massively subsidized already. So the standard welfare calculation of an *increase* n

the subsidy for schooling has to be based on an assessment that currently at the margin the externalities exceed the existing subsidy. So while the academic debate is about whether there are positive growth externalities *at all* (and both he (table 2) and I find that in the median country there are not) the policy question is about *increasing* the subsidy from its already massive level of typically providing school free of charge already.

### **III. An actual proposal on teachers**

Since in comments on the previous version of Professor Orazem's paper (what is still section I) I emphasized the infinite returns to contract teachers he made my point by invoking otherwise unspecified "political" constraints: (a) that contract teachers and regular teachers cannot work side by side over the long term as eventually they will be regularized and (b) whatever long-term arrangement for regularizing teachers will then just produce the same outcome as regular teachers. The most obvious response to that is that one should still do contract teachers whenever one can (where they are win-win on learning and cost) for as long as one can and if it is not "sustainable" one still gets the gains as long as they last so what is the harm from a ROR sense?

On a deeper level I have thought of, and thought through, all of his objections to why "contract teachers" are not a viable proposal for education. In a paper some years ago (Pritchett & Murgai, 2007) we laid out a plan that tried to get the benefits of contract teachers on a sustainable basis. The basics of the plan were:

- The old cadre of the "regular" civil service teachers would be eliminated so that no more appointments could be made under those employment conditions (of course all existing teachers were grandfathered)
- The new cadre of teachers had an entirely different system of compensation and, importantly, assignment, which had four features:
  - A long probationary period of wages equivalent to those of contract teachers. So, rather than beginning at a wage that is fivefold the market wage from day one all teachers have to be "apprentice" teachers at the much lower wage for an extended period (five to seven years).
  - Confirmation as a "tenured" teacher in the new cadre came only after a three-fold criteria was met: (a) a local school had to demand you as a teacher, (b) your performance as a teacher had to be documented administratively (e.g. attendance, training, etc.) and (c) a peer review had to certify your quality as a teacher.
  - At tenure the wage goes up substantially so that over a life-course teaching is an attractive occupation.
  - The new process separates "hiring" from "assignment" so that to be appointed as a teacher in any school requires that one meet the eligibility criteria but to be in a classroom the local school committee has to approve

your appointment. That is, the *assignment* of teachers is controlled at the local level not at the bureaucratic level. During their probationary period teachers not assigned do not get paid (so this is essentially the contract teacher model). So you can only get “tenure” after X years of teaching but you can only teach if you can find an assignment (over which local committees hold veto power). The assignment of teachers with tenure is more complicated since after tenure the district has a obligation to pay irrespective of assignment, but tenured teacher get multiple chances at assignment until their pay is gradually reduced if they are not teaching.

This has several advantages as (a) the gains to lower pay and higher accountability are in full force for all non-tenured teachers and (b) those who make it through to tenure are likely those more likely to comply even once the “high powered” threats are removed. With “normal” turnover patterns for teaching most teachers will be in the “probationary” period and if the filters work well only “good type” teachers make it through to the tenured condition. This gives time for the “norms” of performance to change.

Something very much like this plan was adopted in Bihar India and the same evaluation as quoted above for Uttar Pradesh also found that the “apprentice” teachers in this new plan were both cheaper (though not as much cheaper as pure contract teachers, unfortunately due to political pressures) and their students had higher learning than regular teachers (though not as much difference as the “high powered” threat was apparently less high powered). While it is still too early to tell if in the long-run this reverts to the norms of non-compliance this is at least an administratively and politically feasible plan that lowers costs and raises performance simultaneously. At the margin the returns are infinite (compared to the counter-factual of hiring regular teachers).

Of course the total gains depend on how much the system is expanding and how much turnover there is in the teaching force. In mature systems with little turnover obviously the incremental gains are small as all teachers have to be grandfathered so even though the marginal returns are high the total gains are low. But in expanding systems or systems with turnover the move to a newly designed system of teacher compensation and assignment (and the assignment bit is key as it creates the local accountability) can produce both high marginal and large total returns.



## **Bibliography**

Andrabi, T., Das, J., Khwaja, A., & Zajonc, T. (2009). *Do Value Added Estimates Add Value? Accounting for Learning Dynamics*.

Andrabi, T., Das, J., Khwaja, A., Viswanath, T., & Zajonc, T. (2009). *Pakistan: Learning Achievement in Punjab Schools*.

Bold, T., Kimenyi, M., Mwabu, G., & Sandefur, J. (2011). *Did Abolishing School Fees Reduce Quality? Evidence from Kenya*.

Duflo, E., Dupas, P., & Kremer, M. (2009). *Additional Resources versus Organizational Changes in Education: Experimental Evidence from Kenya*.

Kingdon, G., & Atherton, P. (2012 (forthcoming)). The relative effectiveness and costs of contract and regular teachers in India. *Economics of Education Review*, 1-15.

Muralidharan, K., & Subramanian, V. (2010). *Contract Teachers: Experimental Evidence from India*.

Pritchett, L., & Murgai, R. (2007). Teacher Compensation: Can Decentralization to Local Bodies Take India from the Perfect Storm Through Troubled Waters to Clear Sailing? *India Policy Forum*, 123-168.