International Institute for
Applied Systems Analysis
Schlossplatz 1
A-2361 Laxenburg, Austria

Tel: +43 2236-807-536
Fax: +43 2236-807-503
E-mail: riley@iiasa.ac.at
Web: www.iiasa.ac.at

Finance and Sponsored Research

CESDEG: Education for all Global Monitoring Report (EFA-GMR)

Final Report Projections

Submitted to

United Nations Educational, Scientific and Cultural Organization (UNESCO)

IIASA Contract No. 15-148

August 2016

Modelling SDG scenarios for Educational Attainment and Development

Bilal Barakat
Wittgenstein Centre for Demography and Global Human Capital (IIASA, VID/ÖAW, WU)

1 The educational attainment projection model

1.1 Summary

The scenarios of educational expansion underlying the population projections presented here result from a further refinement of the education model presented in Lutz et al. (2014). In summary, we project the share of the population ever reaching or exceeding a given attainment level. This is done seperately by country, and gender, but with 'shrinkage' within a Bayesian framework (with weakly informative priors). The mean expansion trajectories are modelled as random walks with drift (and potential mean reversion) and independent noise at a probit-transformed scale. The trend parameters are estimated based on reconstructed attainment histories, and extrapolated, subject to additional and some exogenously imposed convergence within regions and between females and males. Under the target scenarios, SDG targets are treated as 'future data' (in other words, target trajectories are modeled looking back from 2030 under the assumption that the target will have been met), with a potential trend break in 2015.

Limitations shared with all existing global projections of educational development include the fact that in the absence of a detailed theoretical basis, they are forced to rely heavily on statistical extrapolations. For example, there is little consensus on whether "higher education is the new secondary education" (as claimed by Andreas Schleicher of OECD), or is fundamentally different from lower levels of schooling (e.g. in terms of institutional framework, its role in the life cycle, economic returns. In addition, global projections can necessarily not account in a satisfactory manner for idiosyncratic policy changes or shocks. In addition, the specific modelling choices outlined above imply a number of trade-offs. Using highest school attainment as the underlying measure solves many problems associated with historic enrolment data by allowing the consistent reconstruction of time series of attainment from relatively recent cross-sectional data, but comes with challenges of its own. While nevertheless preferable overall, the principal disadvantage of attainment measures deserves mention, namely the relatively long time lag with which outcomes can be observed. Late attainment is common in many developing countries, so that attainment cannot safely be assumed to be 'final' until several years above the typical graduation age.

The model operates on 5-year age groups and in 5-year time steps. While the starting (2015) and target (2030) years for the SDGs conveniently line up with this grid, typical durations and graduation ages for different attainment levels unfortunately do not. The target is interpreted such that the cohort aged 15-19 in 2030 will ultimately (not necessarily already at that age, which would be too early for the 15-year-olds with respect to upper secondary) universally attain secondary education. In order to ensure that most late attainment is captured, completed primary attainment is observed at age 15-19, completed lower secondary at age 20-24, and completed upper secondary and post-secondary by 25-29. The latter is likely to underestimate the amount of post-secondary attainment somewhat, but an even higher reference age would come at the cost of an even greater time lag and less current observational data.

The basic model specifies that the inverse probit of the share attaining a given education level or higher among the entire cohort follows a random walk with country-specific drift. In principle, the specification also allows for mean-reversion by partially backtracking an (estimated) proportion of the random shock of the previous period, but in practice no meaningful mean-reversion of this kind was picked up from the data. This is not necessarily surprising, given that mean-reversion on a year-on-year basis will largely be obscured by the 5-yearly data.

Additional complexity is layered over this basic model. Gender convergence is specified such that at each time step, the predicted values for both genders are shifted towards their joint average. An additional level of independent errors of small magnitude that do not persist in the random walk and do not enter the gender convergence is allowed in fitting the observed data, in order to account for exogenous errors at the level of data, rather than in the underlying educational process.

The fitted empirical model is adjusted during projection in the following ways. (Level and gender specific) country trends linearly converge over six time steps to the regional trend. The strength of gender convergence increased in two steps to reach twice the past empirical value. The logical inequality relations between the participation shares (e.g. that the share attaining secondary or higher must be less than the share attaining primary or higher) is enforced by capping participation at the higher attainment at the level of the prerequisite attainment. Projected attainment at the post-secondary level is rescaled to remain below 90 percent, based on substantive reasoning.

For the target scenario, the above forward projection approach is modified. While it would be possible to deterministically calculate the necessary additional drift to reach a given point target level by 2030, doing so would be a lost opportunity to gain additional insight. Instead, SDG targets are treated as "future observations". Specifically, they enter the likelihood by specifying that the drift resulting in the overall upward trend is allowed to increase by whatever amount necessary (with an effectively flat prior) to reach the target, starting in 2015.

Note that this specification of the target scenarios means the target of 97 percent is typically exceeded, not just barely met, in contrast to a typical 'target-achieving path' interpolated deterministically. This behaviour is desired and deliberate. Intuitively, assuming a country did meet the targets, these trajectories represent typical paths of having got there. Retrospectively, the set of countries that meet the targets will have exceeded them on average, given their lack of perfectly exact control over the outcome. An analogy will clarify this: if we invite a group of runners to attempt to run 100 m in 11 s, then the successful group will clearly have taken less than 11 s on average. Since in addition, the target scenarios have the same probabilistic nature as the trend scenario, they allow for arbitrary conditioning. Examples of such conditional perspectives include questions related to the probability of different countries meeting fixed targets by a certain time, to complement the more conventional question of the probability of exceeding certain participation levels in a fixed year. While this is fully analysed elsewhere, for present purposes we focus on the 'minimal' target path traced out by the cross-sectional 0.01 quantile of the target paths that only just reaches the SDG target. In addition to sharing their probabilistic nature, just like the trend scenario, the target scenarios incorporate the nonlinearity of educational expansion as it really occurs. In particular, this includes the likely deceleration of expansion as universal participation is approached, as well as the fact that countries that meet the targets will necessarily have "overshot", on average. This allows us to quantify the risk of failure associated with attempting to monitor whether countries are 'on track' according to simple linear plans.

Fig. 1 displays an illustrative example of a projection for a single country, gender, and education level. One hundred of the simulated trajectories are shown, as well as a particularly high, particularly low, and middling trajectory. In addition, the dashed line connects the cross-sectional medians and represents the 'median trajectory' for projection purposes. Note that it is smoother (by construction) than any individual simulated trajectory. The small amount of uncertainty around the observed points regarding the 'true' past value is evident in the way the projections do not fully coincide at the latest observed point in time.

The following section discussions some modelling decisions in great technical detail and can be skipped by readers more interested in the results.

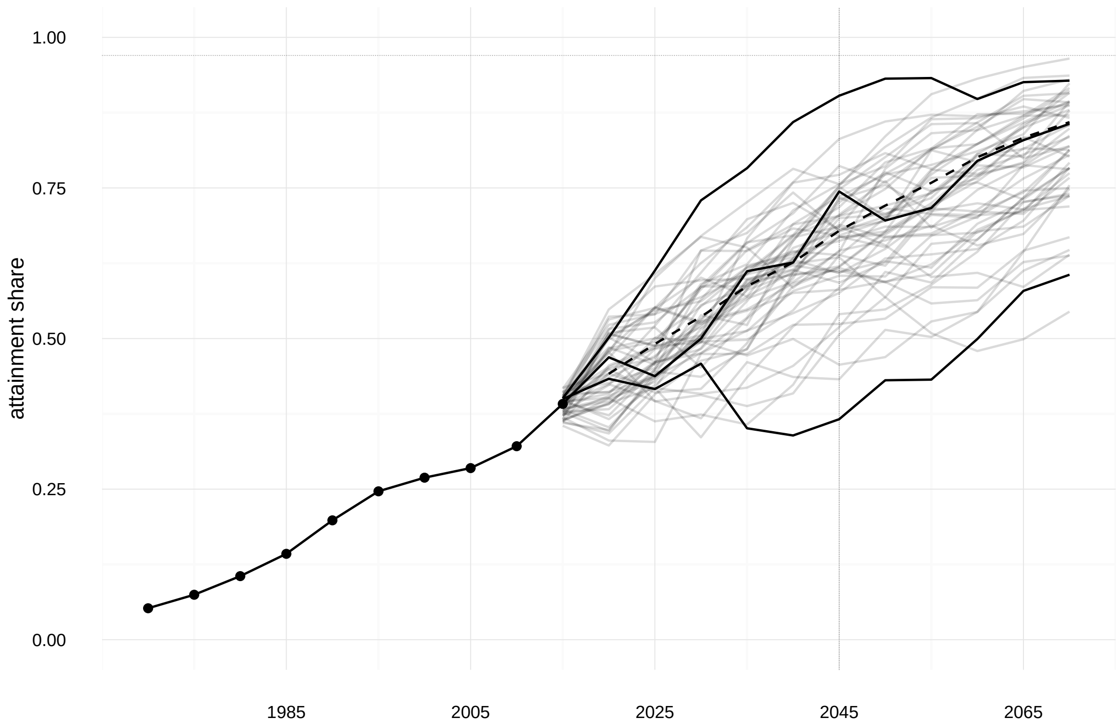

Figure 1: Illustrative ensemble of simulated trajectories. One very high, one very low, and one middling trajectory are highlighted. Dashed: cross-sectional median trajectory. Dots: observed data points.

1.2 Design considerations

1.2.1 Time and age alignment

The model operates on 5-year age groups and in 5-year time steps. While the starting (2015) and target (2030) years for the SDGs conveniently line up with this grid, typical durations and graduation ages for different attainment levels unfortunately do not.

The SDG target is interpreted such that the cohort aged 15-19 in 2030[1] will ultimately (not necessarily already at that age, which would be too early for the 15-year-olds with respect to upper secondary) universally attain secondary education. In order to ensure that most late attainment is captured, completed primary attainment is observed at age 15-19, completed lower secondary at age 20-24, and completed upper secondary and post-secondary by 25-29. The latter is likely to underestimate the amount of post-secondary attainment somewhat, but an even higher reference age would come at the cost of an even greater time lag and less current observational data. Conversely, while these cut-offs may appear unnecessarily conservative (i.e. high), the data clearly show that lowering them by five years would miss significant amounts of attainment. Even without delayed entry or repetition, even in terms of nominal graduation age, lower secondary is not completed until age 15 or 16 in some countries, or upper secondary until age 20. In addition, in the German-speaking countries for instance, higher vocational qualifications that are normally acquired in one's twenties are

[1]For ease of interpretation, no attempt is made at a fractional alignment (that would specify that in terms of the age group 15-19, universal upper secondary attainment would only be reached in 2031, say, because the 19-year-old cohort in 2030 graduated upper secondary *before* 2030).

formally equivalent to upper secondary schooling. As can be seen from the results, even the conservative thresholds given above may be too low in some countries. Unfortunately, in the absence of a large-scale effort to define country-specific thresholds based on a case-by-case analysis of the education system and participation behaviour, fixed thresholds for all countries must necessarily achieve a compromise between timeliness and completeness with respect to late attainment.

One might be tempted to lower the age thresholds, with the intention of modelling *timely* attainment specifically. In other words: upper secondary attainment, for example, that hasn't been achieved by age 20-24 "doesn't count". There are, however, at least four problems with this: a) as in the German example above, it is possible even for "timely" upper secondary students not to have completed at age 20, or for timely lower secondary students not to have completed by age 15, b) for the older cohorts, we have only observed their ultimate attainment, not their timely attainment, so the time series would be inconsistent, c) all the evidence on the effects of educational attainment is for ultimate attainment, not timely attainment, and d) since part of the growth in timely attainment may come from reducing late entry and delayed progression, it makes no sense to extrapolate that trend, because those are finite reservoirs, and ultimately the growth in timely attainment is anyhow bounded by the growth in ultimate attainment (once everyone is on time anyway).

1.2.2 Persistence

Intuitively, it seems highly plausible to expect some degree of persistence of deviations from the long-term trend in educational expansion. In other words, we do not necessarily expect the effects of a negative shock in one period to leave attainment in the following period unaffected. Indeed, the absence of such persistence in the previously used regression specification with independent errors has been one of the most frequently voiced criticism we encountered.

A simple time trend with independent errors (as used by all existing models of educational expansion) exhibits no such persistence, while in the alternative specification of a random walk, all shocks permanently shift the entire future trajectory. In reality, we expect a mixed behaviour. On the one hand, some conditions that would lead to above/below average expansion in one period, such as economic crises or conflict, will often affect multiple five-yearly time spans. Moreover, the education sector is known for its strong inertia. On the other hand, 'excessive' or 'insufficient' growth in one period may be partly compensated by deliberate policy and market forces in the next.

At least two well-established general specifications for capturing such partial persistence need to be considered, namely either a secular time trend flanked by autocorrelated errors, or a random walk with partial mean reversion at each step. Empirically, these two specifications can be notoriously difficult to distinguish, and as time series, the seven observations per country make for time series that are too short to conclusively point to one or the other specification. The choice is therefore based on conceptual considerations. In particular, the autocorrelated error specification would imply, in principle, that during a period of stagnation, an ever increasing 'expansion debt' is built up relative to the expected level. This expectation does not match well how we tend to think about educational development. By contrast, under the random walk specification, expansion trajectories can always be considered to 'carry on where they left off'. In other words, it is *time* that is lost during periods of stagnation, which corresponds to the development community's language around 'lost decades', for example. We therefore prefer the random walk specification.

Note that allowing for persistent shocks leads to more conservative estimated variation in country trends, since spells of above/below average growth can more easily be attributed to the error component of the model than if the errors were assumed to be independent. At the same time, the estimated prediction intervals become wider and more uncertain. We consider both of these effects to be an asset, especially since the existing regression specifications result in prediction intervals that are arguably rather too narrow.

1.2.3 Limiting educational decline

One aspect of extrapolating past trends is that slight historical declines may, if projected sufficiently far into the future, result in complete collapse of education, especially if the non-linear expansion pattern is reversed. There are unfortunately many examples of stagnation and set-backs in the history of educational development, especially as a result of conflict (Syria providing a drastic recent examples). However, there is no precedent for a country permanently going on a 'reverse development trajectory'. Projecting such a development is therefore to be avoided.

One way of preventing this is to include some degree of cross-country convergence. This relies on the fact that most countries do in fact exhibit positive trends. Regardless of whether such convergence is actually estimated, or expected in reality, even a slight amount of convergence is sufficient to avoid any country entering a projected trajectory of long-term collapse. This approch was followed in the education projections underlying Lutz et al. (2014), for example. Here, much stronger convergence at the regional level is assumed, reflecting normative expectations of the educational development community.

Moreover, the random walk specification itself also tends to reduce the risk of the country-specific drift being estimated to be negative. This is because a sustained period of stagnation can be accounted for by a sequence of negative shocks that are all on the order of magnitude of the drift. By contrast, under a simple regression model, ever-increasing negative errors would have to be assumed to reconcile sustained stagnation with a 'true' positive drift.

1.2.4 Cross-country convergence

The decision to include cross-country convergence leaves room for a wide variety of specifications. In particular, this includes the question whether to converge to the regional, global, or some other kind of average, and following what function over time.

Note that it is the drift parameters that converge, not the attainments as such. Here, we converge the country-specific drifts simply by reducing the scale parameter that determines their variation around the regional average. In particular, the scale is shrunk to zero linearly over 6 steps (i.e. thirty years). Such a relatively slow convergence avoids abruptly stopping the rapid expansion among the frontrunners.

This approach implies that convergence is to an *unweighted* regional mean. Whether this is appropriate for regional 'heavyweights' such as China and India, but perhaps also Nigeria, for example, is a matter for debate. We have chosen to model education *systems* as the unit of analysis.

The regional groupings are derived from the GEMR regions, with Australia and New Zealand combined with North America and Europe for purposes of convergence.

1.2.5 Gender convergence

Unlike cross-country convergence, that is only applied to the projection, because its justification is partly normative, the degree of gender convergence present historically is empirically estimated as part of the model fitting. Again, based on normative expectations this is then 'boosted' during the projection period. Another reason for doing so is that — as already mentioned — the completed attainment data may not fully reflect the most recent developments during the EFA period 2000-2015 and may therefore underestimate the amount of gender-convergence.

In particular, gender convergence is specified such that at each step, the gender-specific outcome is shrunk towards the average across both males and females, with an estimated shrinkage parameter. This parameter is allowed to differ across education levels and countries, but is constant across time. During the projection period, this parameter is increased proportionately up to twice times its historical value (with multiplication factors increasing linearly over two time steps), capped at 0.5. This increase was calibrated to avoid actual declines in the outcomes of the higher group as it is shrunk towards the average.

Another reason why gender convergence is specified in terms of levels rather than rates is that if the lagging unit is actually expanding more rapidly, strong convergence in rates actually *delays* convergence in levels. In principle, this applies equally to cross-country convergence, however it is a greater concern with respect to gender convergence because: a) the above situation is very common (female education often lags behind, but is actually growing faster), and b) the assumed convergence is stronger.

1.2.6 Post-secondary ceiling

Substantively, we do not expect post-secondary attainment to become fully universal at any point. To account for this, we impose a ceiling of 90 percent to this level. This is somewhat arbitrary, but reflects the fact that in the most advanced countries, post-secondary participation is already approaching 80 percent. A ceiling much below 90 percent would therefore require a very sudden expansion stop, or even the baseless assumption that this current levels already represent an 'overshoot'.

In principle, an attempt could be made to *estimate* the saturation level. However, for post-secondary, the vast majority of observations are well below the inflection point of the s-curve of expansion. Estimating the maximum level on these data would require excessive confidence in the accuracy of the functional specification. A prior could be put on the saturation level, so that, effectively, some runs would converge to a ceiling of 90 percent, others to 95 or 85 percent, for example. However, again this would then be transformed to a posterior that based on data that may not actually be informative. The alternative is to add uncertainty to the ceiling post-hoc, but doing so would risk 'over-engineering' this adjustment.

1.2.7 Trend Break

In reality, the transition onto a new, target-achieving, trajectory would be expected to occur gradually. While in general it would be feasible to 'phase in' a new drift, in the case of the SDGs, with a target horizon of only 15 years, any trajectory actually reaching the target will have to reach full speed sufficiently rapidly so that in terms of 5-year time steps it can be treated as applying immediately.

The start of the trend break is adjusted by attainment level, since the cohort aged 15-19 in 2010, for example, will already eventually benefit from increased post-secondary participation during the period 2015-2030. Conversely, changes starting in 2015 were largely too late to affect the primary attainment of those already aged 15-19 in 2020.

1.2.8 Target specification

Several components contribute to the justification for such uncertainty around the target. All presume agreement that insisting on a point target of *exactly* 100 percent is unreasonable in practice. Point targets below 100 percent are not meaningful, since no country would be judged to have 'failed' the SDG target operationalised as 98 percent participation if it overachieved and reached 99 percent instead. With this in mind, uncertainty around the model target arises from the following three sources.

Firstly, even if 'universal' is operationalised by a target *range* below 100 percent, there may be reasonable disagreement about the exact value the lower threshold should take, i.e. how close to universal is 'close enough'. A threshold of 97 percent has been used by Unesco in the past, for example, but there is clearly no objective reason why another entity should prefer this over 97.5 percent, 98 percent, or 99 percent, say.

Secondly, given a lower threshold for 'true' participation to be sufficiently close to universal, some allowance must be made for measurement error. This works both as an argument for having a target range in the first place, but, importantly, also as an argument to *raise* the lower threshold with respect to the achievement as measured/observed. In other words, even if we agreed that having 97 percent of children in fact complete secondary school means the target has been met, we may only be sufficiently confident this really is the case if the measured share is somewhat higher, say 98 percent, to account for the possibility this is in overestimate.

Thirdly, even if we accepted nominal 97 percent as sufficient, countries do not, of course, have perfect control over the process. Not only does this suggest that countries need to aim to 'overshoot' if their aim is to cross the threshold. It also means that even if all countries aim squarely at the minimal 97 percent regardless, then post-hoc conditioning on having succeeded will still lead to an average 'overachievement' among the successful countries. This is perfectly intuitive in other contexts. All sprinters taking ten seconds or less over a 100 m dash will on average take strictly *less* than ten seconds.

While the latter two points are related, in the sense that measurement error is itself a form of 'lack of perfect control' over the *measuremed* outcome on the part of countries, it is worth a separate observation this argument applies even if we take all observed indicators at face value. Because of these multiple reasons for it, no attempt was made to endogenise uncertainty around the target, by making it a function of the *estimated* magnitude of the shocks, for example. In this sense, the target of 'universal' participation is more complicated to handle than a simple point target of some other fixed percentage. In that latter case, it would be possible to simply treat the target *exactly* as any real observation for purposes of conditioning on its achievement, with the exact same likelihood contribution.

In the present case, for the reasons just discussed, a different specification should be chosen. For the present projections, the aim is a 'fuzzy' target distribution at the original scale that is practically flat over a couple of percentage points from 97 percent to 99 percent, but drops off rapidly in either direction. A discontinuous cut-off below 97 percent is undesirable for computational reasons, because the implied zero gradient in the likelihood would fail to guide the algorithm towards the target region. In any case, 'meeting the target' is not a perfectly sharp concept in the policy domain either, even once it has been operationalised with a numeric threshold.

In order to achieve the above pattern at the original scale, an exponentially modified Gaussian distribution (with mean corresponding to 0.97 at the untransformed scale and sigma = 0.05, lambda = 0.5) is specified around the target at the transformed scale. The reason for excluding values very close to true unity at the scale of participation shares is that these would translate to values at the transformed scale that diverge to infinity, requiring an unbounded speed-up of expansion.

While the above approach is more meaningful in a probabilistic framework than assuming that 'meeting the target' means all countries reach *exactly* 98 percent in 2030, say, it does represent a communication challenge. The target are extremely challenging as it is, and country representatives are likely to complain if they are deemed to have failed to stay on 'the' target track if the representative target path significantly overshoots the minimum threshold, as both the trace of cross-sectional means and medians do under the above approach. In order not to stray too far from established practice, we therefore identify the 0.01 quantile, i.e. the 'minimal' path, as the benchmark for progress towards the target. The full uncertainty range across target trajectories is employed merely as a pedagogical device to raise awareness of the fact that there is unlikely to be widespread success if everyone merely shoots for the minimum.

Unfortunately, solving one communication problem immediately creates another, because at the beginning of the projection period, the lowes percentile of the target-attaining trajectories naturally lies *below* the median of the baseline trend. Displaying the former as 'the target path' and the latter as 'the trend path' therefore creates the absurd impression that in order to reach a highly ambitious target, we 'ought' to start by slowing down. This leads to the - admittedly ad-hoc - solution of using the trend path as a lower cap for the target path, which therefore effectively deviates from the former not immediately in 2015, but with a delay. In the absence of an established statistical approach to summarising probabilistic target scenarios together with their necessary overshoot, it is not clear how else to deal with the very real communication challenge.

A second subtlety created by the desire to estimate target-driven scenarios probabilistically within a Bayesian setting deserves additional attention. Recall that the proposed set-up corresponds to treating the target as a 'future observation', and effectively selecting target-achieving trajectories by conditioning on the target being achieved. One implication is that, even though these trajectories may make use of a trend-break, the *historical trend* may also be estimated differently in the target scenario. Technically, this is, of course, perfectly correct. By conditioning on target-achievement, we are effectively answering the question: supposing the target is reached, how did we get there? And it is indeed both correct and statistically intuitive that among universes where Thailand, say, reaches universal secondary participation by 2030, those will be over-represented where, historically, Thailand actually has a higher 'intrinsic' expansion rate than historical evidence suggests, and it has to date been underperforming relative to its capabilities. However correct it may be, this implication creates a communication problem, since it is likely to be considered counter-intuitive by a policy audience that the inclusion of a fictitious target should affect our estimates of *historical* dynamics.

This problem is avoided here simply by putting a uniform prior on the amount of trend acceleration, so that it does not affect the marginal distributions of historical parameters. This approach at the same time solves another problem. If acceleration were not 'free' in likelihood terms, the estimation of the random shocks would inevitably be estimated upwards. In words, the results would be shifted towards considering part of the target-attainment to be literally due to luck. The fact that under the current set-up, this effect is avoided, at the same time creates the technical convenience of being able to use the very same simulated sequence of future shocks for different scenarios. Otherwise, doing so would risk creating a spurious upward 'spike' in 2030 even in the 'business-as-usual' trend trajectory.

While it is 'safe' with respect to these problems to impose an improper or proper uniform prior on the amount of trend acceleration necessary to reach the target, and this solution is adequate with respect to modelling SDG target trajectories that are in any case entirely unrealistic (see results section below), this issue deserves additional research in general, since it makes it difficult to include real information on the magnitude of the effects of plausible policy changes in the form of informative priors.

1.2.9 Spill-over effects between education levels

In addition, the target scenarios make explicit that accelerating expansion at one level of the education system will not leave other levels unaffected. In particular, some degree of 'spill-over' to the levels above is to be expected. This effect is modelled by exposing the attainment level above the target level, and the level above that (if any), to an increase in trend drift (at the transformed scale) that is 50 percent respectively 10 percent as large as required at the target level to meet the target.

This can be interpreted as an approximation to cutting the log-odds ratio of transitioning from secondary to post-secondary of the target relative to trend scenario in half for the 'additional' secondary school graduates under the target-achieving trend, and maintaining those new odds into the future. If the model were specified in terms of a logit curve instead of a probit curve, this interpretation would be exact. Parenthetically, as already mentioned above, the reason the model is in fact specified in terms of probits is because this extends more naturally to model elaborations where an underlying Gaussian latent propensity for education is assumed at the individual level. Also recall the preceding discussion concerning the ceiling for post-secondary attainment that maintaining constant transition rates from secondary to post-secondary are not an attractive alternative, because they would imply limiting ultimate post-secondary participation to the level of the current transition rate.

The amount of 50 percent spill-over at the transformed scale was chosen for substantive reasons: there is no reason to expect a targeted boost at one level would actually increase growth at the level above more than the target level itself (suggesting the spill-over should remain below 100 percent), but it seems plausible to expect some upward pressure on post-secondary participation if the pool of eligible upper secondary graduates increases. The reason the spill-over is not specified proportionally to the transition rate from secondary to post-secondary is that doing so would cap a country's long-term participation in post-secondary at the level of the current transition rate, which will often be unreasonably low. If the current transition rate from secondary to post-secondary is 30 percent, for example, and this were held constant, then universal upper secondary attainment would imply merely 30 percent participation at post-secondary, and no further growth or convergence with other countries.

In principle, an attempt could be made to utilise estimated correlations between the drifts at different levels in order to 'endogenise' the amount of spill-over. However, since each country only has *one* past secondary drift and post-secondary drift, these can only be correlated across *sets of countries*. But the spill-over effect will strongly depend on context, and questions such as whether funding for secondary expansion comes at the expense of funding for the post-secondary sector or not. It is not at all clear what the appropriate contextual country sets in terms of spill-over behaviour would be. More importantly, it is clear that the additional secondary expansion associated with a focused effort to universalise that level would be qualitatively different from the past general trend and would not at all represent "business as usual". It is therefore questionable whether the past association between levels could sensibly be extrapolated. It seems preferable, therefore, to make the simple, but transparent, assumptions discussed at the beginning of this section.

As a side note, the same argument explains why there are no secular period effects ('year dummies') included in the model: It is not at all clear that such positive or negative shocks affecting all countries in a single five-year period even exist. This would beg the question whether period effects should not rather be defined at the regional level, for example. At worst, there is a loss of efficiency, as correlation between the idiosyncratic country shocks is not exploited in the estimation. However, from this perspective also, there is no clear reason to expect period effects to be the most important source of such correlation.

1.2.10 Limitations/Constraints

While the above model in many advances the state-of-the-art in long-term education projections, there remain a number of incidental and fundamental limitations.

While most countries of the world are included in the baseline data and the estimation, representing well over 95 percent of the world population, there are some gaps in country coverage. More importantly, these gaps are not random. One category of countries that is difficult to include, but ultimately not consequential in terms of the projections, are small island states. More problematic is the fact that, since the baseline data build on censuses and large-scale surveys, a minimum level of security and state capacity is normally required for countries to be included. Conversely, this means that 'failed states' and countries suffering from violent conflict are underrepresented in the data. Assuming these countries also exhibit below-average rates of educational expansion, this means that overall and regional trends are biased upwards to some extent in their absence. Alternatively, the estimates may be interpreted as being unbiased, but conceptually restricted to representing the range of 'non-catastrophic' scenarios, the same way that even the 'low' projections of global population by the UN do not take into account the possibility of catastrophic disease pandemics.

Another data-related constraint has already been discussed, namely the unavoidable time lag associated with completed attainment. In the present context, this limitation is highlighted further because with baseline data from 2000 to 2010, the inability to pick up on the most recent trends in *enrolments* potentially underestimates the contribution of EFA and MDG related educational expansion to long-term trends. However, the verdict is still out to what extent such a contribution actually occurred at all. While there certainly were cases of strong enrolment growth during that period, the present projections show that significant growth was anyhow to be expected. Indeed, at the aggregate level, our results are broadly consistent with existing extrapolations based on enrolment/attendance, suggesting that, in practice, the time lag of attainment is not particularly problematic.

Perhaps the biggest conceptual constraint is that attainment contains no measure of quality. Neither does enrolment or attendance, of course, and measuring quality is generally recognised as one of the single greatest unsolved challenges in international education statistics, matched only, perhaps, by the challenge of measuring equity. To some extent, this is therefore a data problem, that cannot currently be resolved. Existing efforts to derive general quality indices from international assessments are not without problems, and in any case are currently too limited to country coverage to provide a comprehensive solution. More generalisable perhaps, but even further removed from educational conceptions of quality, are efforts to estimate quality by differences in economic returns to nominally equivalent attainment levels that immigrants from different countries command in the US labour market, for instance. In any case, the challenge of modelling and projecting educational quality cannot feasibly be overcome within the scope of the present exercise. Instead, it is proposed that some sense of the potential impact of quality can be gained from the impact projections. This is discussed further below.

1.3 Implementation

1.3.1 Formal Description

Formally, the core model can be cast in a formula as:

$$y_{c,t,g} = \Phi(\lambda_{c,t,g} + \epsilon_{c,t,g})$$
$$\lambda_{c,t,g} = \lambda_{c,t-1,g} + \tau_{c,g} + u_{c,t,g} - \theta u_{c,t-1,g},$$

where $y_{c,t,g}$ is the share between zero and one reaching a given attainment level (index omitted) in country c at time t among gender g, $\lambda_{c,t,g}$ is the predictor of y at the transfored scale, the ϵ are the 'data error' layer, and the u the random 'shocks' to attainment. The λ follow a random walk, starting from the last position at each step, but potentially retracing a share θ of the previous period's shock. The key parameter of interest for our purposes is τ, capturing the country-specific drift (or 'trend').

The above basic model is complicated further by the presence of gender convergence, which is defined through the expression:

$$\lambda'_{c,t,g=i} = \nu_{c,t} \times \lambda'_{c,t,g=i} + (1 - \nu_{c,t}) \times \lambda'_{c,t,g=-i},$$

and replacing λ with λ' in the definition of y.

In target scenarios, $\tau_{c,g}$ is replaced by

$$\tau'_{c,t,g} = \begin{cases} \tau_{c,g}, & \text{if } t \leq t' \\ \tau_{c,g} + \delta_{c,g}, & \text{if } t > t' \end{cases},$$

where t' is the 'take-off' time for the target scenario (e.g. 2015 for the SDGs) — suitably shifted to account for the age group actually modeled, as described above — and δ is the unconstrained 'boost' required to achieve the target (which is treated as a 'future observation').

1.3.2 Priors and Hyperpriors

In terms of prior distributions, vague priors are specified that only incorporate knowledge of the order-of-magnitude of various effects, as well as logical bounds.

The mean-reversion effect θ has a Beta(1.5, 1.5) prior in the interval (0, 1). The empirical gender convergence factor ν is level and country specific, with prior Beta(1, 5), to ensure a value in the interval (0, 1), strongly skewed towards smaller values. True initial levels are given conceptually uninformative 'flat' priors, but restricted to the interval (-4, 4) to ensure a proper posterior. The idiosyncratic shocks at the probit scale, i.e. the gender, level, year, and country specific epsilons, are i.i.d. draws from a Gaussian distribution with zero mean and standard error σ_ϵ. The additional errors stem from a Gaussian N(0, 0.05) distribution. The (gender, level, and country specific) drift parameters have Gaussian priors centred on regional means (themselves drawn from a Gaussian N(0, 1) distribution), with standard error σ_{trend}. The hyper-priors on variance parameters σ_ϵ and σ_{trend} are Gaussian with mean zero and variance 0.2.

1.3.3 Computational Details

The model was implemented in the 'Stan' software package and posteriors samples generated through MCMC sampling. Chains converge consistently in around 100 iterations, and a total of 500 samples was kept from four chains after discarding burn-in and checking Gelman's 'R hat' split-chain convergence criterion. The number of posterior samples is constrained not only by computation time, but also by the large number of scenario-time-country-level-gender-specific parameters (163 countries, 2 genders, 5 education levels, 2 scenarios, 28 time steps). For each scenario, storage of the results requires more than 5 MB per iteration. However, even 500 samples in fact results in projection quantiles that are sufficiently smooth.

1.3.4 Baseline Data

The empirical historic expansion patterns are estimated on a recent set of global reconstructed time series of completed educational attainment (Lutz et al. 2014). These are disaggregated by country, year in the range 1970-2010, gender, 5-year age groups, and six education levels: none, incomplete primary, primary, lower secondary, upper secondary, post-secondary. The latter is an aggregate category that includes, but is explicitly not limited to, tertiary education. These time series were reconstructed from the most recent available large-scale cross-sectional baseline data. In most cases, that means either censuses or standard international household surveys, such as the DHS. The consolidated and harmonised baseline data were backprojected along cohort lines, accounting for educational mortality differentials. As an illustration of the basic principle, and ignoring said mortality selection, the share of 50-year-olds with at least upper secondary education in the year 2000 informs us of the likely share of *40*-year-olds in 1990. Where possible, these backprojections were validated against historic data sources.

In the present exercise, 163 countries were included that could be nested within GEMR world regions and World Bank income groups. These cover a vast majority of the global population, and most exclusions are small (island) states.

The key advantage of this dataset is firstly its large coverage, that is not limited to countries with historic time series data, and secondly consistency, since all attainment statuses are determined at the same point in time, thus avoiding as much as possible the problem of changing definitions over time. Differences in definitions between countries are harmonised through the ISCED classification scheme and case-by-case validation.

The main disadvantage of this approach is the relatively large time lag. Firstly, the baseline data itself (with censuses normally only conducted every ten years). Secondly, because formal educational attainment can only be assumed to be essentially completed at ages adult ages (depending on the specific level), the effect of very recent or ongoing changes in enrolment trends are not reflected.

2 Attainment projection results

This section presents the results of the educational attainment projection model described above. In all of the following, the reference group in each year are the 15-19-year-olds and their *ultimate* attainment.

2.1 Scenarios

The focus in the previous section was on *how* to handle target scenarios in the projection. Here, we define *what* the scenarios modelled actually assume in terms of attainment dynamics.

The baseline scenario, also called Global Education Trend (GET), has no target, and simply extrapolates the historic country-specific drift. Since the projections are stochastic, this actually represents an ensemble of projected trajectories. We use the path connecting the cross-sectional median values as the single representation of this scenario for visualisation, and to feed into the impact projections in the next section. The full distribution is, however, exploited in some of the analyses below, such as the estimated *probabilities* of reaching the target under the trend scenario.

The SDG scenario assumes that upper secondary education is essentially universal (i.e. exceeds 97 percent) by 2030. In addition, spill-over to post-secondary is assumed, as specified in the model description above.

The 'slow' scenario relaxes the time constraint of the SDG scenario. In other words, it assumes universal upper secondary attainment by 2040, ten years after the SDG target. Note that this also approximately matches a proposed generous interpretation of the SDG target as specifying that in 2030, the *expected* attainment should be at least upper secondary. In other words, that all *transitions* up to upper secondary are universal by that year. This means that upper secondary attainment would be universal among the cohort *entering* school in 2030, not necessarily the cohorts already in school at that time. Since upper secondary schooling typically takes around 12 years, this corresponds approximately to the 'slow' scenario, given our five-year aggregation of ages.

The 'low' scenario relaxes the SDG scenario in terms of the target level, by assuming that *lower* secondary attainment is universal by 2030. This is of interest not only in light of the expectation that the literal SDG target will be unfeasible, but also reflects the fact that until late in the process of negotiating the SDGs, the attainment target was ambiguous as to whether it referred to upper or lower secondary school.

2.2 Aggregate results

Before examining the results for selected countries in the following sub-section (detailed results for all countries are available in an online appendix), we may gain an overall sense of what to expect from aggregate numbers at the global (Fig. 2) and regional (Fig. 3) levels, as well by income group (Fig. 4).

We see that—at this scale—the 'slow' and 'low' SDG target variants have fairly similar implications at the lower secondary level, but that the 'slow' scenario is much closer to the literal SDG scenario at upper secondary. This is not entirely surprising, as the 'low' scenario only gives a boost to upper secondary expansion due to the assumed spill-over effect. However, while this spill-over may seem restrained, note that it is in fact structurally equivalent to the effect of the SDG target scenario on post-secondary expansion, which can be seen to be quite substantial.

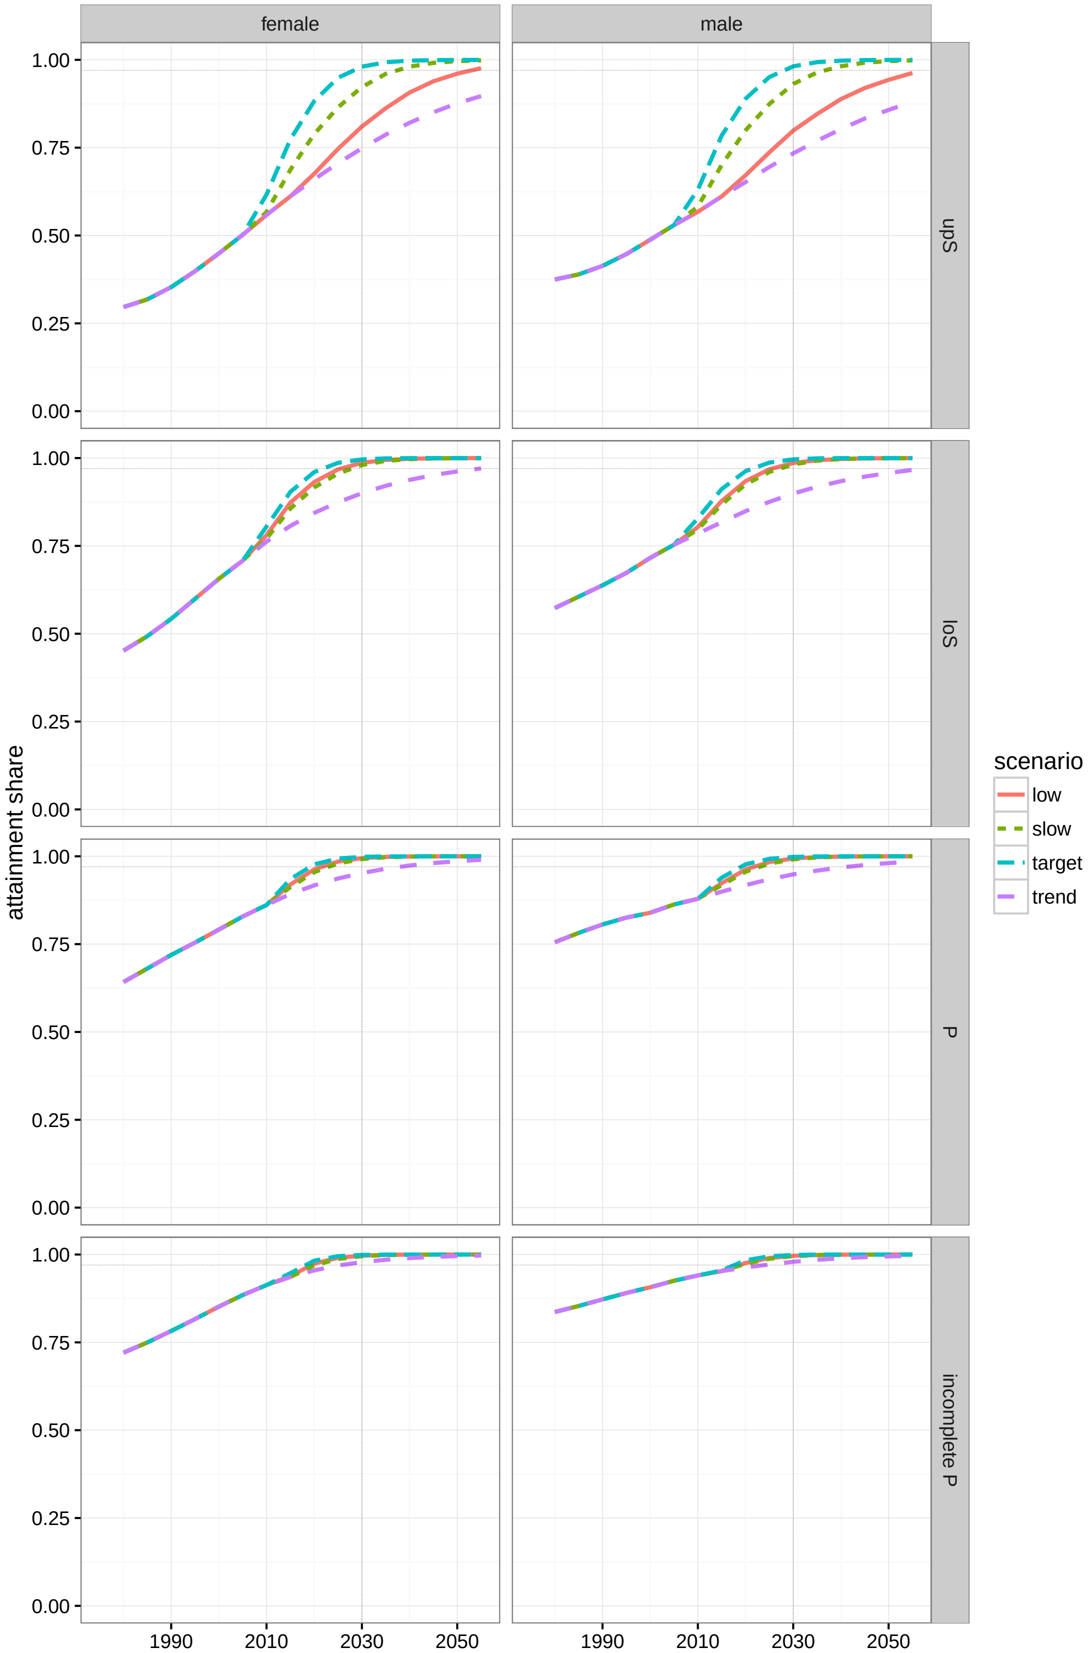

Figure 2: Global projection results. Ultimate attainment of the cohort aged 15–19 at the time shown.

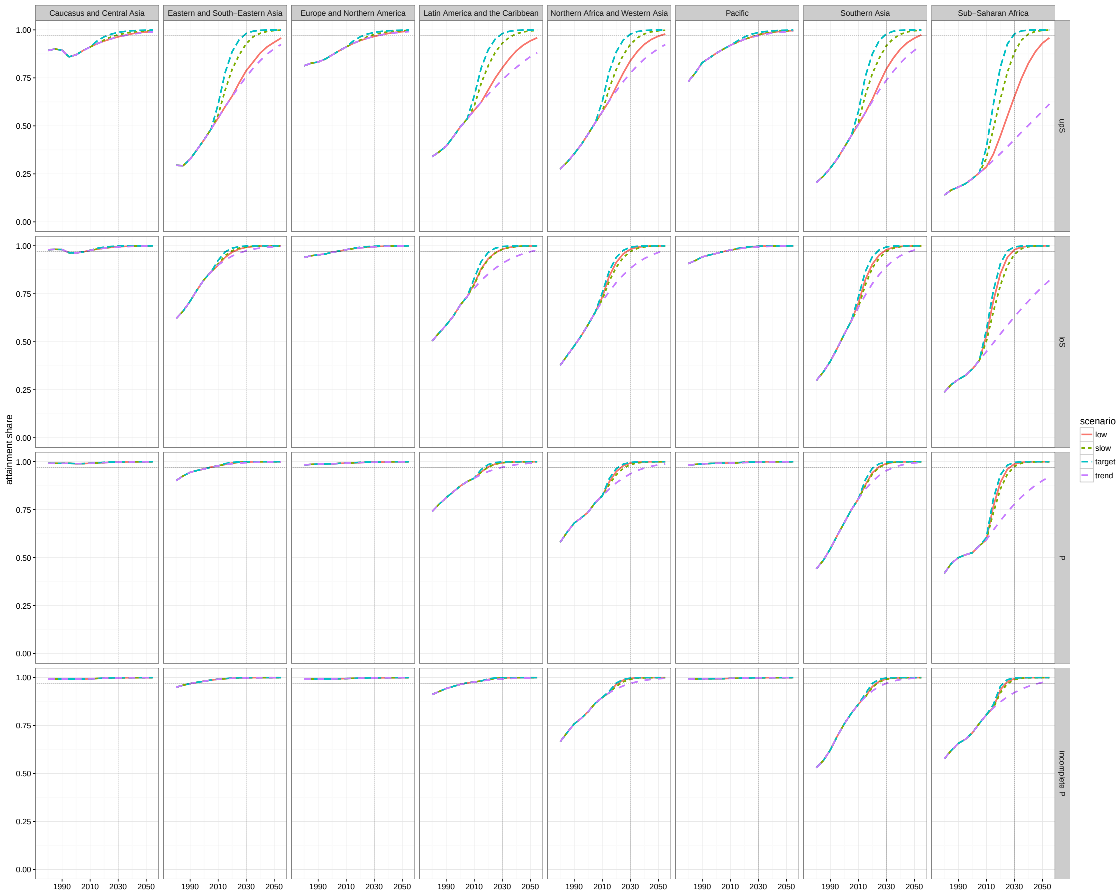

Figure 3: Regional aggregate projection results. Ultimate attainment of the cohort aged 15–19 at the time shown.

What is also evident is that the literal SDG scenario requires additional expansion that sets in immediately, but combined with the expected levelling-off towards the top-end, deviates significantly from a linear 'straight line towards the target'. This will be explored in more detail below.

In terms of the overall magnitude of the challenge, note that under prevailing trends without a sprint towards the targets, not even the old MDG goal of universal *primary* is likely to have been achieved by 2030. Accordingly, the trajectory necessary to meet the secondary school targets of the SDGs involves a decisive break with past trends, else meeting this target is a lifetime into the future.

With respect to the high-income countries that are essentially already 'on track' to universal secondary education, it is worth emphasising once more that just because a given milestone lies in the path of the trend trajectory, this does *not* mean that achieving it is in any way 'automatic'. A significant policy effort will still be required. The message is merely that such a policy effort would not be unexpected in a country in that situation. More generally, in all of the following, the estimated 'probabilities' should not be interpreted as factual statements about the state of the world, but as 'shares of simulated futures' (which, incidentally, is also the correct interpretation of the 'probabilities' issued in the context of weather forecasts). In other words, they are probabilities conditional on countries following future trajectories that are structurally consistent with past behaviour, and neither collapsing into failed states, nor abolishing formal schooling altogether in favour of technology-driven 'on demand' education, which both may have non-zero real-life probability of occurring. As such, a statement such as 'Country A has an x percent probability of reaching the target under current trends' is

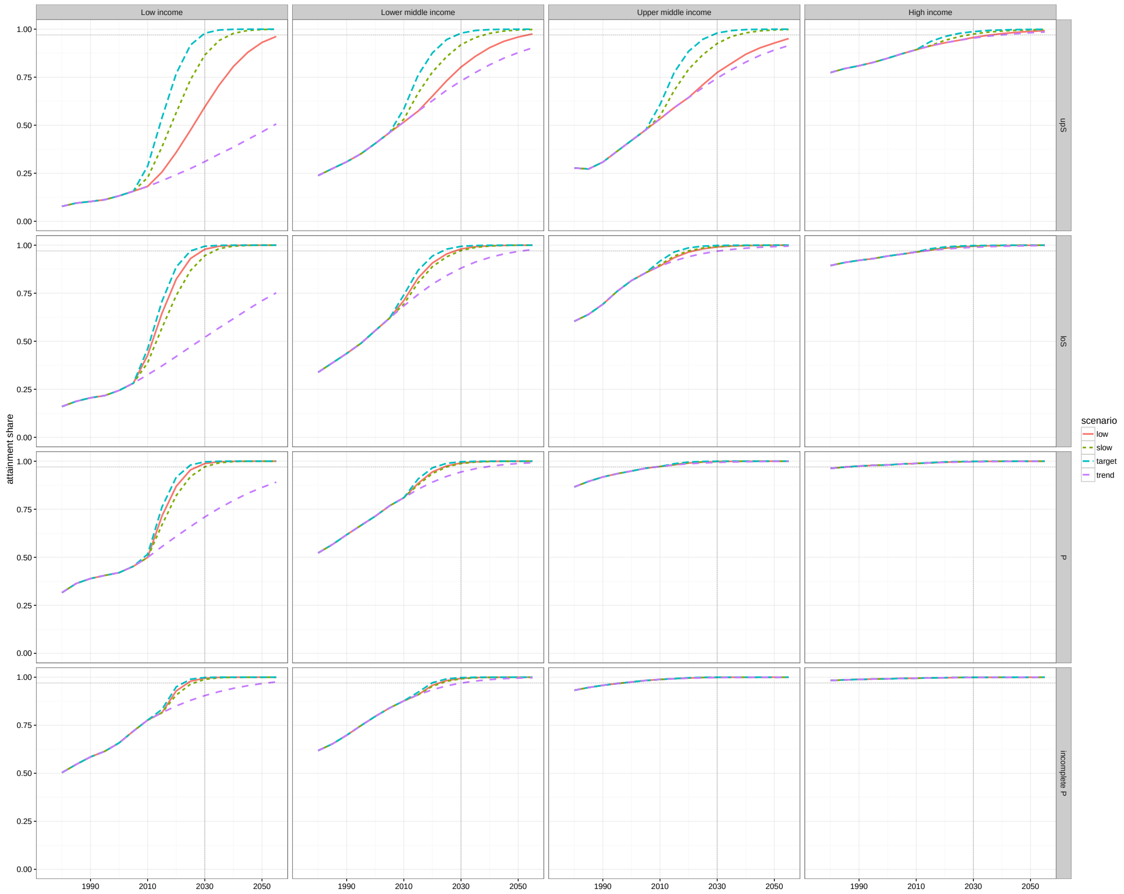

Figure 4: Income group aggregate projection results. Ultimate attainment of the cohort aged 15–19 at the time shown.

correctly interpreted simply as a measure of the degree to which the target deviates from the trend that can be compared across countries. Moreover, the model does not and cannot account for the probability of countries actually adopting additional SDG-motivated expansion policies, or the likelihood of success or failure of such policies. The trend assumes no such *additional* efforts beyond what would be expected based on past record, and the target scenarios assume that the target was indeed attained. Attempted and partial success is not modelled.

As mentioned, attaining a target of universal participation is operationalised as exceeding 97 percent participation. While such 'slack' in the definition is anyhow unavoidable simply because the monitoring indicators are not observed without error, it is clearly problematic from a rights perspective. At some level, it explicitly states that the *most* marginalised and hardest-to-reach literally 'do not count'. However, apart from the fact that there is no alternative at the level of statistical modelling, a tentative answer to this concern is that at that point, the policy focus should perhaps switch from seeing the very hardest-to-reach as a *residual*, to focusing on the positive identification of those excluded.

Returing to the question of missing even the 97 percent threshold, this can occur in rather different ways. Consider a country with 90 percent upper secondary participation that is edging very slowly upwards. It might spend a few decades within a few percentage points of the target. The other extreme is illustrated by a country that somehow manages to go from zero to universal upper secondary participation between 2015 and 2035, but that *in 2030* is still 'only' at 80 percent. In other words, there are potentially large differences between a perspective focused on the *outcome gap* in

2030 relative to the target, versus a perspective focused on the *time gap* between 2030 and actually reaching universal secondary participation.

Table 1: Regional means (population-weighted, see below) of:
(1) probability of universal upper secondary by 2030.
(2) probability of exceeding 80% upper secondary by 2030.
(3) median gap to universal upper secondary in 2030.
(4) probability of universal upper secondary by 2040.
(5) median year of universal upper secondary.
(6) probability of universal lower secondary by 2030.
(7) probability of exceeding 80% lower secondary by 2030.
(8) median gap to universal lower secondary in 2030.
(9) probability of universal lower secondary by 2040.
(10) median year of universal lower secondary.

region	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Caucasus & C Asia	58.1	98.0	3.5	71.6	2020	90.3	99.9	0.7	95.1	1976
E & SE Asia	9.1	33.8	24.5	11.8	2054	71.8	98.2	2.7	85.6	2017
Europe & N America	60.5	97.3	3.4	76.1	2020	94.7	100.0	0.6	97.6	1991
Latin A & Caribbean	0.5	34.6	26.3	2.2	2066	12.0	88.5	9.5	33.5	2039
N Africa & W Asia	3.1	49.7	23.0	12.8	2055	26.5	74.0	12.8	46.0	2038
S Asia	0.1	32.5	26.4	1.6	2058	6.7	82.5	11.1	31.2	2041
SSA	0.0	5.9	57.5	0.2	2085	2.6	32.4	37.6	6.2	2079

Table 2: Income group means (population-weighted, see below).

income group	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
High income	48.3	91.5	5.9	62.9	2028	85.2	99.3	1.3	91.3	1996
Upper middle income	2.2	29.5	25.9	4.4	2059	66.0	97.7	3.2	78.8	2021
Lower middle income	2.6	33.5	27.1	5.7	2059	9.8	79.4	12.2	32.3	2041
Low income	0.0	3.9	69.6	0.2	2092	0.7	12.3	50.3	3.3	2095

Note that in the above, the aggregate figures for groups of countries do not represent the results for the 'pooled' populations directly, but instead are weighted averages of the country results in those regions. One issue this creates is the choice of country weights to use in calculating the weighted averages. Different, equally reasonable, answers may be possible with respect to whether to base the weights on the total or school-age population, and which reference year to use. In the above, the average total population during the period 2010–2030 was used. Note that other summaries or syntheses of our results may prefer alternative weighting schemes, with marginally different numeric results. Projected years of target attainment are more sensitive to this than attainment levels and associated probabilities.

2.2.1 Projected probability of universal upper/lower secondary by 2030

A first indication of which countries face the greatest challenge is to examine the probability of meeting the target under prevailing historic trends. For upper secondary, these figures contain few surprises. Moving through the regions in the (alphabetical) order shown: Many former socialist republics are already 'on track' to universal upper secondary participation, with females outpacing males. In East and South-East Asia, the frontrunners South Korea, Singapore, and Japan will surprise no-one with their square aim at meeting the target. Females in Mongolia, Malaysia, and the Philippines follow far behind this group, but are still fairly well-placed compared to other middle-income countries in the Global South. With the exception of Canada, the top of the list for Europe and North America is dominated by former socialist countries, with traditionally high educational participation. In Latin America & the Caribbean, only Puerto Rico stands a reasonable chance of meeting the target without a historic amount of effort. In North Africa and Western Asia, it is noteworthy that—under current trends—the target is much more likely to be met for females than for males, reflecting the rapid 'catch-up' for females in the last few decades, and that there are very large differences between the Gulf countries, that are similar in many other relevant respects, such as income. Finally, in South Asia and Sub-Saharan Africa, there is practically no chance even for the regional front-runners to achieve the target unless a clear break from past trends is achieved.

With respect to the less ambitious target of universal *lower* secondary, one might say that 'lower secondary is to South Asia and Sub-Saharan Africa what upper secondary is to Latin America and the Middle East', that is: distinctly in the realm of possibility for the frontrunners, but no-one else. Noteworthy rank reversals are that Jamaica, Tunisia, and South Africa rank much higher on lower than on upper secondary in their respective regions.

2.2.2 Projected probability of upper/lower secondary exceeding 80 percent by 2030

It could be that the above paints an excessively pessimistic picture if there are many cases where countries would—under current trends—not *quite* make the target, but get reasonably close. This column is structurally the same as the previous ones, but with respect to reaching 80 percent of lower secondary attainment instead of universal. In this case, many fewer cases have a probability below 1 percent.

On this measure, with respect to upper secondary, the comparison is enlightening between Latin America & the Caribbean, as well as North Africa & Western Asia on the one hand, and South Asia as well as Sub-Saharan Africa on the other. In the former regions, outside of a handful of frontrunners, the projected probability of fully achieving the SDG target of *universal* upper secondary was marginal, and only marginally higher than for countries in the latter group of regions. However, many countries in the former regions are likely to get close to universal upper secondary than in the latter regions. Also worth highlighting, in light of the global aggregate figures, is that India would actually be *expected* to make great strides towards mass upper secondary attainment irrespective of the SDGs, and even though the most recent policy efforts in this direction were not even part of the baseline data.

Considering lower secondary results in a strikingly different picture. Clearly, mass participation (if not quite universal) at this level is vastly more achievable for many developing countries.

2.2.3 Median projected upper/lower secondary attainment gap remaining by 2030

A different way of looking at countries that do not quite reach the target is to examine no the probability of meeting some lower threshold as above, but the actual projected gap. While this gives a similar ranking in many cases, it does not necessarily do so, and even when the top countries are the same, conveys a different piece of information.

In effect, for both lower and upper secondary this column changes little, with almost identical rankings. It does, nevertheless, put the projected probabilities for fully meeting the SDG target into perspective more clearly, if we consider particular examples. For females in Jordan, say, while the projected 12.8 percent probability of meeting the target seems small, the figure in this column shows that this group is actually expected to get within striking distance of the target even under prevailing trends and without additional SDG-inspired drive.

Conversely, the *large* expected gaps provides a sense of the magnitude of effort required in the most challenging settings. There are no big surprises here, perhaps apart from the fact that the largest gaps to be bridged in Europe would still be among the largest in South Asia, or North Africa & Western Asia. While this might be interpreted as a success in defining targets that are still relevant to relatively high-income countries, this comes at the cost of a target that leaves demoralising gaps to what would otherwise be expected to happen in many SSA countries.

2.2.4 Projected probability of universal upper/lower secondary by 2040

It is possible for a country on a rapid expansion path to be quite far below 100 percent in the target year 2030, but to be not very far behind in the time dimension. We therefore consider the projected median probability of meeting the lower secondary target by 2040 instead of 2030.

For upper secondary, this perspective mostly results in a more optimistic view on the Middle East, were several countries that were shown above to have limited prospects of meeting the SDG target outright are actually merely lagging in time. This contrasts sharply with South Asia or SSA, where an extended target date makes little practical difference to making the target more achievable without massive trend breaks.

Similar conclusions apply to lower secondary. We can see that many countries in Latin America & the Caribbean, as well as in North Africa & Western Asia, are currently merely somewhat 'too slow' rather than 'too low' to universalise lower secondary by 2030. And again, South Asia and SSA are 'one schooling level behind', in that their position with respect to lower secondary is comparable to that of the aforementioned regions with respect to upper secondary.

2.2.5 Absolute growth

While we mostly analyse relative changes, at least at the global aggregate level, it is worthwhile to consider the implications of the SDG targets in terms of the *absolute* educational expansion required. fig. 5 shows the total number of 15-19-year-olds that would complete the stated school level in each year, relative to the year 2010. Note the *logarithmic* scale.

Most striking is clearly the extent to which Sub-Saharan Africa faces challenges in terms of capacity expansion that are utterly unprecendented. In every other world region, by the time mass participation in secondary had become part of the policy agenda, the average number of children had already dropped significantly. Combined with low starting levels of secondary participation in many SSA

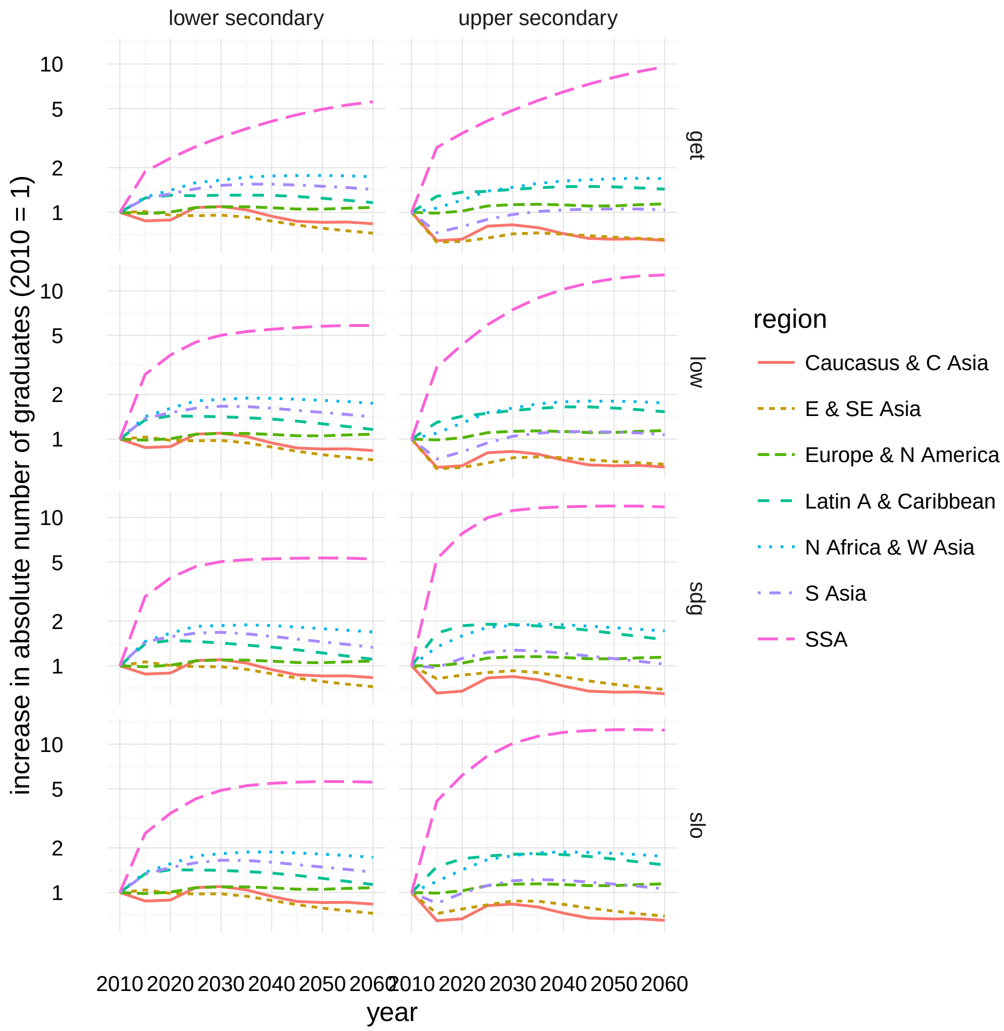

Figure 5: Growth required relative to 2010 baseline in *absolute* number of graduates per 5-year cohort a the level shown.

countries, we see that a five-fold capacity increase would be required at lower secondary, and a ten-fold increase at upper secondary. By contrast, even in North Africa and the Middle East, as well as in South Asia, at most a doubling of secondary school numbers would be required. In Central, East, and South-East Asia, even universal participation can no longer offset shrinking cohort sizes, and secondary schooling is practically certain to contract in the medium term, at least in terms of the number of students.

Notably, while the SDG targets, however interpreted, certainly demand even greater absolute growth from SSA, even the trend scenario implies a multifold increase. Recalling the interpretation of the trend, this means that SSA faces a large capacity challenge just to maintain its position relative to other regions. Conversely, this means that unless upper secondary numbers are increased by a factor of five by 2030, SSA will fall *even further behind*, even if countries everywhere effectively ignore the SDG targets. Nevertheless, although quantitative expansion similar to the SDG scenario is assumed to occur under the trend scenario *eventually*, a tenfold increase over 40-50 years (trend) is a very different proposition than a tenfold increase over 10-15 years.

Careful readers may also notice that in the long run, the 'slow' and 'low' scenarios actually require higher absolute capacity than the more ambitious SDG scenario. That is because by 2060, the latter is already likely to have resulted in smaller birth cohorts.

2.2.6 Country-level results

Graphs of the country-specific trajectories for each scenario are shown in the online appendix.

As mentioned above, different age thresholds are applied at each level for determining which age groups are included in the estimation of past trends. In addition, the most recent observed attainment of the 15-19 age group is plotted for complete/incomplete primary level, but not included in the estimation. These data points clearly show two counter-acting effects, depending on whether they are clearly below or above the otherwise consistent trend, namely late attainment on the one hand, and accelerating improvement in the youngest cohorts that perhaps benefited from EFA on the other. Unfortunately, since both do occur in the data, this age group really does need to be excluded. This fails to capture some recent improvements, but including them would result in too many artificial declines. As a matter of fact, it seems that even the 20-24 suffers such problems at secondary level in some countries, especially in the Caucasus and Central Asia region, but also in some European countries.

2.2.7 Benchmarking against regional front-runners

As is evident, the acceleration that would be required of some countries to actually meet the SDG target of universal secondary education is entirely unprecedented. One simple way of assessing what order-of-magnitude of progress could reasonably be expected.

Obviously this completely disregards differences in income, institutions, tradition, governance, and policy in different countries. Countries within a region are not necessarily 'similar' in the relevant aspects. Even if they were, what worked in a similar country cannot necessarily be successfully replicated elsewhere. We pay attention to these caveats by focusing on the 'negative' interpretation. This interpretation is that the estimates in Table 3 provide further evidence that the SDG education targets are unrealistically ambitious: for the vast majority of countries that are far from the goal, even expanding at the fastest average pace any country in their region has ever done would not be sufficient to meet the SDG education target. This is true even for many industrialised countries, and for *every single country* in South Asia and SSA. In short, the expansion required to meet the SDG target is literally unprecedented. Universalising lower secondary schooling instead is much more achievable in South Asia in relation to historic benchmarks, but in SSA even this less ambitious goal would require unprecedented expansion.

Table 3: Percentage of countries in each region for which matching the highest rate of educational expansion previously achieved in their region would be sufficient to meet the stated target.

region	universal upper secondary by 2030	universal upper secondary by 2040	universal lower secondary by 2030
Caucasus and Central Asia	11	18	22
Eastern and South-Eastern Asia	3	5	14
Europe and Northern America	9	16	27
Latin America and the Caribbean	0	1	7

region	universal upper secondary by 2030	universal upper secondary by 2040	universal lower secondary by 2030
Northern Africa and Western Asia	1	4	14
Southern Asia	0	0	3
Sub-Saharan Africa	0	0	1

Note that both here and in the following section, the results are with respect to the simulated paths towards the targets that result in 2030 attainment levels spread over the range from 97 to 100 percent, as would be expected across all countries if even the slowest expanding succeed in crossing the 97 percent threshold. With a more forgiving interpretation, specifically if the probabilities of each country barely reaching the 97 percent threshold are considered instead, the numbers would be somewhat higher, especially in higher-income regions. However, the qualitative conclusion that the historic benchmark is insufficient for practically all countries in South Asia and SSA remains unchanged. These alternative figures are available on request.

2.2.8 Nonlinear expansion and intermediate targets

As discussed above, a deliberate choice was made to model the path towards the SDG target variants as 'organic' expansion trajectories that are accelerated, but in terms of their general pattern nevertheless match typical past experience. In other words, we examine what it would look like if countries, on their way to achieving the target, expand education the way rapidly-expanding countries do, in fact, expand. In particular, this includes a 'levelling-off' as universal participation is approached, whereby the 'last few percent' become more and more difficult to reach. In addition ta matching past empirical patterns, this also fits well with general expectations. This issue of nonlinear expansion is important, because ignoring it potentially results in misleading policy assessments. In particular, determining whether countries are 'on track' to achieving the target based on whether they have — so far — followed a *linear* path towards the target is likely to overestimate their chances of actually reaching the target on time, since such intermediate progress leaves no room for later deceleration.

This effect is quantified in the following table, that shows the share of simulated target-achieving trajectories that passed within 2 percentage points of linearly-interpolated milestones at 5 and 10 years into tho SDG time horizon. As is evident, outside of the high-attainment regions, and including such a relatively advanced region as Latin America, reaching the linear milestones for universal secondary education does not at all indicate that a country is 'on track' to reach the target. On the contrary, countries that pass these milestones are already so far behind that it is already extremely unlikely they can reach the target without even further acceleration beyond whatever speed-up they already achieved during the initial part of the SDG period. This also applies to a marginally lesser degree to the target of universal lower secondary.

Moreover, the differences between the regions imply that the least advantaged regions are potentially disadvantaged further by the application of linear milestones, because when they are judged to be 'on track' according to the linear milestones and perhaps therefore less in need of additional support, they are actually still less likely to actually reach the target.

Table 4: Share of target-reaching trajectories in each region that exceed by at most 2 points the linearly interpolated "stepping stones" between starting point and target.

region	5-year linear stepping stone towards universal upper secondary by 2030	10-year linear stepping stone towards universal upper secondary by 2030	5-year linear stepping stone towards universal lower secondary by 2030	*10-year* linear stepping stone towards universal lower secondary by 2030
Caucasus and Central Asia	38	35	78	77
Eastern and South-Eastern Asia	15	14	31	29
Europe and Northern America	22	21	81	79
Latin America and the Caribbean	1	1	12	11
Northern Africa and Western Asia	1	2	12	10
Southern Asia	0	0	5	4
Sub-Saharan Africa	0	0	1	1

As noted in the preceding section, these probabilities are with respect to the interpretation that a group of countries reaching the target means their attainment levels will be spread over the range from 97 to 100 percent, rather than every country barely crossing the lower threshold, but even the latter, laxer, specification would not change the substantive conclusions.

2.3 Summary

We can observe that at the country level, just as at the regional level already shown, the extent to which meeting the SDGs, however interpreted, requires a deviation from past trends varies considerably. This ranges from getting close to meeting the target being almost expected based on past trends in many OECD countries, to requiring a noticeable, but perhaps not infeasible push in Indonesia or Iran, for instance, to requiring trend changes, as in Mexico, that perhaps strain credibility but are not entirely without precedent, to 'scenarios' that at best serve to bring into stark relief the utter implausibility of universalising upper secondary education in Chad, say, within 15 years.

At the risk of over-simplification, we might say that in Latin America & the Caribbean, as well as in North Africa & Western Asia, the main challenge is to accelerate existing dynamics, by contrast with large parts of South Asia and Sub-Saharan Africa, where the challenge is better framed in terms of raising the level. With respect to these two groups of regions, we also note that a 'more level playing field' that respects different starting levels would be to aim for universal *upper* secondary in the former, but universal *lower* secondary in the latter.

It is nevertheless of interest to gain a sense of what the contribution of educational expansion to other SDG goals could be, because even partial progress towards the education target may be worthwhile (including in purely economic terms) even from a purely instrumental perspective.

3 Projections of other development dimensions conditional on education

As the review in the first part of this report showed, education has strong linkages with virtually all other dimensions of sustainable development, both direct and indirect. Most (though not all) of these are positive reinforcements. Accordingly, trajectories of educational progress towards the SDG's education targets are not only of interest on their own terms (as analysed in the preceding chapter), but also in terms of their implications for other development outcomes, which is the focus of this chapter.

3.1 Modelling educational development impacts

3.1.1 Second-stage models selected

As outlined at the end of the literature review, not all development dimensions are equally suited to be analysed from the perspective of the potential contribution of educational expansion.

Firstly, goals and targets associated with well-defined indicators are required for a quantitative model of the relationship with educational attainment. Secondly, the existence of strong theoretical arguments and evidence for a relationship with education makes for more meaningful analyses, and takes us closer to a causal interpretation. Thirdly, there must be some amount of historic data on which the quantitative *strength* of the relationship with education can be estimated. Finally, in supporting the overall argument that education is strongly-connected within the 'network' of SDG interactions, it is desirable to model outcomes drawn from different dimensions of sustainable development.

These criteria justify the selection of the following outcomes. Two of these, namely health and economics, are areas where the effect of education is well-known in principle, but the specific difference that achieving the education SDGs could make has not been known until now. The third, disaster vulnerability, is an area where a strong relationship with education can be identified, but there remains little understood among actors in that field.

Specifically, in the following we present model results for how infant and child mortality (relating to SDG 3.2), specifically survival to age 5, differs according to education scenario, as well as adult life expectancy (reflecting SDGs 3.3 and 3.4). In the economic dimension, we do the same for aggregate national economic growth (relating to SDGs 8.1 and 8.2), the absolute extreme poverty headcount rate (SDG 1.1), and — more tentatively — the relative position of the 'bottom 40 percent' (SDG 10.1). Finally, we move to modelling disaster deaths (relating to SDGs 1.5, 13.1, and 11.5).

The domain-specific models of the relationship between education and these outcomes are discussed in the respective sections that follow.

Not explicitly modelled are other SDG targets that fully meeting the SDG education target of universal secondary education would contribute to or even achieve 'by definition' even though they are outside of 'the education SDG 4', such as SDG 8.6, calling for a substantial reduction in the proportion of youth not in employment, education or training.

3.1.2 Notable exclusions

The selection according to these criteria clearly does not amount to a judgement on the relative 'importance' of different SDGs. This follows immediately from the very first criterion above, quantifiability, which is a necessary requirement for modelling, but most definitely not a measure of importance.

Notable exclusions of outcomes that would be desirable to model in principle include health impacts in terms of Disability-Adjusted Life Years (DALY), for example, but also the contribution of rising educational attainment to climate change itself, either in the form of carbon emissions, or even at the scale of degrees centigrade directly.

Unfortunately, explicit numeric estimates of the contribution of an SDG education scenario to these outcomes is infeasibly, or not prudent, for different reasons.

With respect to DALYs, which are based on disease-specific prevalences, data on education differentials is unfortunately only available for very few specific diseases. Moreover, the way 'DALYs saved' are conventionally calculated is partly already based on *projections* of how disease-specific morbidity/mortality rates will develop in the future. This implicitly already accounts for expected improvements to education, and applying education differentials on top of the published DALY rates, in combination with compositional change in terms of attainment, would effectively be 'double-counting' the education effect. Even if this issue were ignored (and it appears to be routinely ignored in estimates of the benefit of technological progress in terms of DALYs, which strictly speaking suffers a similar problem), education is likely to have quite different effects on morbidity and mortality. This means such a model would be highly complex and is not feasible as one outcome study among several.

Obtaining a customised estimates of the effect of SDG-driven educational expansion (through income and consumption) in terms of numeric increases in global average temperature is not feasible because, at the time of writing, it would require running the entire chain of IPCC models.

3.1.3 Education 'quality' as strength of association with outcomes

It would be desirable to be able to investigate the potential impact of education *quality* on other development outcomes. However, it is not practicable to model this explicitly. To begin with, there is a lack of agreement on what exactly 'educational quality' is, and how to measure it. For any particular measure, data coverage across countries and time is severely limited. It is unsound to assume that the quality of the education today's 50-year-olds in some country received as children is equal to whatever quality measure is — at best — available for a recent year. Time series of quality indicators would therefore definitely be required that are simply not forthcoming. It is also not at all clear that the same quality markers would be relevant for different development outcomes.

Some sense of the role of education quality can arguably be gained within the existing framework, by examining the strength of the association between education and other development outcomes as a proxy for quality. Intuitively, if education makes us live longer, then *better* education should make us live even longer.

For impact models that simply use educational attainment as a regression input, and where the strength of the education effect is a single number, such as the disaster deaths model further below, this is straightforward to implement by varying the magnitude of the relevant regression coefficient.

Unfortunately, this approach is not always possible when each attainment level has its own coefficient, especially when these are defined in dependence on each other and underly a time trend. The way the association between a mother's attainment and child mortality is currently modelled in the Wittgenstein Centre population projections, for example, is such that the effect of different attainment levels is defined relative to the effect of upper secondary, which in turn follows a defined trend. If the reference level were 'no education', it would be possible to vary the strength of 'the effect of education' as a proxy for education quality, simply by proportionately re-scaling the coefficients associated with higher attainment levels. However, similarly reducing the education differential by reducing the size of the coefficients when the reference level is secondary education would effectively *raise up* the health outcomes of the less educated, rather than diminish the health premium of those with secondary schooling. At the same time, redefining the reference level is complicated by the fact that the secular trend is specified in terms of the upper secondary baseline. As a result, varying the strength of the education effect in the desired way involves recalculating the secular mortality trend.

However, in such situations where varying the education coefficient is not straightforward or impracticable, the comparison between the different attainment scenarios may serve as an indication of the role of quality instead. For instance, if the full SDG scenario is met at the expense of making upper secondary only as good as lower secondary is now, then that impact has already been estimated by the 'low' scenario. And conversely, if super-quality lower secondary conferred the same advantage in the future as upper secondary does now, this is already estimated by the 'sdg' scenario. This interpretation would suggest that the potential role of quality is of a similar magnitude as the quantity: SDG with deteriorating quality could easily halve the benefits, generaling from the comparisons between the universal upper ('sgd') and universal lower ('low') secondary scenarios explored below.

3.2 Health

Health is a crucial dimension of sustainable development, and health goals are prominent among the SDGs. In addition, the relationship between education and health is known to be robust, and is relatively well-studied. This is reflected in the fact that, unlike the link with economic outcomes and disaster vulnerability discussed further below, mortality differentials by education are an intrinsic component our population projections, rather than being modeled post-hoc. Accordingly, the differences in the simulated outcomes include interactions such as the fact that the education-induced improvements in average under-5-mortality are attenuated by lower average fertility of the more educated mothers. The exact assumptions included in the model, and the evidence based on which they rest, are fully documented in Lutz et al. (2014). Note that a general SDG population scenario would take into account the implications of the health goals, for instance. Here, the 'SDG scenario' deliberately maintains existing trend assumptions for fertility and mortality, in order to isolate the potential contribution of educational expansion.

With infant and child mortality in mind, we first examine how the education profile of women of child-bearing age develops over time under the SDG scenario for education compared to the trend Fig. 6. The most important difference for the issue at hand is the complete disappearance of women with very low education under the SDG scenario. The impact is greatest in SSA, where the largest share of women with little education would otherwise be expected to remain under the trend. As mentioned, the education profiles of *mothers* is likely to improve less rapidly, due to the higher average fertility of the less-educated. However, there is also some evidence that child health also benefits from community-level effects and the general diffusion of healthy practices and behaviours. Such effects in turn would suggest that the benefits would be greater than suggested by the changing education distribution of individual women.

In terms of the outcomes in the form of non-survival to age 5, fig. 7 displays the estimated impacts at the same two period cross-sections: once in 2030, and for 2050. The rationale for an extended time horizon is that the later 'SDG cohorts' who complete their schooling close to 2030 will have most of their children some time after that. Indeed, it is evident that the education effect continues to grow betwene 2030 and 2050 even in absolute terms, despite the fact that the overall mortality levels are significantly lower at the latter time. We see the first evidence of a recurring theme, namely that the benefits of universal lower secondary schooling are roughly half of those of universal upper secondary either. Achieving either can be expected to make a meaningful contribution to reducing

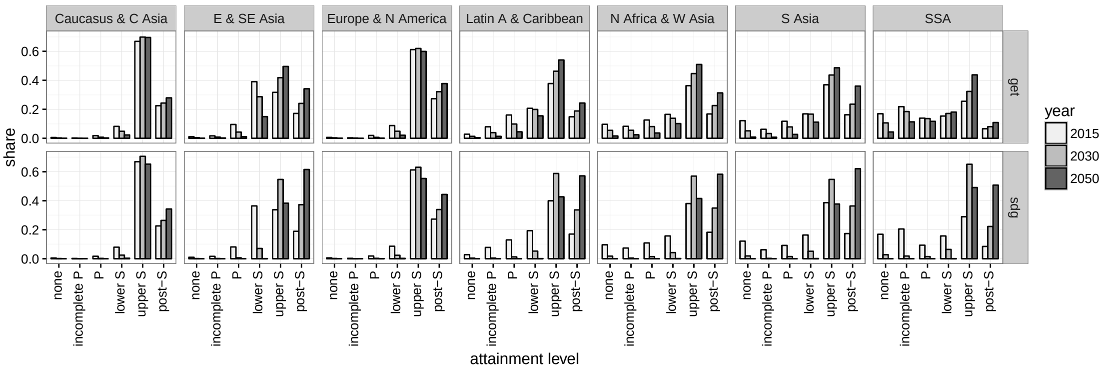

Figure 6: Change in attainment profile of females aged 20–29 over time.

infant and child mortality. This includes even regions with relatively low overall child mortality, such as East & Southeast Asia, where even a numerically small average difference of 1–2 deaths per 1,000 represent a 10–20 percent difference amounting to tens of thousands of avoided deaths.

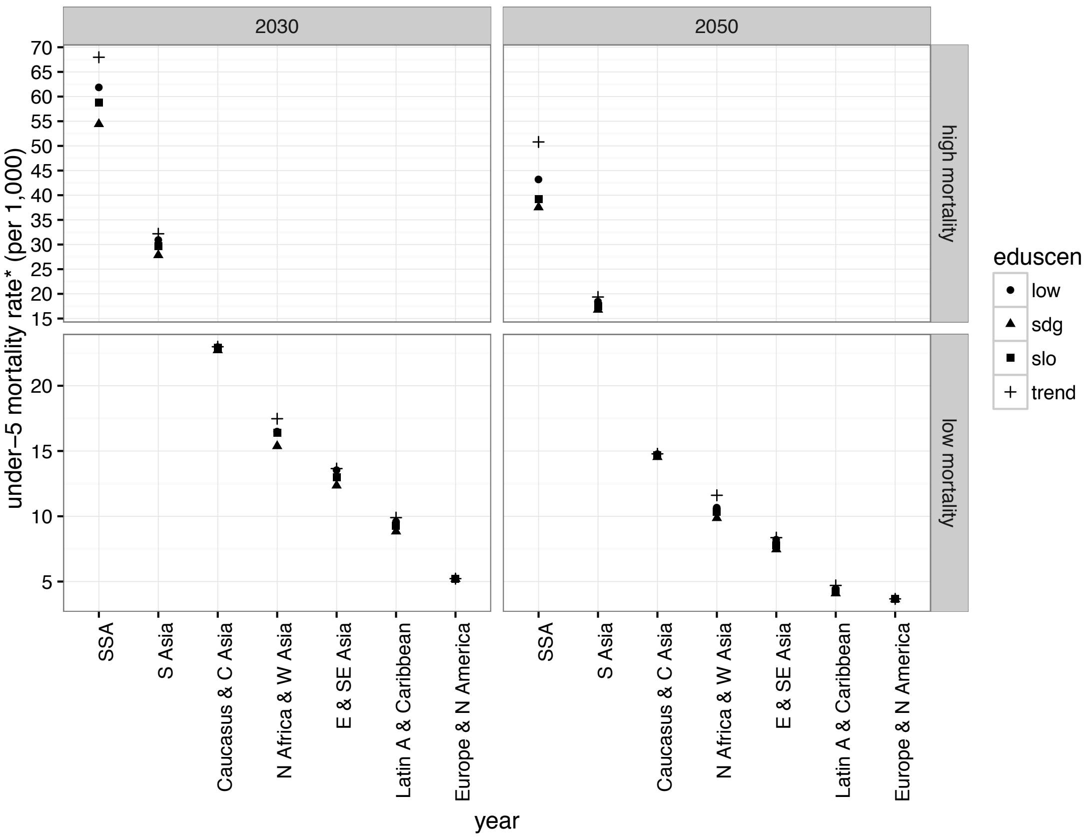

Figure 7: Under-5 mortality by region and scenario.

Nevertheless, it is the high-mortality settings that are our greatest concern, and fig. 8 shows the trajectories for those in greater detail, and in terms of relative changes compared to the baseline in 2015. Up to 10 percentage points in the drop of under-5-mortality may be added by achieving the SDG education target of universal upper secondary schooling, even by 2030, before all the beneficiaries have actually begun their childbearing. Moreover, in SSA, the decline in child mortality may well begin to slow down in the absence of additional educational expansion.

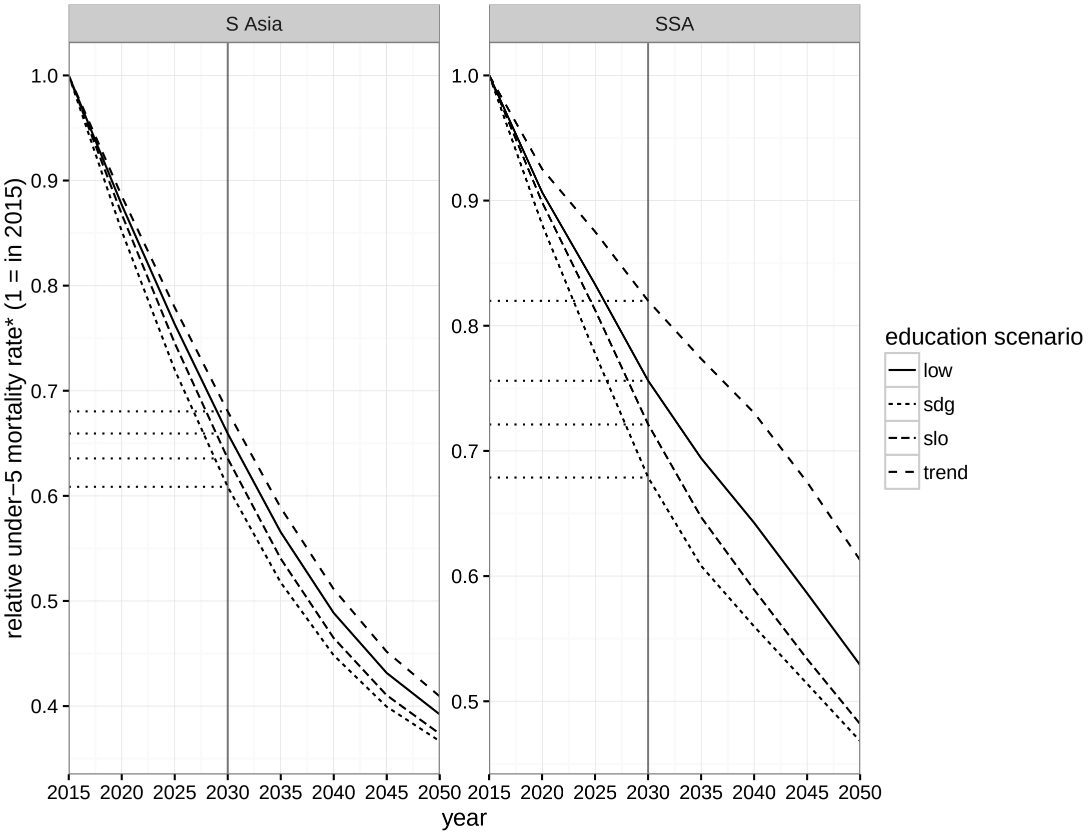

Figure 8: Under-5 mortality over time for highest-mortality regions.

While these effects may be considered to be only moderate, it is important to note that they represent the differences between SDG scenarios and a baseline that in itself assumes quite significant educational expansion based on existing trends. Moreover, much of the *additional* expansion implied by the SDG targets occurs at upper levels of education, while the most dramatic gains in terms of child health occur at the bottom end of the education ladder.

A direct comparison with the estimates of the historical contribution of education to declining child mortality of Gakidou et al. (2010) is not possible, because their analysis concerns years of schooling rather than attainment, and is benchmarked against *no* educational improvement. With these caveats in mind, the additional benefit of the SDG scenario over and above prevailing trends estimated here for SSA is similar in magnitude to the contribution of educational expansion during the period 1990 to 2010 estimated by Gakidou et al., at around 15 per 1,000 child deaths less.

In terms of the potential contribution of the education SDG to improving *adult* health, fig. 9 displays modelled trajectories of average remaining life-expectancy at age 15 in SSA. In other regions the impact is marginal. Despite significant differences in attainment-specific life expectancy even in high income countries, it is unsurprising that the change to overall life expectancy is very modest, given the relatively short time horizon and modest increases to the stock of total population attainment. Even in SSA, the cohort benefitting from the SDG scenarios have not reached ages of high mortality even by 2050. Extending the time horizon further does not provide additional insight, because in the very long run, convergence assumptions drown out most of the education differentials.

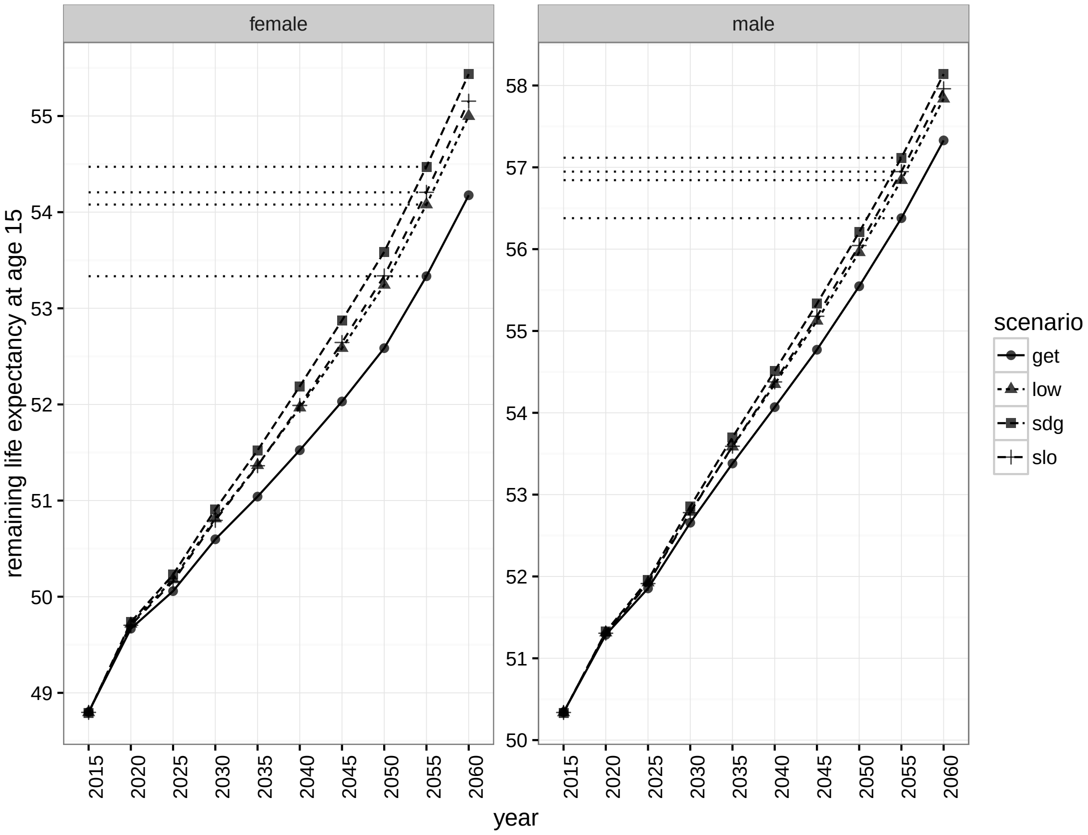

Figure 9: Projected adult life expectancy (remaining life expectancy at age 15) by scenario, for highest-mortality regions.

3.3 Economic growth and poverty reduction

The economic effects of human capital expansion are among the most widely-studied, and have already been commented on in the literature review. However, only relatively recently has this literature become sufficiently sensitive to the importance of accounting for age structure in determining the expected economic benefits of educational expansion. Of particular interest in light of the SDG goals, and also in terms of readily-available models which can be applied to our education scenarios, are aggregate economic growth, and extreme 'dollar a day' poverty (currently defined as USD 1.25 at 2005 PPP, or equivalently at USD 1.9 in 2011 PPP).

In the following, national income projections are obtained using the model by Crespo Cuaresma (2015). The income projection framework combines population projections by age and educational attainment level with an aggregate production function estimated using historical data. Human capital dynamics are assumed to have two distinct effects on income per capita. On the one hand, improvements in educational attainment affect labour productivity. On the other hand, total factor productivity is also affected by human capital through its effect on technology creation and adoption. The results of applying this model to our SDG education scenarios are shown in fig. 10. While the main scale is in logged GDP per capita relative to its level in the year 2000, the difference between the trend and SDG scenarios in 2050 is additionally translated into straightforward percentages.

Because the individuals benefitting from educational expansion during the period 2015–2030 have to enter the labour force in significant numbers before being able to make much of an impact, it is

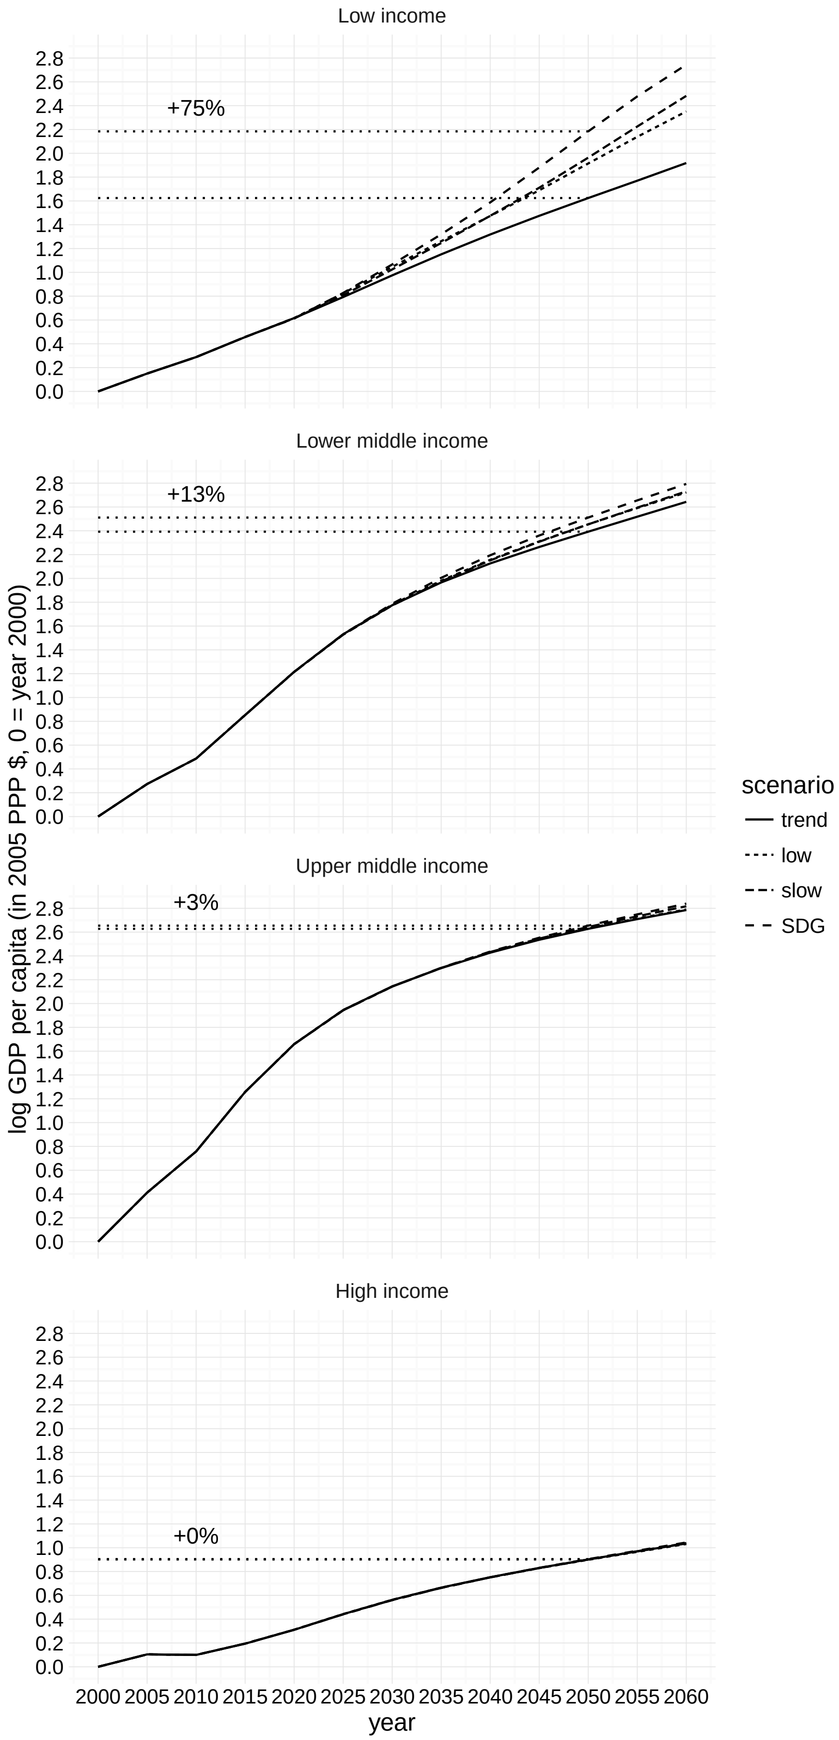

Figure 10: Projected GDP growth by scenario.

unsurprising that meaningful growth effects are delayed until long after the SDG target year of 2030. Indeed, in high and upper-middle income countries, the additional growth expected from SDG-driven educational expansion is minimal. Not because education does not matter, but because many of these countries anyhow have high and increasing levels of secondary participation even under the trend scenario, so that the SDG effect is at most marginal.

This plays out rather differently in low-income countries. While universalising upper secondary education would in principle be expected to make a large difference to their growth prospects in the long run, to the tune of an additional 75 percent by 2050, these countries are of course also farthest from realising this goal. However, even universal secondary education may actualise around half of this gain. Again, this understates the contribution of educational expansion as such, because significant educational growth is expected even under the trend scenario that provides the baseline here.

Arguably of greater importance in terms of the SDGs than overall economic growth is the goal of poverty reduction. Poverty impacts here are approximated by combining the growth rates of the above model to with the economic growth elasticities for poverty obtained by Ravallion (2012, AER). The underlying assumptions is that (mean) income per capita growth leads to poverty reduction, but that poverty itself has an effect on the elasticity of poverty to economic growth. This implies that econometric models for poverty change should include an interaction term of (lagged) poverty and income growth in addition to the standard income per capita growth variable. The results are shown in fig. 11. Only the lowest two income groups are shown here, because 'dollar a day' poverty is rare in higher income countries almost by definition.

Assuming the overall pattern stays the same, it seems that while accelerated educational expansion can be expected to make a sizeable contribution to overall growth, given the lag time, the education SDG might be 'too late' to contribute much to eliminating extreme poverty in terms of an absolute threshold, other than in the very poorest and currently least-educated countries. In those, however, the SDG scenario could 'save 10 years' in poverty elimination. That alone is nevertheless insufficient to 'eliminate' extreme poverty by 2030. Because the main constraint on education making a larger contribution to this goal appears to be timing, no analysis based on varying the *strength* of the relationship as a proxy for education quality was attempted.

To put these results into perspective, however, note that by 2030, the current absolute extreme poverty threshold will become increasingly irrelevant, and the question whether accelerated educational expansion can still contribute to raising large numbers of people to higher income levels becomes increasingly important. Indeed, the current target can be criticised as far too minimalist, as living on even two or three dollars a day amounts to dire poverty, all the more so by 2030 when average incomes are likely to have increased substantially. Estimating the effect on arbitrary poverty thresholds would, unfortunately, no longer be able to draw on the existing published research on the elasticities between growth and the current threshold.

In important side note concerns the question to what extent the above effects are due to the assumed spill-over to *post-secondary* attainment, rather than due to the achievement of the literal SDG target as such. While the spill-over effect should indeed be included in principle, its specific magnitude as included in the model is an assumptions that is to some extent arbitrary. It is prudent, therefore, to investigate the sensitivity of the results to this assumption. This is was done by performing additional estimations where the growth contribution of the post-secondary graduates attributable to the spill-over effect was reduced to that of upper-secondary graduates. The general conclusions stated above still hold for these more conservative estimates. In other words, most of the difference is in fact accounted for by secondary attainment.

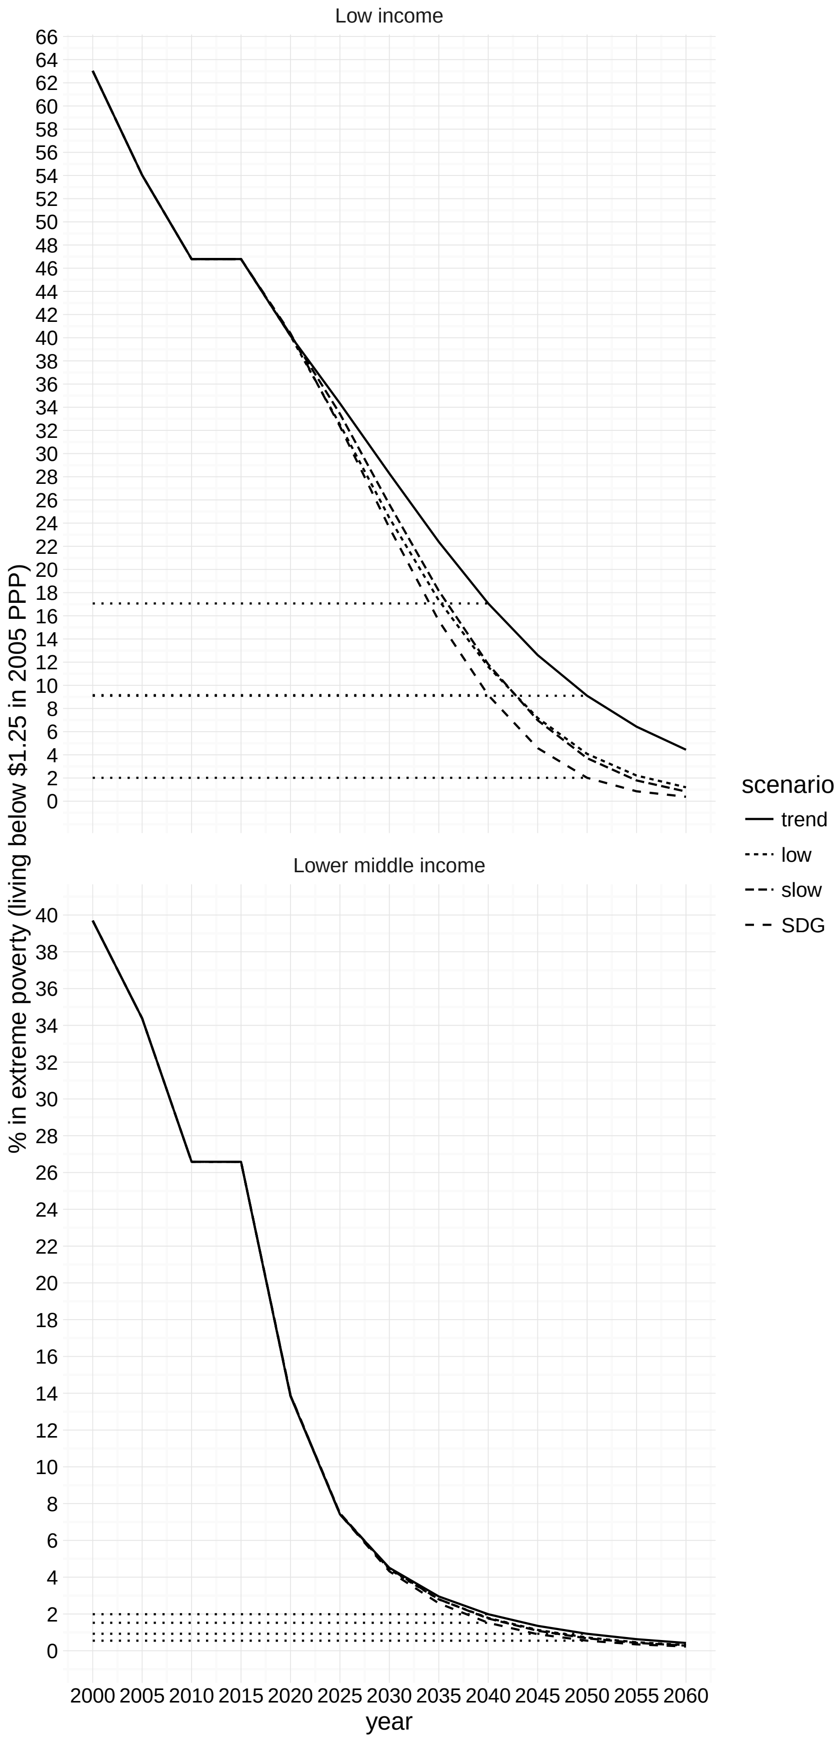

Figure 11: Projected absolute extreme poverty by scenario.

An explicit analysis of economic *inequality* is not feasible here. As the literature review demonstrated, the relationhsip between educational expansion and inequality is highly contingent on institutional arrangements, and does not lend itself to modeling an 'average effect'. As a proxy, it is possible, however, to changes in education itself from the perspective of the 'bottom 40 percent'. The rationale is that the bottom 40 percent in terms of income must be *at least* as well educated as the bottom 40 percent in terms of education. Improvements in the latter therefore provide a sense of the extent to which the 'rising educational tide' contributes to improving the relative position of the economic bottom 40 percent. For example, if in the future, only 20 percent of adults have only primary attainment or less, then at least half of the income bottom 40 percent must have completed lower secondary school at least.

To shed some light on this question, we have analysed the improvement in educational attainment in the 'top 60 percent' and 'bottom 40 percent' of the education distribution under the SDG scenario compared to the trend. Specifically, we model the probability for a random individual (at a given time, scenario, region) in the bottom/top group to have a strictly higher level of education than a random individual from the corresponding group in 2015. The reason this is not 50 percent in 2015 itself is because of ties. So if half have no education and half have primary, the probability in question is 25 percent. These results are summarised in fig. 13.

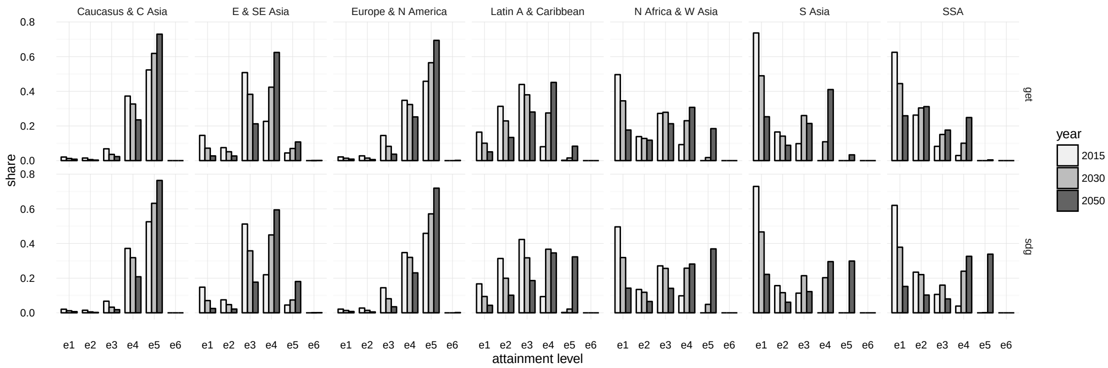

Figure 12: Attainment profile of the 'bottom 40%' of the education distribution.

We see that the bottom 40 percent are becoming much better educated, in the sense of a high probability of 'out-attaining' the baseline. For example, while a random member of the bottom 40 percent in 2020 in South Asia in terms of educational attainment has only a 25 percent chance of possessing strictly higher attainment than a member of the bottom 40 percent in 2015, this probability rises to almost 75 percent by 2050 under current trends.

In addition to this rise against the 2015 baseline, we see that the projected educational expansion can be interpreted as being moderately inequality-reducing, in that the bottom group improves more rapidly on this measure than the top 60 percent. However, the SDG target scenario is actually slightly *less* inequality-reducing than the baseline trend. This is to be expected, because the extra push compared to the baseline trend is stronger at a higher level, namely upper secondary.

Returning to the motivating question of *economic inequality*, given that in the above sense the bottom 40 percent are likely to see larger improvements than the top 60 percent in terms of education, we may be tentatively optimistic that educational growth may contribute to promoting above-average income group for the bottom 40 percent also. This is not clear from the above results alone, however, as the picture may well be different in terms of the gains in years of schooling, rather than the probability of reaching a higher attainment level.

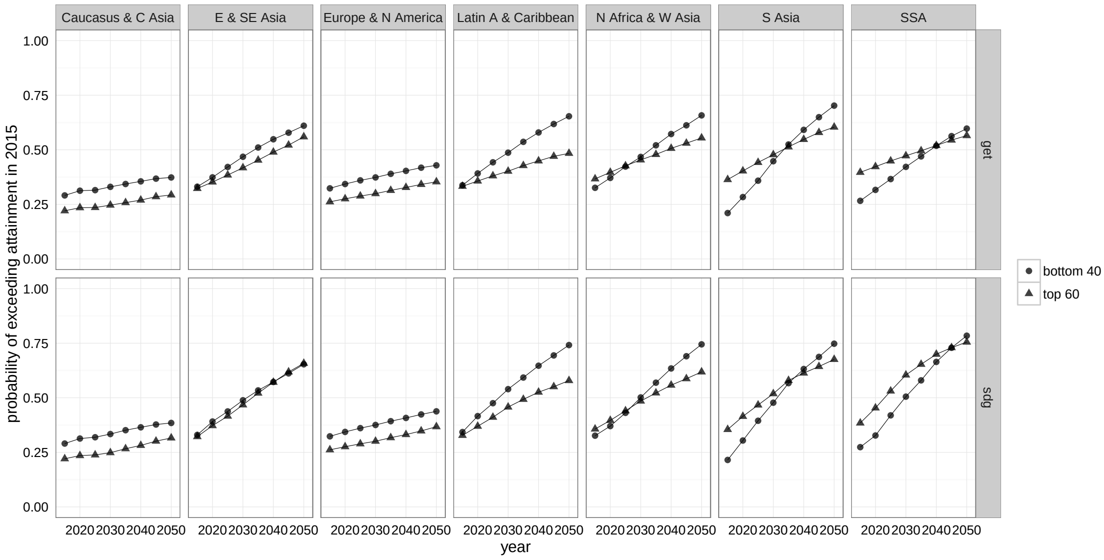

Figure 13: Probability of a random individual in the group (bottom 40 / top 60) & year strictly exceeding the attainment of a random individual from the same group in the baseline year. See main text for details.

3.4 Disaster deaths

As was demonstrated in the review in the first part of this report, there are strong linkages between education and climate change. On the one hand, this concerns both the ambivalent relationship between higher education and higher incomes and consumption, that potentially increase emissions overall, even if the more educated may be more supportive of institutional reforms and interventions aimed at climate change mitigation. On the other hand, there is evidence that higher education levels make a positive contribution to reducing vulnerability and increasing resilience to natural disasters generally.

While there is a notion that natural disasters strike 'indiscriminately', this contradicts the basic assumption underlying Disaster Risk Reduction, namely that information and preparedness can make a difference to survival and/or loss of assets and livelihoods. Unfortunately, information on educational characteristics of affected populations is often not collected in disaster-response situations. However, where it is available, the data suggest that the more educated do indeed tend to exhibit a greater awareness of risks, of appropriate preparation and response, and suffer smaller average losses in case of actual disaster. This justifies the expectation that vulnerability vis-a-vis climate-change induced disasters in the future may benefit similarly. This is of particular interest because 'combatting climate change and its impacts' is an important addition to the SDG agenda in comparison to previous development frameworks, and because this topic is among the least-well studied.

The potential positive contribution of education (and future educational expansion) to reducing vulnerability to climate disasters has been found both at the micro-level and in aggregate data. The following model and the study on which it is based fall in the latter category. Nevertheless, the fact that the findings fit well with corresponding micro-level evidence, and that the causal mechanisms outlined above are plausible, it may serve to at least illustrate the potential magnitude of the contribution education can make.

In this simulation, the predicted decadal numbers of deaths from natural catastrophes (storms, floods, droughts, landslides, and extreme temperature events) are obtained from the model presented in

Lutz et al. (2014). This model uses past information on the relationship between human capital and disaster deaths (as available from the Emergency Events Database (EM-DAT 2010)) controlling for other relevant covariates to project the impact of climate change in terms of future vulnerability according to the different Shared Socioeconomic Pathways (SSP). For the present purposes, the SDGs were simply added as another narrative of future adaptive capacity to the set of SSPs and several different scenarios were calculated for future environmental hazard. In the first scenario, the future number of disasters experienced within a country over a decade was assumed to remain what it was during the 2000-2010 period over the entire 21st century. This is then contrasted with a climate change scenarios assuming a 20 percent increase in the decadal number of disasters in comparison to the previous decade, respectively. The uncertainty ranges around the predictions indicate 95 percent confidence intervals.

This figure is obtained from the panel regression model with time fixed effects predicting the log of disaster deaths by climate-related disasters i.e. hydro-meteorological hazards such as floods, droughts, storms and extreme temperature. The estimated results are then transformed into predicted number of disaster deaths (measured as the logged number of deaths per million of population) according to different socioeconomic development pathways (SSPs) which are highly relevant for population dynamics and composition, and different climate change scenarios. For climate change scenarios, we make an assumption of an increase in hydro-meteorological extreme events of an average 10% and 20% per decade respectively. Although this is a rather simple assumption, even among the climate modelling community, there has not yet been a consensus on how climate change-induced extreme weather events would look like (Schleussner et al. 2015). However, the IPPC report and other scientific papers have confirmed that the current increase in the frequency, intensity and severity of extreme climate events observed today is due to anthropogenic climate change and these events are likely to rise in the future. The IPCC is particularly highly certain about the increase of longer and/or more intense heat waves, heavy precipitation events and increased incidence and/or magnitude of extreme high sea level (IPCC 2014).

In terms of the general pattern, the convexity in the expected number of disaster deaths over time is explained by the fact that some mitigating factors, including the diffusion of secondary education, become saturated while severe disasters are still becoming more common. With respect to the strength of the estimated linkage with education, this is large relative to the overall prediction range for a given climate scenario. For example, the median predicted outcome in 2040 under the SDG education scenario is close to the lower 2.5 percentile under the trend scenario. In absolute terms, the predicted number of annual disaster deaths is some 10–20 thousand lower under the SDG education scenario in the medium term, at constant disaster frequency and severity. Under a climate change scenario of more frequent disasters, the difference between the education scenarios widens to some 30–50 thousand annual disaster deaths.

From the perspective of the potential contribution of educational expansion to *off-setting* climate change with respect to its disaster toll, note that even under the more severe climate change scenario, baseline educational expansion may contribute to keeping the actual number of disaster deaths approximately constant well into the second half of the century. Educational expansion as envisaged by the SDGs would even keep the predicted value under the severe climate change scenario within the expected range under a no-change climate scenario until 2040 or so.

In this particular study, the effect of education is modeled as a coefficient on the share of the adult population with at least lower secondary education. Accordingly, the SDG, 'low', and 'slow' education scenarios essentially coincide in this case. Only the SDG scenario is shown, therefore.

In terms of regional variation, because these are absolute numbers of deaths, the global pattern is strongly dominated by the experience of Asia, which is not only home to some of the largest popula-

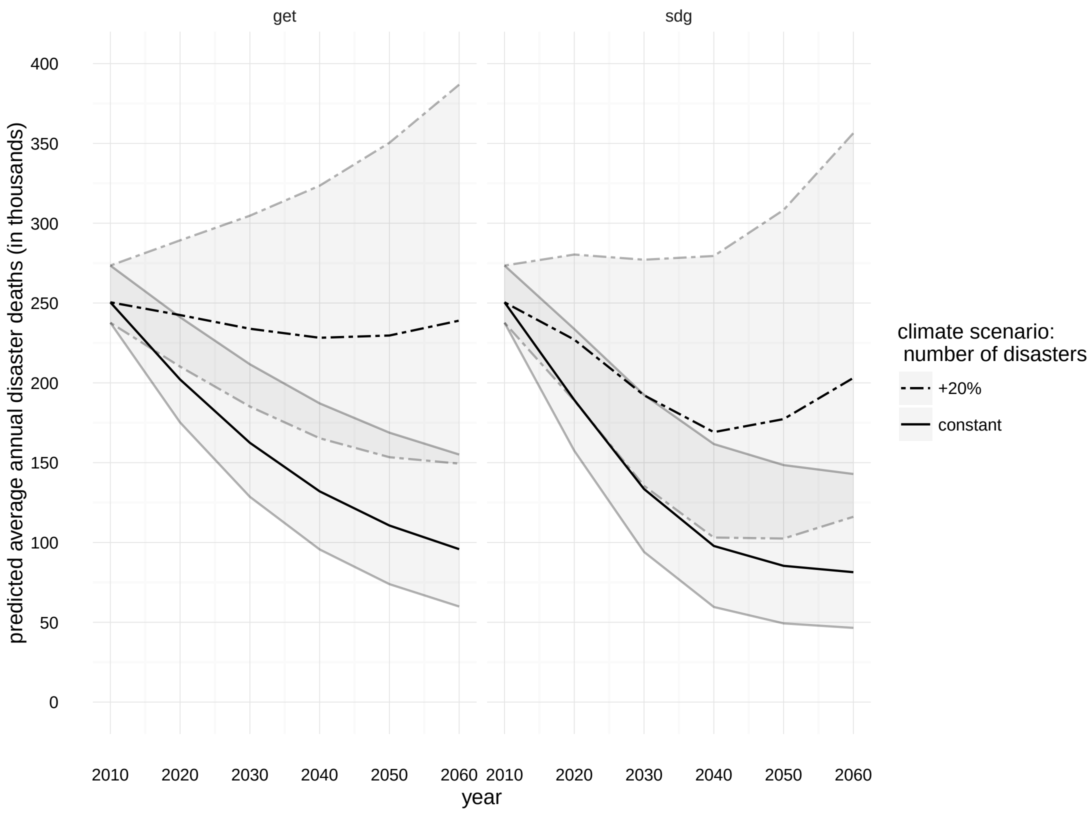

Figure 14: Projected annual deaths from hydro-meteorological disasters by scenario.

tions, especially coastal, but at the same time the locus of many disasters.

References

Cuaresma, J. C. (2015). Income projections for climate change research: A framework based on human capital dynamics. Global Environmental Change, page (online first). doi:10.1016/j.gloenvcha.2015.02.012

EM-DAT, CRED. 2010. "The OFDA/CRED International Disaster Database." Brussels: Université Catholique de Louvain, Center for Research on the Epidemiology of Disasters (CRED). http://www.em-dat.net.

Gakidou, E, Cowling, K., Lozano, R., and Murray C.J. 2010. "Increased Educational Attainment and Its Effect on Child Mortality in 175 Countries Between 1970 and 2009: a Systematic Analysis" 376 (9745): 959–74. doi:10.1016/s0140-6736(10)61257-3.

Lutz, Wolfgang, Raya Muttarak, and Erich Striessnig. 2014. "Universal Education Is Key to Enhanced Climate Adaptation." Science 346 (6213): 1061–62. doi:10.1126/science.1257975.

Ravallion, Martin. 2012. "Why Don't We See Poverty Convergence?." American Economic Review, 102(1):504-23.