Credit Where Credit Is Due: An Approach to Education Returns Based on Shapley Values

Bilal Barakat[a,b] *(corresponding author), bilal.barakat@oeaw.ac.at*

Jesus Crespo Cuaresma[a,c,d]*, jcrespo@wu.ac.at*

[a] *Wittgenstein Centre for Demography and Global Human Capital (IIASA, VID/ÖAW, WU)*

[b] *Vienna Institute of Demography (VID), Austrian Academy of Sciences, Welthandelsplatz 2, 1020 Vienna, Austria*

[c] *Vienna University of Economics and Business (WU), Welthandelsplatz 1, 1020 Vienna, Austria*

[d] *Austrian Institute of Economic Research (WIFO), Arsenal Objekt 20, 1030 Vienna, Austria*

Abstract

We propose the use of methods based on the Shapley value to assess the fact that private returns to lower levels of educational attainment should be credited with part of the returns from higher attainment levels, since achieving primary education is a necessary condition to enter secondary and tertiary educational levels. We apply the proposed adjustment to a global dataset of private returns to different educational attainment levels and find that the corrected returns to education imply a large shift of returns from tertiary to primary schooling in countries at all income levels.

Keywords: Returns to education, Shapley value.

JEL codes: I26, I25, C71.

1 Introduction

Numerous studies have aimed at assessing the relationship between education and earnings and obtaining estimates of the private returns to different attainment levels. Typically, such estimates derive from earnings regressions and are interpreted as the expected premium for an extra year of schooling in a particular attainment level, keeping other personal characteristics constant.

Assigning returns to years of schooling corresponding to each one of the different educational attainments is not a trivial exercise. Since achieving primary education is a prerequisite for enrolling in secondary education, the returns to schooling corresponding to primary should to a certain extent incorporate the contribution of such an educational achievement to pursuing higher attainment levels. This issue has been noted and discussed in the literature often. Colclough, Kingdon, and Patrinos (2009), for instance, when discussing the dynamics of private returns to primary education, note that '*primary education is a necessary input into further levels of education which may have higher economic returns. If the benefit that primary education confers by permitting access to more lucrative levels of education is taken into account, its true return will increase.*' (Colclough, Kingdon, and Patrinos 2009, 3)

Past theoretical attempts at taking this effect into account have focused exclusively on the 'option value' of completion at a lower level (Weisbrod (1962); Appleton, Hoddinott, and Knight (1996)), in other words: on the forward-looking prospects from an individual perspective based on expected earnings from further educational progression. Arguably, this answers a question different from the one of greatest interest with respect to public policy. From the latter provider perspective, it is more pertinent to assess the retrospective re-attribution to lower levels of returns that have actually been realised after all the choices about educational progression have been made. The difference between these perspectives is highlighted by the argument put forward by Weisbrod (1962) that applying his adjustment simultaneously to returns at all attainment levels does not amount to double-counting. However, this reasoning clearly fails from a provider perspective, since in the distribution of realised returns, each dollar can only be attributed once.

In any case, the prospective approach has seen limited application in practice, perhaps partly because of the recognised challenges of identifying the right transition probabilities between attainment levels. The prospective individual perspective strongly depends on what is assumed with respect to heterogeneity. To the extent that educational attainment is at least partly determined by self-selection based on ability and other factors that also affect earnings, it is not clear that the returns achieved by tertiary graduates really do represent the returns that secondary graduates would likewise achieve if only they progressed further, or that all secondary graduates even have the option of progressing. Indeed, it is precisely for this reason that

Appleton, Hoddinott, and Knight (1996) warn their estimates '*should be treated with caution and [be] regarded as illustrative*' (Appleton, Hoddinott, and Knight 1996, 217). By contrast, from the retrospective provider perspective investigated in this piece, this issue does not arise because no such counterfactual earnings are involved: the returns that tertiary graduates do in fact achieve are partly re-attributed to lower education levels that these tertiary graduates did in fact need to pass. Accordingly, the population shares at different attainment levels are just that, shares, and need not be interpreted as probabilities. To the best of our knowledge, no adjustment to conventional educational returns has been proposed to date that adopts such a retrospective (re-)attribution approach from a provider perspective.

Here, we fill this gap by using an analogy between the problem of (re-)attributing returns between education levels on the one hand and cost allocation based on the Shapley value on the other. Such a reinterpretation of the problem yields returns at *all* levels of schooling that account for their prerequisite relationships. Using the global dataset on private education returns recently developed by Montenegro and Patrinos (2014), we illustrate the calculation of adjusted returns for all countries with available data. To fully implement this reasoning requires knowledge of the population shares reaching each attainment level. In light of the fact that such data are not always available, we also derive an approximation that can be applied to raw return estimates without additional data on population shares, and assess the quality of the approximation. We then discuss the implications of extending the approach to social returns. Our results indicate that adjustments based on the Shapley value tend to equalize returns across attainment levels and result in a shift of returns towards primary schooling in low income countries. The results appear robust when separately comparing returns to schooling for males and females.

The paper is structured as follows. Section 2 provides the conceptual setting to the assignment of returns across educational attainment levels based on the Shapley value. Section 3 presents empirical estimates of adjusted returns. Section 4 discusses the implications and concludes.

2 The Shapley value and returns to education

2.1 Attribution of education returns as a cooperative game

As a motivating thought experiment, consider a syndicate of three teachers, each one able to teach and certify only one level in a hierarchical sequence of qualifications. A student is willing to pay the syndicate a total fee x to complete levels 1 through 3. How should the teachers split the fee between them? As stated, this is an example of a cooperative game.

In cooperative game theory, the Shapley value provides a conceptual basis to such problems allocating collective pay-offs in coalitions. Its logic of 'fair shares' derives from the expected marginal contribution of players as they enter a coalition. Using the Shapley value as an instrument to divide the pay-off is the canonical choice because it presents particularly desirable properties. In particular, it is the only division rule that satisfies additivity, efficiency (that is, the full pay-off of the coalition is distributed), symmetry (equal pay-off for players that are substitutes) and the so-called null-player axiom (players that do not contribute to the pay-off of any coalition are assigned a zero pay-off).

Of course, attainment levels are not perfectly comparable to players in a coalition game, and we are interested in general returns, not fees actually paid to the public provider. Nevertheless, the institutional relationship between different levels of publicly-provided schooling more closely resembles a coalition cooperating to generate some value than competing agents in a market environment. This makes the Shapley value suitable for the computation of returns to education at different levels in a way that take the hierarchical relationship between levels into account.

2.2 Definition and calculation of the adjustment

Assuming a strictly linear dependent ordering of successive schooling levels $i \in \{0, 1, \ldots, N\}$, the Shapley approach yields particularly simple results. Denote by r_i the marginal return to attainment level i over level $i-1$. When derived from coefficients in a log-earnings regressions, these returns are in units of the baseline wage for an individual with no education.

In general, the Shapley value corresponds to the sum

$$\frac{1}{|N|!} \sum_{R} \left[v(P_i^R \cup i) - v(P_i^R) \right],$$

over all possible orderings R of the N players, where P_i^R is the set of players preceding player i in R, and $v(S)$ is the worth of coalition S. In anticipation of our empirical application, from here on we assume $N = 3$, in analogy to the three educational attainment levels we consider. For a given ordering R, $v(P_i^R)$ is the value of coalition before player i joins, and accordingly $v(P_i^R \cup i) - v(P_i^R)$ is the marginal increase in the value of the coalition when player i does join. The normalizing factor $\frac{1}{|N|!}$ ensures that all orderings receive the same weight and these sum to unity. The Shapley value, through the above definition, captures the intuition that the bargaining power of each player in claiming their share of the collective pay-off can be determined by the extent to which their marginal contribution to increasing the collective pay-off depends on the prior

cooperation of other players.

Note that the ordering does not refer to the order in which players contribute to the task, but to the order in which they join the coalition, i.e., the order in which they commit to cooperation. In the context of returns to different edcational attainment levels, this means that we still need to consider all possible orderings of the three attainment levels, including ones that are 'out of sequence' in relation to the educational hierarchy. Indeed, the fact that the primary education 'player' P always increases the value of the coalition regardless of the order of entry into the coalition (since it allows for the returns of higher attainment levels to be realised) is precisely what reflects P's superior bargaining power. The resulting marginal payoffs of the different coalitions are shown in Table 1.

To walk through the interpretation, first consider the first row for the order PST. This order matches the educational sequence, and so as each 'player' joins the coalition, the value of the coalition increases by the corresponding marginal wage premium. This is trivially true as we move from the empty coalition to primary 'player' P alone. When secondary education 'player' S joins next, the coalition of P and S represents a complete educational sequence up to S, so the secondary wage premium can be realised. Similarly, when tertiary education 'player' T joins the coalition next, the educational sequence up to tertiary is complete and the tertiary wage premium is realised.

By contrast, the situation is more complicated in the other rows. For example, in the second row, when the P joins the empty coalition, its worth increases by the primary wage premium r_P, as before. When the T joins next, this does not increase the coalition's worth at all, since exploiting the returns of tertiary education depends on (yet absent) secondary and there is no complete educational sequence up to tertiary represented in the coalition at that point. The wage only jumps from a primary-graduate wage to a tertiary-graduate wage when S joins the coalition, and the educational sequence from primary through to tertiary is completed. The logic leading to the entries in the remaining rows follows analogously.

Recall that the ordering is that of declaring cooperation, not of educational sequencing (which always remains primary, then secondary, then tertiary), in order to assess implicit bargaining power. Summed across orderings, the tertiary marginal wage premium r_T is shared equally between all 'players' (i.e., across all educational attainment levels), r_S is shared equally between the primary and secondary 'players', and the primary 'player' retains all of r_P.

Order	Pay-off to P	Pay-off to S	Pay-off to T
PST	r_P	r_S	r_T
PTS	r_P	$r_S + r_T$	-
SPT	$r_P + r_S$	-	r_T
STP	$r_P + r_S + r_T$	-	-
TPS	r_P	$r_S + r_T$	-
TSP	$r_P + r_S + r_T$	-	-

Table 1: Marginal payoffs to primary ('P'), secondary ('S'), and tertiary ('T') education 'players' as they join the coalition in every possible sequence. $r_P, r_S,$ and r_T: marginal wage premiums.

Labeling the fair share of the marginal return to attainment level j that is attributable to level i as r_{ij}, the above derivation generalises to the straightforward expression

$$r_{ij} = \begin{cases} {}^{r_j}/_{j} & \text{, if} \quad i \leq j, \\ 0 & \text{otherwise.} \end{cases}$$

Accounting for all returns requires combining the return to primary education for those with primary only with the returns attributed to primary for those with higher attainment. Let p_i be the share of the population that completed level i. The conditional probabilities q_{ji} of having completed at least level j among those having completed at least level i are: $q_{ji} = \sum_{n=j}^{N} p_n / \sum_{n=i}^{N} p_n$ for $j \geq i$. The adjusted return is simply given by the weighted sum $\tilde{r}_i = \sum_j r_{ij} q_{ji}$.

Alternatively, consider $\tilde{r}_i^- = r_{ii} = {}^{r_i}/_{i}$, the return attributed to level i for an individual with attainment i. This excludes the share of the return to level i that is 'attributed away' to lower levels, but does not include the share of returns to higher levels that are 'attributed towards' i. As a result, $\tilde{r}_i^-$ is strictly less than $\tilde{r}_i$, and we therefore refer to these as 'decremental' adjusted returns, indicated by the minus-sign superscript. The calculation of the $\tilde{r}^-$ offers the practical advantage of not requiring knowledge of the participation rates, i.e., the values of p_i and the implied q_i. However, by reducing the returns to higher levels without increasing the returns at lower levels, some returns remain un-attributed and unaccounted for when using the adjustment provided by $\tilde{r}_i^-$.

2.3 Interpretation

Procedurally and numerically, the above adjustment is agnostic as to the kind of unadjusted returns it is applied to, although the interpretation will naturally differ. Adjusting private returns yields adjusted private returns, adjusting social returns yields adjusted social returns. The adjustment of wage premiums follows the same process as the adjustment of internal rates of returns. However, it is worth noting that this approach amounts to implicitly re-attributing costs along with the benefits. This is defensible in terms of the underlying conceptual motivation as equating 'the coalition pay-off' with the *net* pay-off. Nevertheless, a case could be made for treating costs differently from benefits in the re-attribution, or even performing no re-attribution of costs at all. This would make the adjustment much more challenging to perform, since it could no longer operate on published unadjusted returns directly, but would require getting back to their raw input data in order to explicitly disentangle costs and benefits. In addition, such an approach would result in an even larger adjustment, especially under typical conditions of higher costs at higher levels of attainment.

Both adjusted returns are comparable to conventional marginal returns in the sense that they capture the benefit attributed exclusively to a given attainment level, net of the contribution of lower levels. They are, however, *pseudo*-marginal in that they do not relate directly to an individual's marginal choice of continuing or not continuing education to the next level. For example, a primary graduate's choice to continue with secondary education would 'retrospectively' change their return to primary.

As is evident from the definitions above, the adjustment is performed on the returns to complete education attainment levels, in other words, on the return to completed primary/secondary/tertiary education. These estimates can be translated to a 'per-year' scale, based on the number of years at each level of schooling. This is done in the empirical analysis below, for instance, for consistency with the convention of existing published estimates. However, note that adjusted returns do not decompose with respect to nested education levels. This reflects a general property of the underlying Shapley values. In particular, adjusted returns for individual years of schooling do not aggregate up to the adjusted returns for whole attainment levels shown here. This means that the results would differ if the adjustment had been directly performed on the returns to individual grades within attainment levels.

This is less of a limitation than it might seem. The most comprehensive and widely-used current dataset of returns to education at different levels (Montenegro and Patrinos 2014) is likewise based on estimates of attainment-level returns that are rescaled down to, not aggregated up from, returns to individual years at the given level. Indeed, while it would be possible in principle to estimate the marginal returns to individual grades using sufficiently detailed datasets, this is seldom, if ever, done in practice. Studies that estimate

single coefficients for the returns to particular ears of schooling do not tend to distinguish between different levels of schooling to begin with. This practice reflects the fact that levels are the natural scale at which differences in returns at different stages of educational progression are used to inform or guide public policy. There tends to be little interest in the differences between individual years at a given level of schooling. To the extent that policy tends to focus specifically on the transition between specific grades, this is motivated by removing a 'bottleneck' in the progression of students through the system, not by considerations of differences in returns to specific grades.

3 Empirical application: Revisiting returns to education

3.1 Data

Montenegro and Patrinos (2014) present homogenized estimates of private returns to schooling, by sex, for 139 countries, based mostly on data since 2000, and for all world regions (although high income countries and Latin America & Caribbean are overrepresented in terms of the number of data points). These are derived from micro-level Mincerian earnings functions, with dummy indicators for attainment levels capturing completed-level wage premiums. These are re-scaled according to the number of school years at each level, and corrected for the opportunity cost of foregone earnings.

Unfortunately, a comparable current dataset of *social* returns is not available. Certainly not of comprehensive estimates including non-monetary externalities, but not even of 'narrow' social returns accounting only for total (private and social) costs. Indeed, Psacharopoulos (1994) continued to serve as the source for *social* returns for Psacharopoulos and Patrinos (2004), and is—to the best of our knowledge—still awaiting an update at the time of writing.

The adjusted returns proposed above additionally require information on the proportion of population completing each one of the different attainment level. These are sourced from the World Bank's World Development Indicators (World Bank 2013), specifically the 2010 values from the Lutz, Butz, and KC (2014) time series on attainment of the population aged 25 and above, which covers more countries than the Barro-Lee data (Barro and Lee 2013).[1] In general, attainment data are more relevant in this context than enrolment rates (whether net or gross), since the returns themselves are estimated for the adult, not the school-age population. Accordingly, we used the attainment shares of the population aged 25 and above to construct the corresponding adjustment.

[1] Our conclusions remain entirely unchanged if the Barro-Lee data are used instead. Results using this alternative source are available from the authors on request.

3.2 Adjusted returns to education: A comparison

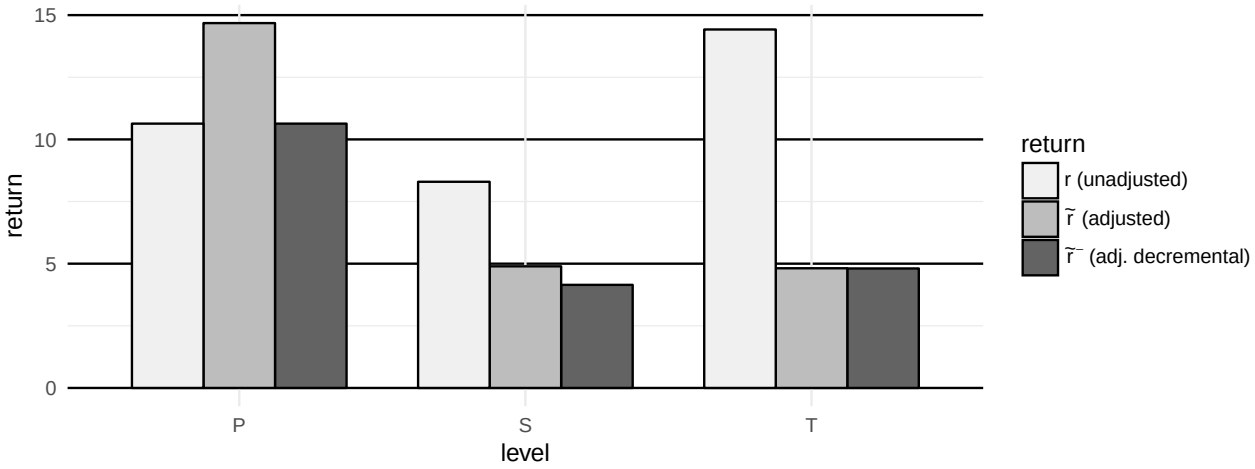

Figure 1: Average returns by adjustment type: world sample (unweighted)

Using the data described above, we compute adjusted returns to education for three attainment levels and both sexes. Unweighted averages of the unadjusted and adjusted returns aggregated over sex and all countries are depicted in Figure 1, annualised to match the scale of the conventional figures.

Figure 1 shows that the position of primary and tertiary clearly flip after the adjustment, with adjusted primary returns far exceeding adjusted secondary and tertiary returns, which are approximately equal. Returns to secondary education are affected both positively and negatively by the re-attribution implied by the adjustment proposed, with a negative net effect reflecting the relatively low participation in, and therefore gains flowing down from, tertiary education at the global level.

The large reduction in the return attributed to tertiary education implied by the proposed adjustments is not balanced by the increasing returns to primary and secondary. For the decremental adjustment, this is unsurprising, given its 'leaky' definition. Moreover, as a general principle, any increase to the (pseudo-)marginal return to primary (whether actual or through re-attribution, as with the adjusted returns) accrues to a larger number of individuals than the marginal return to higher levels. Expressed as averages over beneficiaries, a large decrease in the return to tertiary is therefore indeed balanced by a relatively smaller increase in the return to primary.

Table 2 shows the unadjusted and adjusted returns by gender and country income level. The effect of the adjustment on the relative magnitude of the returns at different levels is fairly consistent across the whole range of overall levels of returns and educational development. Accordingly, the main patterns characterising the unadjusted returns, with the exception of the relationship among education levels, are largely preserved:

returns are higher for females in general and particularly high at lower levels of economic development. The adjustments lead to significantly less variation of returns to secondary and tertiary schooling across countries by income group as compared to their unadjusted counterpart.

The purely decremental adjustment, which can be calculated even without knowledge of participation rates, represents a surprisingly good approximation to the full adjustment in terms of indicating the overall shift in balance. Taking the 'total' figures for the low income group as an example, the unadjusted tertiary return of 19.4 is approximately 1.2 times higher than the unadjusted primary return of 16.4, while this ratio already drops to 0.4 for the decremental adjustment and only exhibits a moderate further change to 0.3 through the full adjustment. The decremental adjustment represents a similarly good indication of the overall shift in other cases.

Table 2: Average (unweighted) returns by adjustment type, gender, and income group.

Gender	Level	Return	High	Upper Middle	Lower Middle	Low
Female	Prim.	Unadjusted	10.1	10.9	14.3	17.4
		Adjusted	13.3	14.5	17	21.7
		Adj. decr.	10.1	10.9	14.3	17.4
	Sec.	Unadjusted	6.5	6.9	6.8	15
		Adjusted	4.7	4.5	4.5	7.9
		Adj. decr.	3.2	3.5	3.4	7.5
	Tert.	Unadjusted	10.3	18.5	17.1	18.5
		Adjusted	3.4	6.2	5.7	6
		Adj. decr.	3.4	6.2	5.7	6.2
Male	Prim.	Unadjusted	6.8	8.4	10.2	13.7
		Adjusted	10.3	11.6	13	18.3
		Adj. decr.	6.8	8.4	10.2	13.7
	Sec.	Unadjusted	5.9	6.3	6	13.7
		Adjusted	3.8	4	3.9	7.5
		Adj. decr.	3	3.1	3	6.9
	Tert.	Unadjusted	10.8	16.7	13.3	16.7
		Adjusted	3.6	5.6	4.4	5.6
		Adj. decr.	3.6	5.6	4.4	5.6

Gender	Level	Return	High	Upper Middle	Lower Middle	Low
Total	Prim.	Unadjusted	8	9.4	10.3	16.4
		Adjusted	13.8	12.8	13.6	20.1
		Adj. decr.	8	9.4	10.3	16.4
	Sec.	Unadjusted	9.7	6.1	7.2	11.5
		Adjusted	5.4	3.9	4.4	6.4
		Adj. decr.	4.8	3	3.6	5.8
	Tert.	Unadjusted	9.9	15.7	13.6	19.4
		Adjusted	3.3	5.2	4.5	6.5
		Adj. decr.	3.3	5.2	4.5	6.5

4 Discussion and policy implications

We propose a method based on Shapley values to adjust returns to education across attainment levels, thus accounting for the hierarchical nature of lower educational attainment levels as necessary conditions to achieve higher ones. This adjustment can be practically applied in a simple manner both to private and social returns.

In some ways, an application to social returns would be the most natural. As noted, the underlying reasoning is based on a provider perspective. One consequence is that the Shapley-adjusted returns are pseudo-marginal in that they do not directly relate to the individual decision problem of educational progression. From a public policy perspective, the difference in interpretation between 'proper' marginal and pseudo-marginal returns is not clear-cut, however. Neither refer to the use of 'marginal dollars of public funds', after all. Even conventional marginal returns are therefore correctly interpreted as merely motivating general development priorities.

Despite this natural affinity between the Shapley adjustment and a public policy perspective, adjusting private returns following the rationale proposed in this piece is a useful exercise. To begin with, to gain insight into the effect of the adjustment across a wide range of profiles, there is no alternative, given the absence of a current set of global estimates of social returns. For the same reason, for better or for worse, public policy is in fact frequently based on estimates of private returns, and this mismatch will be ameliorated by at least adjusting these private returns from a public provider perspective. Moreover, there is a policy dimension to private returns in their own right, not least in terms of the potential of educational expansion as a means of poverty reduction. Finally, given the treatment of costs implicit in the adjustment as implemented

here (see the discussion in Section 2.3), the relative magnitude of the shift in returns towards lower levels is independent of whether the unadjusted returns are private or social. Accordingly, the comparison of unadjusted and adjusted private returns, rather than the absolute values of the latter, can inform public policy actions even in the absence of social returns that they should ideally be based on.

The level estimates obtained for secondary education, which are significantly lower than their unadjusted counterpart, need to be interpreted with caution. The sizeable decrease in private returns to secondary education which results of the adjustment should not be interpreted as evidence that this educational attainment level appears less important in terms of its contribution to aggregate income. In fact, recent empirical contributions highlight the central role played by secondary education as an engine of economic development through its effect on total factor productivity. At the macroeconomic level, investment in secondary education has been shown to be particularly benefitial for developing economies in terms of GDP per capita growth (Lutz, Crespo Cuaresma, and Sanderson 2008) and as a determinant of vulnerability to climate change risks (Striessnig, Lutz, and Patt 2013; Lutz, Muttarak, and Striessnig 2014). To the extent that such spillovers are not explicitly accounted for in estimates of private returns with or without adjustment, direct comparisons of returns of schooling are not able to offer a perfect instrument to inform public policy on expenditure distribution and should thus be complemented with alternative empirical evidence. Moreover, even in terms of the present estimates, both unadjusted and adjusted returns show the smallest amount of *variation* at the secondary level, suggesting that secondary education is in fact the most reliably rewarding investment, and that an analysis taking risk-aversion into account would increase the estimated value of secondary relative to the other levels.

With these observations in mind, we advance the following interpretation of the results. The adjusted returns should not replace unadjusted returns as a basis for policy decisions. Instead, the *difference* between unadjusted and adjusted returns identifies the extent to which returns are generated *jointly* at interacting levels of schooling. From this perspective, our finding is that these jointly-generated returns that cannot unambiguously be assigned to one level or another in isolation (i.e. the off-diagonal returns $r_{i,j}$ with $i \neq j$ that are realised at a different level than they are attributable to) are sufficiently large to reverse the ranking of levels depending on where they are included. This insight suggests that high returns at one level simply do not serve as a sound economic justification in educational or general development policy for privileging investment at that level at the expense of other levels. Instead, increasing these jointly-generated, systemic returns calls for balanced development of the entire educational ladder.

Acknowledgements

The authors would like to thank Felix Mendoza for research assistance, as well as two anonymous referees for helpful comments on an earlier draft of this paper.

References

Appleton, Simon Mark, John Hoddinott, and John B Knight. 1996. "Primary Education as an Input into Post-Primary Education: A Neglected Benefit." *Oxford Bulletin of Economics and Statistics* 58 (1). Department of Economics, University of Oxford: 211–19.

Barro, Robert J, and Jong Wha Lee. 2013. "A New Data Set of Educational Attainment in the World, 1950–2010." *Journal of Development Economics* 104: 184–98.

Colclough, Christopher, Geeta Kingdon, and Harry Anthony Patrinos. 2009. "The Pattern of Returns to Education and Its Implications." *Policy Brief 4, Research Consortium on Educational Outcomes and Poverty.*

Lutz, Wolfgang, William P Butz, and Samir KC. 2014. *World Population and Human Capital in the Twenty-First Century.* Oxford University Press, Oxford (UK).

Lutz, Wolfgang, Jesus Crespo Cuaresma, and Warren Sanderson. 2008. "The Demography of Educational Attainment and Economic Growth." *Science* 319 (5866). American Association for the Advancement of Science: 1047–8.

Lutz, Wolfgang, Raya Muttarak, and Erich Striessnig. 2014. "Universal Education Is Key to Enhanced Climate Adaptation." *Science* 346 (6213). American Association for the Advancement of Science: 1061–2.

Montenegro, Claudio E, and Harry A Patrinos. 2014. "Comparable Estimates of Returns to Schooling Around the World." *World Bank Policy Research Working Paper*, 7020.

Psacharopoulos, George. 1994. "Returns to Investment in Education: A Global Update." *World Development* 22 (9). Elsevier: 1325–43.

Psacharopoulos, George, and Harry Anthony Patrinos. 2004. "Returns to Investment in Education: A Further Update." *Education Economics* 12 (2). Taylor & Francis: 111–34.

Striessnig, Erich, Wolfgang Lutz, and Anthony Patt. 2013. "Effects of Educational Attainment on Climate

Risk Vulnerability." *Ecology & Society* 18 (1). The Resilience Alliance: 1–16.

Weisbrod, Burton A. 1962. "Education and Investment in Human Capital." *Journal of Political Economy* 70 (5, Part 2): 106–23.

World Bank. 2013. *World Development Indicators 2013*. The World Bank.