COPENHAGEN CONSENSUS 2008
EDUCATION
PERSPECTIVE PAPER

COPENHAGEN CONSENSUS CENTER

Copenhagen Consensus 2008 Perspective Paper

Education

By

Lant Pritchett
Kennedy School of Government, Harvard

As a development economist I love the idea of the Copenhagen Consensus

As an economist, I naturally love the idea of comparing possible policy actions across domains to judge their attractiveness. Too often public policies are only discussed by advocates/experts in particular sectors, so educators discuss education, doctors/public health experts discuss health, transport engineers discuss highways. The result of this what I find nearly every time I get my hair cut—the people who cut hair are more interested in my hair than I am and often recommend I "do more" to my hair. This is not surprising as one would expect, for a variety of reasons, people who devote their careers to hair care to care about hair. So naturally if one gathers together experts in any field the discussion presumes "more should be done" in their sector and the only discussion is which of the many things in their sector are highest priority (and how to convince the rest of the world, mostly not true believers, to cough up the cash). Since economics is a theory about resource allocation (and since we are short on technical expertise in the specific sectors) a cross sectoral comparisons of allocations of public sector resources—weighing the gains to resources devoted to global warming against those devoted to disease eradication or improving nutrition—brings economics to the fore.

As a *development* economist I love the idea because comparing alternatives, some of which affect poor people and some of which affect rich people, is the only way to highlight the stark consequences of the existing distribution of wealth, productivity, and income across the world. Spending (and of course imposing regulatory costs that reduce output is the equivalent of spending) a dollar of public resources has the effect of reducing the amount available for private consumption. In the richer countries of the world there might be the sense that reducing consumption to devote to public purposes has a low cost, because what would have been consumed is not that important—it would have been spent in crass materialistic, keep-up-with-the-Jones, advertising driven consumerism anyway so "we" might as well take it away and spend it on what is truly important like protection the environment or improving schools or raising incomes of the poor. However, recently Banerjee and Duflo (2006) examined household surveys to create a statistical profile of "the poor"—and of those living on $2/day 2/3 of every dollar was spent on food—and this standard of poverty is roughly the median in poor countries. So, in thinking of devoting resources to global warming or education or health if the alternative is a dollar's worth of private consumption for the median poor country citizen, roughly 2/3 of this would go to spending on food.

The genius of the Copenhagen Consensus exercise is to create a feasible, evidence based, alternative for the use of incremental resources that takes into account both the effectiveness of resources across sectors and across the globe. This is a useful corrective to the confusion of small percentages—the costs of attacking global warming might be only 1 percent of GDP in a rich country. One percent sounds small, and one could easily think of ways to reduce consumption by one percent that must have trivial consequences for well-being (although this is usually more of one's neighbors than ones self). But one percent of high income country GDP is (in PPP$) roughly 320 billion dollars. The total poverty gap in the world (the difference between actual consumption of the poor and the "dollar a day standard" is roughly 35 billion dollars (in 2001 PPP)[1]—an *order of magnitude* smaller.

To preserve the logic of the exercise one has to use the hard headed logic of economics all the way through. That is, "beating something with nothing"—comparing the costs to an unspecified alternative--is difficult but maybe we want to compare foregone losses from attacking global warming, which generates a great deal of concern in rich countries, with something else that people can get warm glow from—like programs to attack malnutrition or educate girls. But if we are to use economics to compare these alternatives, we must really use economics. The essence of a normative "cost effectiveness" exercise for public sector actions for an economist is to start from some objective function (usually an inequality averse (or at least neutral) aggregator of underlying utilities (self-assessed individual well being) and from a model in which the choices of individual agents produce an equilibrium outcome. Then we ask, relative to the counter-factual of the existing equilibrium (which includes all existing public sector interventions) by how much does our objective function increase from public sector action A with cost (direct or indirect) P.

The problem with comparing public sector actions across two sectors, say global warming and education, is that it is *not* the efficacy of "cost effectiveness" of spending in the two sectors that should be compared, particularly of course if one is a pure public good (non-rival and non-excludable) and one is a pure private good (rival and

[1] These are just simple calculations from Chen and Ravallion (2004), there are 1089 "dollar a day" poor, the poverty gap is 6 percent of the poverty line, times 365 (to get to annual) times an adjustment for inflation to 2001 dollars and the total is 35,191 millions.

excludable). We economists have two welfare theorems about that. Under a very stringent set of assumptions we know that governments cannot tax money away from people (and any regulatory imposition can be thought of as an equivalent) and give it back to them in private goods and make everyone better off—since if spending on X were a "priority" then that would have already been reflected in the choices of households. Of course the government could "improve" the distribution of income—but the second welfare theorem tells us something like that this need not involve anything other than redistribution of cash (no sectoral specific interventions required).

Of course this is not true when there are "market failures"—so it is possible for the government to tax money away from people and then spend that money in ways that makes everyone, by their own evaluation, better off. The classic examples are pure public goods, which are non-rival and non-excludable which means no profit maximizing producer will produce them.

It is with considerable trepidation I lay out this simplistic account of normative public economics to an august panel of economists, but my key objections to this paper lie not so much in the empirical evidence they present, but rather what to make of it for policy. In my view, there are three fundamental issues with this paper, and issues that are relevant to comparing there results to something like investing in global warming which is, at this stage of human history, something like the mother of all public goods (the quantitative magnitude of the returns can be debated but no one questions that greenhouse gas emissions are a global public good/bad and hence their reduction a legitimate matter for public policy).

First, it does not grapple sufficiently with the fact that education is predominately a private good.

Second, it does not present a compelling positive theory of schooling that is consistent with a welfare theoretic interpretation of their results.

Third, they do not consider alternatives *at the margin* and *one by one*—their results with claim to be highly cost-effective with "demand side" transfers mostly combine two policies—and it is the inessential one that actually produces all the gains.

And let me preface everything I am about to say with the note that this is an excellent paper of its type, these three issues are, more or less, endemic to this branch of the literature.

I) Education is predominantly a private good

The authors make things very hard for themselves right up front. Lets think of what I would want to demonstrate if I were to justify public sector spending on schooling.

Since the optimal Pigovian tax/subsidy on a specific item depends on the *difference* between private returns and the public returns (which include the private returns). "Having an externality" is not like virginity, it is not a discrete, in fact, unless all markets are perfect then likely nearly all goods have *some* element of an externality—if only in using a good which is a pure public good in production. The question is the *magnitude* of the externality compared to the *magnitude* of the private benefits.

The authors make their life very difficult by starting with table 1 which shows there are very large *private* returns to schooling—on average across the 42 developing countries which they estimate an additional year of schooling is associated with an 8.2 percent increase in wages for men and a 10.3 percent increases in wages for women, and a 9.2 percent increase in urban areas and 8 percent in rural areas. Why does this make their life difficult? Because if I really want to justify *public* spending on something I want the *private* gains to be small and the *public* gains to be large. But what they show is that the *private* gains are considerable, which means that, if, say, one were to justify public spending as a large proportion of the private spending then one needs very large externalities to education.

One way into that question is to ask: *if* the common policy of (near) complete subsidization of all instructional costs of schooling *were* to be justified exclusively on the basis of externalities to schooling, how big do those externalities need to be (measured in a way consistent with the Mincer return) [2]? To calculate this I assume the standard Mincer framework that wages are a function of experience, its square, and years of schooling, a 45 year working life, 15 percent tax rate, and a discount rate of 11.5 percent. Why a discount rate 11.5 percent?—because that is the discount rate at which a 15 year old would choose to complete ninth grade at a Mincer wage increment of 9.9 percent if the only cost to the individual were the opportunity cost of the foregone wage. Now to calculate instructional costs of primary school we assume a teacher wage based on 15 years of schooling (12 plus three years teacher training) at the Mincer return of 9.9 percent and 20 years experience (with a 2.5 percent experience premium and a quadratic term such that experience premia peaks at 25 years). We explore a range of class sizes to get per student cost and assume that teacher wages are only 60 percent of total instructional costs (as construction and maintenance costs of the buildings, plus administrative costs, plus all instructional materials need to be included). I assume secondary school instructional costs are 50 percent higher than primary. These assumptions give an estimate of instructional costs that has the main virtue of consistency, with a *patina* of plausibility.

[2] These assumptions draw on Heckman and Klenow (1997) who do a similar calculation for college costs in the United States, with the result that the externality would need to be about 3 percentage points to justify the instructional cost subsidy at a typical public university and the calculations are reported in Pritchett (2006).

In this simple framework the "externality inclusive Mincer"—the impact of a year of schooling on aggregate well-being with all externalities monetized—should exceed the micro-Mincer—the impact of a year of schooling on an individual's wages—by between 3.5 to 6.5 percentage points at the primary level (because opportunity costs are low, instructional costs are a higher fraction of the total costs) and 1.7-3.3 percentage points at the secondary level (since opportunity costs are higher, instructional costs are a lower proportion of total costs).

Table 1: What rate excess social over private rate of return to schooling would rationalize full subsidization of instructional costs?

	I	II	III
Teacher years of schooling=	15	15	17
Teacher experience=	20	20	20
Primary Class Size= (secondary assumed half as large)	30	40	25
Age 6/ Schooling=0	**4.5**	**3.4**	**6.5**
Age 10/ Schooling=5	**2.2**	**1.7**	**3.2**
Age 15/Schooling=8	**2.3**	**1.7**	**3.3**

The calculations assume the only private cost is wage foregone, a working life of 45 years, 15 percent tax rate, and 11.5 percent discount rate. At these assumptions a 9.9 percent wage increment is sufficient to induce a 15 year old to complete a ninth year of schooling at zero instructional cost.

Source: Pritchett 2006.

Most of the discussions of the returns to education are like that in the challenge paper: (a) there are large *private* returns to schooling and (b) there are plausibly *some* externalities to schooling (e.g. economic spillovers, health effects, reduced crime, etc.). But that there are *some* externalities to education only justifies *some* subsidy. It is already the case that most governments in the world offer highly subsidized schooling (at all levels in fact)—usually free primary schooling. In fact, just take the fact that public spending on education is roughly 4.1 percent of GDP in lower and middle income countries suggests that roughly 500 billion is already being spent on education in poor countries.

Therefore the question is not whether there are externalities or not, we are not looking for *some* externality we are looking for *huge* externalities. Just a simple example, suppose that in a country an unskilled worker, with zero years of schooling, supports a family of 3 (herself and two others) at the "two dollar a day" standard and hence earns PPP$2,190 a year. With an 8 percent wage premium a worker with 6 years of school would make P$ 1285 more and with 10 years of schooling P$2500 more. These differences should be glaringly obvious—and they are. As the authors show, that workers with higher education on average have higher earnings rivals Engel's curve for the best and most widely documented empirical regularity.

By the same token, if there were externalities of the equivalent to a 4.5 percent social return then *each* worker with a primary degree (5 years of schooling) should contribute P$539 of benefits to the economy (over and above the wage returns)—30 percent of the unskilled workers total earnings. Each worker with junior secondary (9years of schooling) at a 3.5 percent social return (weighted average of 4.5 and 2.3) should return P$794 of benefits to other people in the economy.

Effects of this magnitude should be easy to find. But on this section (pp 11-13) of "why don't parents choose the optimal amount" there are just two arguments. One is credit constraints (on which, more below) and the other is "market failures" for which we get the usual suspects—educated women have fewer children (see important footnote below)[3], make markets more efficient—and a heartening acknowledgement from the authors that most of the benefits often cited as "externalities" are in fact internal to the household (e.g. healthier children) and so are not very "external." But no one doubts that there are some externalities. But do the externalities exceed the existing subsidies? This is important since below the authors will argue for "demand side" actions that *increase* the subsidy for enrollment. If the subsidy is already too large (e.g. crudely, the subsidy as a proportion of total costs already exceeds externality as a proportion of total benefits) then these demand side transfers—even if effective in raising enrollments—are welfare worsening as they exacerbate the welfare losses from the existing distortion. This would be like subsidizing driving when there are net externalities to driving over other modes of transport.

No one likes to hear this, but there has never been a single case in which it was empirically demonstrated that the externality effects of schooling were sufficient to justify "free schooling" as optimal normative economics—much less free schooling plus. I have recently written a review of the evidence from aggregate output data on the evidence for a positive spill-over of education on aggregate output (over and above the private wage effect) for the Handbook of Education Economics and I fail to find any evidence at all for a positive externality to schooling in output/productivity—much less evidence of one the magnitude needed to rationalize full cost subsidy.

Perhaps the authors could justify an excess of social over private returns of sufficient magnitude, but they have definitely laid out the hardest case—that private returns are universally high and external returns are not quantified at all.

---

[3] Since the authors return again and again to the example of lowered fertility as an external benefit, one has to make it clear that this depends on the view that children are like littering—private benefits but *impose net negative social costs.* I would be happy to see this "children are pollution" position defended as I do not believe it, nor do I believe there is any rigorous evidence for this view, and there are certainly counter-examples (e.g. positive externalities to economic agglomeration).

Now, there are many other ways that spending on primary education can be justified—that education is a universal human right, education is a merit good, the demands of political socialization demand universal education. I suspect that the actual positive theory of education has more to do with those than with the economic returns. But for the purposes of the present exercise of comparing alternative uses of public funds across sectors one cannot invoke "human rights" as a reason to spend on schooling without a counter of "intrinsic values" of an unchanged natural environment, which returns the debate to non-quantifiable values.

II) Positive Theory of Schooling

The most interesting part of the most interesting table in this paper gets almost no mention. What do you think the most important reason given for why children who are not in school are not attending school?

The "supply side" view was that students lack access—cannot get to the buildings. While that may have been true some many many years ago, and probably remains true in some sparsely populated regions of very poor countries—the average is 1.9 percent in Urban areas and only 4.9 percent in Rural. This is consistent with evidence from Filmer (2004) on the lack of importance of distance and with the importance of drop-out as a cause of the lack of schooling completion.

The "demand side" view is that children drop out of school because they either (a) have very attractive alternatives or (b) are credit constrained. Let's add up as reasons (saving for later how to parse these responses) the answers "work outside the home" "housework" and "poverty"—we get 32.9 percent in urban areas and 33.9 percent in rural areas. This is at least plausible. But those three *together* are not as big as the most common reason.

The most common reason given why children are not attending school is "lack of interest"—which is 47.3 percent of urban children and 44 percent of rural children not in school. Why do these children "lack interest"? This is hugely important to understand, as it also influences how we understand the other responses. If someone says the reason a child is not in school is because they "work outside the home" that may just be begging the question as the question is "why do they work outside the home and not go to school?" While it is obvious that children not in school work more, it is not so obvious how much is cause (children drop-out of school in order to work) and how much is effect (once children have dropped out (for other reasons) they work more).

Here is my conjecture. Going to school reveals two things. First, it reveals your adeptness for formal schooling (not some catch all like "intelligence" but just how good at school you are). Second, it reveals the quality of your school. By the time most children reach, say 14 or 15 years old, many of them "lack interest" in schooling because either (a) they have realized they are not adept at schooling (and hence do not like it) or (b) they realize the school they are in is miserable and/or no learning is going on or both (a) and (b). Rather than the model of parents pulling children out of school to work (in the market or at home) I would suspect the much more common phenomena is children *pleading* with their parents to not have to go to school.

Just as an example, Tyack (1974) tells that Helen Todd, a factory inspector in Chicago interviewed 500 children working in factories (often in dangerous and unpleasant conditions) and asked the question: "If your father had a good job and you didn't have to work, which would you rather do—go to school or work in a factory?" Of those 500 fully 412 said they would choose factory work. She recounts asking one fourteen year old girl in one particularly unpleasant factory (lacquering canes, involving heat and turpentine) why they did not go to school and got the response "School is the fiercest thing you can come up against. Factories ain't no cinch, but schools is worst."

I think this is a much more plausible view of much of the drop-out phenomena than is "credit constraints." First, strictly speaking, "credit constraints" is not a very good description of the problem. Let us take the author's numbers seriously that the return to schooling is, say, 8-10 percent. Let us suppose that families in developing countries could borrow at the prime interest rate. The real interest rate in many countries in the world is around 8 to 10 percent. So given the opportunity to borrow at prime to finance schooling many households would rationally decline the offer. Second, imagine one relaxes the pure credit constraint—would that money flow into education? Returns on investments from micro-credit programs (which typically have lending rates between 12 and 20 percent per annum) are profitably at those lending rates. One would need more research but the range of investments for which households usually borrow have much higher returns (and quicker) than 8-10 percent.

Third, one needs to be clear about the difference between a "credit constraint" and a "budget constraint." The fact that school enrollments are much lower to poorer people does not prove there is a "credit constraint"—consumption of all types of goods is lower for poor people than rich people because of a budget constraint. I would think evidence about the nature of a credit constraint could be inferred from a large temporary windfall. While there is some evidence of credit constraints from the South Africa evidence they cite, there is a much more widespread (non-experimental) evidence on the use of remittances which my reading of is that the marginal propensity to spend on education is about what we would expect if education decisions were budget constraint decisions not credit constraint decisions.

Kids who drop-out can be classified into two types: (a) those who wish to drop out (and parents do not object) because their individually assessed returns to more schooling are low because either (a.1) their personal adeptness is low and/or (a.2) the quality of available schools is low (e.g. a flat learning profile, increment to achievement per year is low) and (b) those for whom anticipated returns to schooling are high (adeptness and school quality not the key issues) but (b.1) they have very high marginal valuation of non-schooling goods because they are from poor households and/or (b.2) they are credit constrained. Alternatively put, people drop out of school because marginal costs are higher than marginal benefits which could be because benefits are low (a) or costs are high (b).

The problem with assessing the implications of "demand side transfers"—particularly conditional cash transfers is that you are inducing people to return to school who have (or would have) chosen to drop out of school. If one does not distinguish between (a) and (b) reasons it is difficult to believe the returns to this are high, and certainly are not as high as alternative policies that make some attempt to differentiate—suppose by demanding high performance standards.

I realize there are some studies, cited by the authors, in which the returns to compulsory schooling appear to be as high as from chosen schooling—but this is actually a puzzle rather than something we would have expected from an underlying choice based theory—and I suspect these type of results are quite context specific to the schooling system (high average and uniform quality) and labor market (credential effects).

The results from Progresa—of impacts on attendance but not on measured learning outcomes—is consistent with the view that marginal students were forced back into marginal schools. If we really believe the returns to school is returns to skills this just cannot be the case that this has the same returns as a child choosing to attend. The argument in the paper that "drop-outs are disproportionately going to be in households facing liquidity constraints" is perhaps true, but drop-outs are also likely to disproportionately weaker students and disproportionately facing poor school options. Even if we add up "poverty" and "working outside the home" this is less than half of "lack interest" in rural areas from their own "all world" column in table 3.

III) Analysis of Policy Options

A final quick point is that policy options that implicitly consist of a number of policies should be analyzed independently. This is particularly important for two of the demand side issues considered.

First, with "conditional cash transfer" programs it is vital to consider when the intervention is making an existing cash transfer conditional or to large a new transfer that is also conditional. In the Progresa case for instance the cash transfer (in various forms) existed and hence the incremental step against which the incremental enrollment benefits are to be gauged was the making the transfer conditional—so the bulk of the resources were sunk costs in this calculation. If however one is launching a new program then the overall program needs to be justified on its entire range of benefits, of which increased schooling is only one part.

This consideration cuts both ways. For instance, the authors point out that giving a transfer at an age at which all children are in school has low marginal returns—but it may well be the main objective of the *cash transfer* part of the scheme is to transfer income to poor households and hence excluding households with children in those ages from the design would defeat that purpose. On the other hand, it is going to be very hard to justify CCT on their education benefits alone unless they are very sharply targeted to the ages and situations in which the enrollments are very low.

Second, one question with "demand side" interventions is whether the costs of an additional school place (say, the average cost per child) is added to the cost of the program or not. If it is not, on the argument that class sizes are such that increases in class size have no effect so that the marginal cost of an additional student is zero and therefore only the incremental cost of the inducement need be included, this is essentially combining two policies. One is "increase class size" and one is "induce additional students to attend." If the assumptions that allow the exclusion of marginal costs from the demand side calculations is correct (there is not net loss from increasing class size) then the "increase class size" reform has enormous economic returns—that, to scale, dwarf anything about the demand side. Perhaps one could argue that the bundle of the policies is cost effective but we should at least be clear where the action is—if we include the full average cost of an additional school place it must be that these returns are considerably lower (as they have to be lower than for the marginal student who attends with no incentives).

Third, and returning somewhat to the two points above (about market failures and behavioral theories) combined with a sense of the evaluation of policy alternatives, the authors have an embarrassment of riches. According to them iron supplements to secondary school students has a BCR of 32. What is the market failure? One can just buy iron supplements on the market right? So if parents knew of these *private* benefits from a *private* action then they should willing adopt, right? So consider policy of "paying for iron supplements" why not a policy of "publicizing the benefits of iron supplements" from which one should get much higher uptake with less cost.

Similarly with something like *Balsakhi* tutoring. There is a huge flourishing market for private schools in India and for tutoring. If you really believe the BCR is 528—you can get 5,152 in benefits for only 9 dollars—which not the policy of simply making parents aware of the potentially massive benefits of a particular type of tutoring? The behavioral model is that parents will turn down a low total cost hugely *private* benefit intervention? If not, then spreading the information (the creation of which is a public good) should be enormously more cost effective as a policy than scaling up the program (or alternatively the scaling up of this program should be an enormous cash cow the private sector would willingly scale up).

Conclusion

This is an excellent paper, laying out all the issues. But, if the game of normative public economics is to be played pitting one sector against another then all should have to play hard by the same rigorous rules. Those rules are that the counter-factual for evaluating a public sector intervention has to be against the market equilibrium with a plausible behavioral model. On this score education scores super-high on the *private benefits* which makes public policy advocacy that much harder because one has to ask (a) what is the *magnitude* of the market failures that justify the intervention and (b) if there really are huge private benefits to low cost interventions—why are these not (or could not be) scaled up already?