



The allocation of World Bank Group resources to leave no one behind

Paddy Carter

Key messages

- IDA18 will give poorer countries more aid because the budget has increased, but the allocation is not significantly more pro-poor.
- A benchmark allocation created by DFID economists is much more pro-poor and could be approximated by adjusting the weights IDA's allocation rule.
- The allocation of WBG resources could be reformed to remove the need for an IDA/IBRD graduation threshold and to introduce more flexibility to tailor finance to country circumstance.
- Allocation decisions will inevitably be taken without knowledge of absorption constraints. Ambitious new global goals should bias allocation risks towards trying but failing, rather than failing to try.
- The ambition to leave no one behind suggests that the WBG should invest its resources in building its in-country capacity to relax absorption constraints in poorly-performing countries

Acknowledgements

The author thanks Stefan Emblad, Romilly Greenhill, Paul Isenman, Hubertus De Leede, Sarah Orzell and Andrew Rogerson for helpful discussions and comments. Funding from the Bill and Melinda Gates Foundation is gratefully acknowledged.

Table of contents

Acknowledgements	ii
Abbreviations	3
Executive summary	4
1 Introduction	8
1.1 New objectives, new allocation?	9
1.2 Performance-based allocation and reaching the poor	10
1.3 Do we know enough about absorption constraints?	11
2 IDA18	16
3 The DFID aid allocation benchmark	21
3.1 Two innovations	21
3.2 Ability to self-finance	22
3.3 Absorption constraints	23
3.4 Simplifying allocation rules	24
4 The way ahead: towards an ideal allocation of World Bank Group resources	27
4.1 Start with transfers	27
4.2 Absorption constraints, again	28
4.3 Using CPR to guide investment in capacity-building, not volumes	30
4.4 GNI per capita or poverty?	31
4.5 Computing the grant element and differentiated pricing	32
4.6 Two regimes	33
4.7 Setting the menu	34
4.8 Allocating the grant element	35
5 Conclusion	37
References	39
Appendix 1: Aid, growth and poverty in allocation models	42
Appendix 2: The allocation of transfers from the WBG	44
Figures	
Figure 1: The distribution of GNI across IDA members (including blend countries)	15
Figure 2: The distribution of CPR across IDA members	15
Figure 3: Allocations under IDA17 to IDA18 (annual \$ per capita)	20
Figure 4: Comparisons between PBA and IDA18 allocations and the Lea–Dercon benchmark	26
Figure 5: Financing menus	35
Tables	
Table 1: Mean allocations across income quartiles	18

Boxes

Box 1: The IDA Performance Based Allocation rule	9
Box 2: IDA member categories	11
Box 3: Understanding the weights in the PBA rule	14
Box 4: Doing Development Differently	31

Abbreviations

Abbreviation	Description
CPIA	Country Policy and Institutional Assessment
CPR	Country Performance Rating
DFID	Department for International Development (UK)
FCV	Fragility, conflict and violence
IBRD	International Bank for Reconstruction and Development
IDA	International Development Association (IDA)
IFC	International Finance Corporation
MIGA	Multilateral Investment Guarantee Agency
PBA	Performance Based Allocation system
SDG	Sustainable Development Goals
SDR	Special Drawing Rights
WBG	World Bank Group
WGI	World Governance Indicators

Executive summary

The World Bank Group will soon announce the results of its Forward Look exercise, to map out its medium to long-term strategy. It will also shortly finalise IDA¹18, the first replenishment of its concessional lending arm during the Sustainable Development Goals (SDG) era.

This paper is intended to inform the response of World Bank Group (WBG) shareholders and stakeholders to the resource allocation elements of IDA18 and the Forward Look. We ask whether the allocation of WBG resources is well aligned with its twin goals of eliminating extreme poverty and boosting shared prosperity, and the fundamental ambition of the SDGs: to leave no one behind.

The paper has three parts: an analysis of the proposed allocation under IDA18; an exposition of an allocation benchmark model created by DFID economists and a comparison of IDA18 against that benchmark; and a look ahead at how the allocation of WBG resources could be reformed.

Performance-based allocation and reaching the poor

The WBG has described the SDGs as a ‘trajectory shift’ in development dialogue, but while the allocation of WBG resources is responding to ambitious new global goals, the shift has been incremental. At the heart of the SDGs is the ambition to leave no one behind. The WBG’s twin goals of eliminating extreme poverty and promoting inclusive growth also entail reaching the hardest to reach.

There is inherent tension between the goal of leaving no one behind and the need to allocate funds efficiently to maximise impact. For its main concessional window, the IDA, the WBG uses a performance-based allocation system² (PBA), which adjusts allocations to account for how effectively the recipient is expected to make use of funds. There is nothing about the goal of leaving no one behind that means wasting money is a good idea, but at the same time pursuing a goal that will require reaching those resident in low-performing countries does not sit easily with an allocation rule that heavily penalises low-performing countries.

The basic idea behind optimal resource allocation is to equalise the marginal impact on WBG goals across countries. That impact depends on the constraints on each country’s ability to use WBG resources effectively. Little is known about these absorption constraints. The most influential research behind PBAs is based on the estimated cross-country relationship between aid and growth, which may be a poor guide to how the impact of WBG resources on its goals varies across countries. The WBG is not in the business of trying to stimulate aggregate GDP growth. Higher costs and a lower probability of success in low-performance countries may be offset by the fact that improvements to basic infrastructure and public services could have a much greater impact on the lives of people living in extreme poverty than similar investments in higher

¹ The part of the World Bank that helps the world’s poorest countries. IDA aims to reduce poverty by providing loans (called ‘credits’) and grants.

² The formula used to allocate IDA resources, by placing weights on each member country’s GNI per capita and CPR score.

performing countries. Any judgement about the optimality of WBG allocations requires taking a stand on the nature and extent of absorption constraints. This paper compares IDA18 allocations against a benchmark created by economists from the UK’s Department for International Development (DFID), which incorporates a particular view of absorption constraints.

IDA18

The most important fact about IDA18 is that the overall quantity of resources, under the management’s base scenario, will increase significantly. Core concessional IDA will increase from \$39.1 billion under IDA17 to \$52.3 billion under IDA18. Proposed replenishment contributions from shareholders are flat (in national currency terms) so over the course of IDA18 the pot (equity) will shrink as the grant-equivalent value of allocations exceed shareholder contributions. The main innovation in IDA18 is that IDA will leverage its equity and raise fund on capital markets, responding to demand for larger volumes on harder terms from IDA countries that have attained or are approaching middle-income country status. However, those leveraged funds will not be earmarked just for non-concessional lending (which is why equity will shrink). A number of changes will be introduced in IDA18 to increase resources allocated to countries affected by fragility, conflict and violence³ (FCV). These changes include eliminating the discount for 100% grants, increasing the minimum allocation, reducing the weight placed on performance in the PBA, and a new risk mitigation scheme for four countries, to complement an existing ‘turnaround’ regime (which will continue and is expected to see larger disbursement). However, the allocation of WBG resources will remain roughly flat with respect to recipient income. Table 1 shows per capita allocations over the three-year period, averaged over income quartiles, based on core concessional IDA allocated by PBA, ignoring various carve-outs and other changes.

Table 1: Mean allocations across country income quartiles

Income quartile	IDA17 (\$ per capita, mean)	IDA18 (\$ per capita, mean)
1	22.1	29.1
2	23.6	30.4
3	28.1	33.6
4	24.2	31.8

Source: author’s calculations

The DFID benchmark

DFID economists have created an aid allocation benchmark model (Lea and Dercon, 2016) as a tool for thinking about how well multilateral allocations are aligned with DFID’s objectives. The allocation they derive maximises the reduction in extreme poverty, as determined by the model they propose, so their benchmark is only aligned with one of the WBG’s twin goals. The main innovations that Lea and Dercon introduce

³ . The World Bank Group’s Fragile, Conflict and Violence Group annually releases the Harmonized List of Fragile Situations.

are that the model is forward-looking, so that a country where poverty is expected to persist for longer will receive more aid (all other factors being equal), and that the allocation accounts for a country's ability to self-finance its own efforts to eliminate poverty. Under the baseline allocation, the authors define that point as when the poverty gap has fallen beneath 1% of GDP. This method produces a sharp cut-off, so countries that have reached that point, such as India and Ghana, receive zero aid. Countries see their allocations curtailed quickly as that point draws near. The square of the WBG's Country Policy and Institutional Assessment index⁴ (CPIA) score is used to capture variation in absorption constraints across countries (so it has a lower weight than in the PBA). The end result is an allocation that is much more pro-poor than that proposed under IDA18.

The existing IDA PBA places weights on two country characteristics: mean GNI per capita, intended to capture the need for aid; and a Country Performance Rating⁵ (CPR) index, intended to capture absorption constraints. For IDA18 those weights are -0.125 and 3 respectively (they were -0.125 and 4 under IDA17). To an extent the innovations introduced by Lea and Dercon can be accounted for by appropriate adjustment to the weights in the PBA, because information on both the expected persistence of poverty and a country's ability to self-finance is contained, to a degree, in a country's GNI per capita and CPR: for a given poverty rate, countries with higher GNI tend to have more financial resources, and CPR should have some predictive power concerning the future path of poverty. This paper locates the weights on the PBA that come closest to achieving the Lea–Dercon allocation. They are -1.26 on GNI and 0.65 on CPR, implying a much more pro-poor allocation with a much smaller adjustment for performance.

The way ahead

What would the allocation of WBG resources look like if we were starting with a blank sheet of paper? The scarce resource that the WBG must allocate to maximise the impact on its twin goals is its equity, or grants and grant-equivalents. The same grant-equivalent transfer from donor to recipient can be conferred via a menu of different flows, starting with a pure grant and then progressing through various combinations of volumes and terms (interest rates and repayment schedules). This suggests a two-stage procedure: in a first stage, the WBG allocates the transfer it wishes to make to each country; then, in a second stage, in consultation with the WBG, each country chooses the combinations of price and volume, consistent with its allocated transfer, taking into account potential projects and respecting debt sustainability constraints. The attraction of this arrangement is that it offers greater flexibility and removes arbitrary discontinuities in the volumes and terms available to countries as they develop. If desired, this could be achieved by two organisations, blending pure grants from an organisation such as the IDA with zero grant element loans from an organisation such as the International Bank for Reconstruction and Development⁶ (IBRD). Further research is needed to compute the grant allocations under various scenarios and compare how available prices and volumes would compare to the status quo, but it may be possible to shift the allocation of grant resources further in a pro-poor direction while compensating other countries with larger volumes on harder terms. This may require the use of differentiated discount rates, to account for variation in risk across countries, when computing grant elements.

Other possible innovations could include using a poverty index that reflects more of the full distribution of incomes within countries, in place of GNI per capita in the PBA rule,

⁴ A World Bank diagnostic tool that is intended to capture the quality of a country's policies and institutional arrangements. There are 16 criteria grouped in four equally weighted clusters: Economic Management, Structural Policies, Policies for Social Inclusion and Equity, and Public Sector Management and Institutions. For each of the 16 criteria, countries are rated on a scale of 1 (low) to 6 (high).

⁵ Used in the IDA PBA. A linear combination of clusters A–C of the CPIA ratings (with a weight of 24%), Cluster D of the CPIA rating (with a weight of 68%) and portfolio performance rating (with a weight of 8%).

⁶ Provides loans and other assistance to middle-income and credit-worthy low-income countries.

without going as far as to base allocation on extreme poverty headcounts. However, continued use of GNI per capita could be justified on the basis that a country's ability to self-finance its own development also matters. Another possibility is to give the CPR score less of a role in determining the volume of resources that a country receives, and more of a role in determining what resources are spent on. WBG resources should be investing in relaxing absorption constraints and increasing the future impact of resources on its goals. Even if it is not always possible to change the underlying political and economic environment, if the WBG was to 'invest in investing' in country, it could be possible to improve performance in the given environment. This would entail reforming the current system for deciding allowable overheads in country.

Conclusion

The inescapable fact is that resource allocation decisions will always be taken in the face of considerable uncertainty about absorption constraints. A fruitful avenue for future research could be formal analysis of resource allocation using the theory of decision-making under uncertainty. These theories often select actions based on the regret a decision-maker would experience under alternative scenarios. In this era of ambitious global development goals, we may regret failing to try more than we would regret trying but failing. On that logic, the WBG should take a less conservative stance towards absorption constraints, invest more in capacity-building and introduce a more pro-poor pattern of resource allocation.

Box 1: key points

1. The proposed IDA18 replenishment will increase the quantity of resources allocated to poorer countries because the size of the pot will increase, but the allocation is not significantly more pro-poor in the sense of giving a larger share to poorer countries.
2. A benchmark allocation created by DFID economists, which incorporates absorption constraints, is much more pro-poor. The DFID benchmark is forward-looking and so allocates more resources to countries where poverty is expected to persist. IDA18 could approximate this benchmark allocation, while retaining its existing PBA rule, if it significantly increased its responsiveness to poverty and decreased the weight placed upon CPRs.
3. Looking ahead, the allocation of WBG resources could be reformed by first allocating a transfer to each country – the grant element of its financing – and then letting countries choose from a menu of prices and volumes that confer this transfer. The aim would be to maximise the impact relative to WBG goals, while respecting debt sustainability constraints and the Bank's own credit worthiness. Such a system would remove the need for an IDA/IBRD graduation threshold and would introduce more flexibility to tailor instruments to country circumstance.
4. It is impossible to know how best to allocate WBG resources without knowing how the ability to use aid effectively varies across and within countries. The existing evidence base, largely comprising estimated relationships from aid and growth regressions, offers little guidance to the impact of WBG resources on its twin goals. Allocation decisions are inevitably taken without knowledge of absorption constraints, so should be approached as a problem of decision-making under uncertainty. When outcomes are uncertain, it may be better to try and fail, rather than fail to try. Rather than seeing the problem as allocating resources given absorption constraints, the WBG should explore ways of allocating its resources to relax those constraints, tolerating higher overheads in poorly performing countries to reach those being left behind.

1 Introduction

There is one ambition that cuts across all the Sustainable Development Goals (SDGs): to leave no one behind. In 2013, two years before the world's governments signed up to the SDGs, the World Bank Group (WBG) set itself the twin goals of eliminating extreme poverty by 2030 and boosting shared prosperity. Since the adoption of the SDGs, the ambition to reach the hardest to reach and lift them out of poverty is now at the heart of both global development policy, as seen by WBG shareholders and clients, and the avowed purpose of the WBG itself.

The allocation of WBG resources has responded to this increased emphasis on tackling extreme poverty, but change has been incremental. In 2013, when the WBG adopted its twin goals, the allocation of resources by the Bank's International Development Association (IDA) in its seventeenth replenishment round, IDA17, shifted somewhat towards poorer and more vulnerable countries, and also extended transitional support to India, where large numbers continue to live in extreme poverty.⁷ Now the next replenishment round, IDA18, is almost finalised, and proposed changes to the allocation system will again see a modest shift towards poorer countries (although there will be a more dramatic shift towards countries afflicted by fragility and conflict, which could be described as more than 'incremental').

The purpose of this paper is to examine whether the allocation of WBG resources is well aligned with its twin goals and the ambition behind the SDGs, and to look ahead at how its allocation procedure could be reformed. The WBG is also currently undertaking a 'Forward Look' review of its operations and strategy and this paper may inform the response of shareholders and stakeholders to the resource allocation elements of that review.

The analysis proceeds in three parts: first, an examination of the proposed allocation under IDA18; second, an exposition of the Bank's Performance Based Allocation (PBA) rule, including a comparison with the allocation benchmark produced by the UK's Department for International Development (DFID) as part of its multilateral aid review; third, a depiction of an idealised WBG allocation system. The remainder of this introductory section discusses the fundamental problem of incorporating absorption constraints into allocation decisions.

⁷ IDA17 created an exceptional allocation regime for fragile and conflict-afflicted countries facing 'turnaround' situations, reduced the weight placed on the Country Performance Rating in its Performance-Based Allocation system from 5 to 4, to increase the poverty-orientation of its allocation, and raised the minimum base allocation from SDR3 million to SDR4 million to the benefit of small countries.

Box 2: The IDA Performance Based Allocation rule

Core concessional IDA resources are allocated according to a formula that places weights on a country's population and two characteristics designed to capture the country's need for aid and its ability to use it effectively: GNI per capita (calculated using the World Bank's Atlas method) and Country Performance Rating (CPR) (a weighted average of the World Bank's Country Policy and Institutional Assessment and ratings on past project performance). Because the weight on population is 1, it is possible to express the rule in per capita terms:

$$\text{resources per capita} = \kappa (\text{GNI}^\beta \text{CPR}^\gamma)$$

where κ is a constant capturing the size of the budget, β is the weight on GNI and γ is the weight on CPR. In IDA18, the weights are $\beta = -0.125$ and $\gamma = 3$. The final allocation includes various adjustments, such as a minimum base allocation and allocations from various carve-outs, such as the 'turnaround regime' and a 'scale-up facility'. Once the allocation is decided, the terms on which it is received is determined by the WBG's evaluation of credit risk. For more information, see Annexe 2 of WBG (2013) and WBG (2016).

1.1 New objectives, new allocation?

How the WBG should allocate its resources depends on what it is trying to achieve. Formally we can think of the optimal allocation as maximising some objective function – in this case the elimination of extreme poverty and the promotion of inclusive growth. The twin goals and the SDGs have given the WBG new objectives; new objectives imply a new allocation. How should such a rule respond to the goal of leaving nobody behind?

Carter (2014a) performs an experiment that could be interpreted as increasing the importance placed upon leaving no one behind.⁸ The results show that doing so makes little difference to the *optimal* weights in an allocation rule, but it should be emphasised that this does not tell us anything about how the *actual* weights in the PBA should change (because we do not know how far they are from optimal). The optimal weights do not much change because they will always give the poorest countries as much aid as they can effectively use. This happens when the global aid budget is sufficiently large, relative to the size of recipient economies, and when helping the poorest matters more than helping the better-off. In that case, the optimal allocation always pushes aid up against the absorption constraints in poorer countries, to the point where further aid does almost no good. So further increasing the importance of helping the poorest countries does not imply changing the weights to direct more aid towards them, because unless absorption constraints have been relaxed, doing so would be wasteful. Hence if the allocation of WBG resources was roughly right before the adoption of its new twin goals (under IDA16), meaning that the poorest countries were already receiving as much as they could usefully use, then the relatively minor changes to the allocation under IDA17 and IDA18

⁸ Carter (2014a) uses dynamic simulations in which aid is allocated across countries over time using a PBA (performance-based allocation) formula as used by the WBG's International Development Association (IDA), to show what determines the optimal choice of weights in such rules. In this experiment, countries are the unit of analysis and the model takes no account of the distribution of income within them, and so does not address the problem of leaving behind poorer citizens in richer countries. The experiment aims to simulate the impact of increasing the importance placed on reaching those left behind. It does this by increasing the curvature of the utility function so that the marginal utility of consumption diminishes more rapidly. In this experiment the objective of the donor is to maximise the sum of welfare across recipients; therefore, increasing the curvature of the utility function implies that helping the worst-off has a greater impact on the donor's objective, and less is gained by helping the better-off. See Carter (2014b) for a non-technical summary of this paper.

could be rationalised, despite the adoption of much more ambitious goals that put a premium on reaching the poorest.

Carter (2014a) also shows that when the quantity of aid being allocated increases, the optimal weight on poverty should become smaller. This is because, again, the optimal allocation always gives the poorest countries as much aid as they can use, so, starting from an optimal allocation, if the budget is increased without changing the weights, those countries would then receive too much aid. Hence the weight on poverty must decrease to keep aid to poorer countries at the right level. IDA18 has increased the budget being allocated without changing the weight on poverty, so that the poorest countries will receive more aid in absolute terms. That only makes sense if they were receiving too-little aid under IDA17.

Unfortunately, whilst it is desirable to know how the optimal weights in a PBA rule would respond to changed circumstances, such as placing more importance on reaching the poorest or increasing the quantity of aid being allocated, this knowledge is of limited practical use to the WBG for the simple reason that we do not know how far the current weights are from being optimal.⁹ It could well be that changes that have taken place between IDA17 and IDA18 would imply that the optimal weight on poverty should shrink, but in fact the weight on poverty in IDA18 should be much larger because it was far too small to begin with. Carter (2014a) shows that under plausible absorption constraints, the weight on poverty should be much larger than it currently is.

So what can we say about the allocation of WBG resources in this era of global goals built on the ambition to leave no one behind? One thing we can say for sure is that if the allocation of WBG resources is insufficiently pro-poor (the PBA weight on poverty is too small) then the cost of that error, in terms of WBG objectives, has risen. In the presence of uncertainty about absorption constraints, the new emphasis on reaching the poorest might imply a shift towards erring on the side of being overly generous, rather than overly conservative. Formal analysis of resource allocation problems under uncertainty is beyond the scope of this paper, but could a fruitful avenue for future research.

1.2 Performance-based allocation and reaching the poor

There is inherent tension between the goal of leaving no one behind and the need to allocate funds efficiently to maximise impact. The WBG remains committed to the idea of performance-based allocation, where funds are not just allocated in simple proportion to recipient need, but are adjusted to account for how effectively the recipient is expected to make use of funds it is allocated. There is nothing about the goal of leaving no one behind that means wasting money is a good idea, but at the same time pursuing a goal that will require reaching those resident in countries with, typically, low-quality economic environments does not sit easily with an allocation rule that penalises low-quality economic environments.

The basic idea behind optimal resource allocation is to equalise the marginal impact of allocated resources upon WBG goals, across countries. The PBA rule uses an index called the Country Performance Rating (CPR) to capture variation in the effectiveness of aid across countries.¹⁰ If Bank funds would do more to end extreme poverty and foster inclusive growth in one country than in another, the WBG would make more overall

⁹ Carter (2014a) shows that the optimal weights vary greatly according to what assumptions are made about absorption constraints, so it is not possible to say how far the actual weights in the PBA are from being optimal, without first knowing what the correct assumptions are about absorption constraints.

¹⁰ The CPR combines a weighted average of World Bank indices designed to measure various aspects of institutional quality with evaluations of performance on past WBG projects. Allocative efficiency is not the only justification for the use of the CPR in allocation formulae; there is also the possibility that countries will respond to the incentives inherent in the rule and take actions to improve their institutional quality. Collier and Dollar (2004) refer to a 'moral hazard constraint', which stipulates that aid allocation must not reward poor policy choices.

progress towards its goals by adjusting allocations until that disparity has been erased.¹¹ An implication of this is that if people living in extreme poverty are unfortunate enough to live in a country that the WBG allocation rule identifies as offering a low impact on poverty alleviation and inclusive growth, those people must be left behind for the sake of efficiency. This is the tension between performance-based allocation and the goal of leaving no one behind.

One interpretation of a PBA rule is as an admission that some people cannot be reached and that the SDGs and WBG goals will not be attained everywhere. This may not be publicly admitted by official development actors, but it is implicit in the logic of PBA. Were it not, aid could be entirely allocated in proportion to need, without consideration of the effectiveness of aid in some contexts.

If the reality is that WBG resources would not achieve much in some countries, nothing is gained by acting in denial of that reality. Lofty ambitions do not justify wasting resources, and the principle of PBA is sound. But do we know what those absorption constraints look like? It is possible that the assumption of excessively tight constraints in the PBA rule could result in allocations that are worse than if it is assumed that there are no constraints at all.

Box 3: IDA member categories

The set of countries eligible for access to IDA resources can be divided into three main categories: (a) core IDA countries (formally known as IDA-only non-gap countries), with income per capita below the operational graduation threshold (\$1,195 under IDA17) and who are unable to access credit on private sovereign debt markets; (b) gap countries (IDA-only), whose income level is above the operational threshold to receive allocations in the form of loans on harder 'blend' financing terms, but who are not considered creditworthy for IBRD financing; and (c) blend countries, who are creditworthy but have incomes below the threshold, hence have not graduated from IDA. World Bank support to blend countries takes the form of a blend of IDA credits, IDA hard-term credits and regular IBRD loans. Core IDA countries are subdivided into three groups: (i) countries with a high risk of debt distress, which receive 100% of their allocation in the form of grants; (ii) countries with a moderate risk of debt distress, which receive 50% in the form of grants; and (iii) countries that receive support in the form of loans on regular IDA terms. Under IDA17, grants currently receive a 20% volume reduction, which IDA18 proposes to remove.

Source: (IMF, 2015, WBG, 2016)

1.3 Do we know enough about absorption constraints?

It is natural to think that WBG resources will have more impact in some countries than in others, and also that at some point further increases in aid will achieve little. Many people also believe that too much aid can be positively harmful.¹²

What evidence we have on absorption constraints comes from estimated non-linear relationships in aid-and-growth cross-country regressions. This is typically done by interacting aid with country characteristics (such as an index of institutional quality) to capture how the effectiveness of aid varies across countries, and by adding a quadratic

¹¹ This is principle behind optimality: equalising the marginal product across countries (where marginal product is interpreted as the forward-looking marginal impact upon WBG goals).

¹² Deaton (2013) is the most recent high-profile articulation of this view.

term (aid squared) to the regression to capture how the effectiveness of aid varies within countries as the volume received rises.¹³

These estimated relationships are unsatisfactory for WBG's purposes for two main reasons. First, the WBG is no longer in the business of simply trying to stimulate GDP growth. The PBA rule needs to capture the ability of countries to translate WBG resources into, for example, improvements in social service delivery to marginalised communities or the extension of infrastructure to poor regions or urban neighbourhoods. A leap of faith is required to assume that the predicted short-run impact of aid on GDP growth, conditional on volume and country characteristics, is a good guide for predicting the impact of WBG resources on its goals. Although it is not clear that the choice of weights on the PBA rule is based on estimated empirical relationships in growth regressions, those estimates are an important part of the intellectual basis for the idea that WBG resources will have more impact in countries with better institutions.¹⁴

Second, and more importantly, even if these estimated relationships between aid and growth were statistically robust and stable (and they are not), finding a statistically significant estimated coefficient on aid squared in a growth regression would not constitute a test of a quadratic relationship against alternative non-linear relationships. There may be other non-linear relationships where aid exhibits diminishing returns but does not become harmful at higher levels that fit the data equally well.¹⁵

This may sound like an obtuse technical point, but it is an absolutely crucial question for the WBG. If aid exhibits diminishing returns but does not start to actively do harm at higher levels, that would imply that a much more pro-poor allocation of WBG resources is optimal. The reason is that parameterised formulae such as the PBA are actually quite blunt instruments and lack the flexibility to fully adjust allocations to individual countries. Any choice of weights in the rule will give some countries too much and others too little. If aid starts to do harm at higher levels then avoiding giving too much aid to any country becomes of paramount importance and so the optimal rule is calibrated to avoid over-aiding the very poorest countries, which quite possibly could result in not nearly enough being given to other, somewhat less poor, countries that could absorb more. On the other hand, if aid only becomes progressively less effective but never harmful, the optimal PBA rule would be much more pro-poor, because not much harm is done by over-aiding a handful of countries.¹⁶ Perhaps surprisingly, if aid is not thought to become harmful at higher levels, this does not have so much impact on the optimal weight on the indicator of absorptive capacity, CPR, but rather implies that the weight on GNI per capita in the PBA should change. That is because if aid does not become harmful, the whole distribution should shift towards the poor, whereas the weight on CPR merely adjusts aid across countries that have the same level of income.¹⁷

¹³ Benyon (2003) is a good introduction to and analysis of the work of Collier and Dollar (2001, 2002, 2004); Wood (2008) includes some history of the use of allocation rules and the relationship between optimal allocation solutions such as those of Collier and Dollar and less-flexible parameterised PBA rules.

¹⁴ More recently Bulman et al. (2015) have found that country characteristics only explain between 10% and 25% of the variation in the quality of project outcomes for the World Bank and Asian Development Bank. The rate of GDP growth and the CPIA index of institutional quality are positive correlates of project outcomes. The fact that most variation is explained by characteristics such as the track record of project managers suggests that the WBG could achieve good results in low CPIA countries by dedicating extra resources to ensuring implementation of best practice.

¹⁵ Based on the analysis of WBG projects, Presbitero (2016) finds that quality only suffers slightly during rapid accelerations of public investment.

¹⁶ Carter (2014a) shows that when absorption constraints are logistic, so that the marginal impact of aid tends to zero, the optimal weight upon GNI per capita is much larger than when absorption constraints are quadratic. One way to get around the inflexibility of PBA rules if aid is thought to do harm at higher levels is to place a large weight on GNI per capita, combined with a cap on the proportion of aid in GDP that a country can receive.

¹⁷ In more detail: if the shapes of alternative absorptive capacity relationships are similar in the range where aid has positive marginal returns but diverge at higher levels of aid (such as when comparing a logistic and a quadratic form) then the weight on CPR will be similar in each case. This is because most countries will be allocated aid in the range where aid still has positive marginal returns and the two functional forms are similar. What happens when countries are given too much aid does not greatly affect the optimal weight on CPR, so long as the weight on GNI is chosen so that most

It is hard to evaluate the allocation of WBG resources without taking a position on the type and degree of absorption constraints, or without explicitly incorporating uncertainty into the analysis. Section 2 will evaluate the allocation of WBG resources against the benchmark created by DFID economists.¹⁸ However, two points can still be made. First, while further research to produce better estimates of absorption constraints would be valuable, the WBG should also conduct research into formally analysing resource allocation as decision-making under uncertainty (Brock et al., 2007; Manski, 2013). Information on absorption constraints is always likely to remain very imprecise, and the WBG should establish how it responds to this uncertainty. Does it make sense to err on the side of allocating too much or too little to the poorest and lowest-capacity countries?

Second, WBG resources can be employed to eradicate poverty and promote inclusive growth given existing absorption constraints, but its resources can also be dedicated to building capacity in countries and relaxing those absorption constraints to enable greater progress towards its goals in the future. Seen in this light, indicators of absorptive capacity such as the CPR score could play more of a role in determining the share of WBG resources dedicated to building capacity in-country, and less of a role in determining the overall volume of WBG resources allocated. This idea is taken up again in Section 3, and the topic of absorptive capacity also comes up again, in the context of the DFID benchmark model.

countries are not getting too much aid. However, if no harm is done by giving countries too much aid, that makes it possible to shift the whole distribution towards poorer countries, which is achieved by changing the weight on GNI per capita. See Carter (2014a) for how optimal weights change.

¹⁸ Lea and Dercon (2015) assume a log relationship so that aid has diminishing returns but never becomes harmful.

Box 4: Understanding the weights in the PBA rule

The effect of a weight in the PBA allocation rule depends not only on the sign and magnitude of the weight but also the amount of variation in the data: if a weight is placed upon a statistic that does not change much from country to country, then a large weight will be required to produce an allocation that varies greatly across countries.

The bulk of IDA recipients have per capita incomes in the range \$500 to \$2,000 (Figure 1), so it is quite common for one country to have an income between two and three times higher than another. Placing a weight of -1 on GNI per capita in the PBA rule would mean, all else being equal, that a country that has twice the income of another will receive half as much aid, and hence its aid intensity (aid in proportion to the size of the recipient's economy, or aid/GNI) would be a quarter of that in the poorer country.¹⁹

The current weight in the PBA allocation rule of -0.125 implies that the poorer country would receive 9% more aid than the country with twice the income, and therefore the aid intensity is 2.18 times higher than in the richer country. A country with a GNI per capita of \$718, at the top of the first quartile, would receive roughly 18% more aid than a country with GNI per capita of \$2,720, at the top of the third quartile, implying that the aid intensity is 4.5 times as high (aid intensity is aid/GNI).

The range of CPR scores across IDA recipients is much tighter: half of the countries have scores between 2.9 and 3.4 (Figure 2). Placing a weight on 4 on CPR in the PBA rule, as in IDA17, means that a country with a CPR score of 2.9, at the top of the first quartile, would receive roughly half as much aid as a country with a CPR score of 3.4, at the top of the third quartile, all else being equal. Reducing that weight to 3, as proposed for IDA18, means that the 2.9 CPR country would receive 0.6 times as much aid as the 3.4 CPR country.

There is no particular reason to desire symmetry in the effect of income and CPR on allocations, but, for the sake of illustration, if we wanted the \$2,720 country to receive 0.6 times as much aid as the \$718 country, all else being equal, to mirror the impact of moving from the top of the first to the top of the third quintile of the CPR distribution, when the weight on CPR is 3, we would need to increase the absolute magnitude of the weight on GNI per capita from the current -0.125 to -0.39 .

Under the current choice of weights, CPR is responsible for more variation in aid per capita terms across IDA recipients than income; however, if aid/GNI is the relevant measure of aid allocation, then the picture changes.

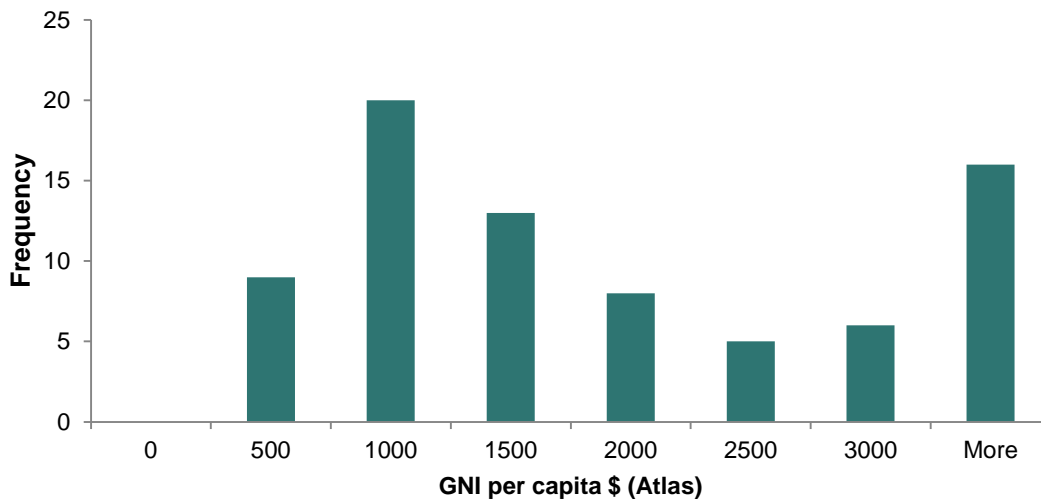
Based on the work of Collier and Dollar (2002) the 'performance based' aspect of the allocation rule is there to capture absorption constraints or variation in the ability to use aid effectively across countries. In theory the optimal weight on CPR in the allocation rule is there to adjust allocations so that the marginal impact of Bank funds is equalised across countries, for a given level of GNI per capita.

There is nothing in the PBA that captures absorption constraints within individual countries as the volume of aid increases. So, if there is some upper limit to the quantity of WBG resources that a country can use effectively, including the possibility that excessive dependency on outside donors destroys political accountability and has deleterious effects on national institutions, the only way to respect that limit is to calibrate the weight on GNI so that the poorest countries do not receive too much aid.

Source: WBG and author's calculations.


¹⁹ A weight of zero on GNI per capita – so that income makes no difference to allocations – would mean, all else being equal, that a country twice as poor would see aid/GNI twice as high. If we wanted to hold aid/GNI constant across countries, all else being equal, we would need a weight of $+1$. If aid/GNI is the right way to think about how much aid a country receives then perhaps rather than thinking of a weight of -0.125 as being only slightly different from zero, we should think of -0.125 as being -1.125 different from $+1$, the weight that would keep aid/GNI constant across countries.

Figure 1: The distribution of GNI across IDA members (including blend countries)



Source: author's calculations

Figure 2: The distribution of CPR across IDA members



Source: author's calculations

2 IDA18

This is the second replenishment since the WBG adopted its twin goals, but the first since the adoption of the SDGs. Setting out its proposals for IDA18, the Bank says the SDGs represent ‘a trajectory shift to the development dialogue’ to which the IDA ‘needs to respond with an innovative paradigm shift’ (WBG, 2016).

These changes have the following implications for resource volumes and the broad pattern of allocation:

- A large increase in total disbursement, from \$52 billion in IDA17 to \$75 billion in IDA18, under management’s base scenario. In percentage terms the grant element increases slightly, implying a large increase in the real economic transfer from WBG to IDA members.
- A shift in the allocation of concessional finance towards countries affected by fragility, conflict and violence (FCV). The absolute value of grants and the grant element of loans received by FCV countries will rise by 85%, and for IDA-only non-gap countries the increase is 50%. The share of the overall grant component allocated to ‘gap’ and ‘blend’ countries will fall by 20%, compensated for by larger gross flows.
- There is also a low scenario of \$65 billion and two high scenarios of \$80 billion. Under the first high scenario, higher volumes would be achieved by reducing the grant element, and under the second, higher shareholder contributions would be required.

The main specific changes to the allocation rules that will deliver this overall package are as follows:

- A reduction of the weight on CPR in the PBA rule from 4 to 3; the elimination of netting out of Multilateral Debt Relief; the elimination of the 20% allocation reduction when receiving 100% grants (all of which tend to disproportionately benefit FCV countries).
- Transitional support during IDA18 and an extension of the non-concessional scale-up facility introduced in IDA17. Both of these schemes will be financed by leveraged resources and not come at the expense of reductions elsewhere. A proposal to extend transitional support for graduating countries, phased out over three replenishments (nine years), was not adopted.
- A large proportional increase in the base allocation from SDR4 million to SDR15 million, to the benefit of small countries.²⁰ The concessional terms offered to small *island* states will be extended to all small states.
- The turnaround facility introduced in IDA 17 to support qualifying fragile states “on the cusp of major positive change” will be complemented by a new risk mitigation scheme with core IDA for four pre-identified countries (Guinea, Nepal, Niger and Tajikistan) that present increased risks of fragility, giving them up to a third more than their base PBA allocation.

²⁰ Base allocations are denominated to Special Drawing Rights, an international reserve asset, created by the IMF, based on a basket of four major currencies (soon to be five). One SDR is currently worth around \$1.4.

-
- A sub-window within the regional window for countries hosting refugees, and a new private sector window in collaboration with the International Finance Corporation (IFC) and the Multilateral Investment Guarantee Agency (MIGA).

The key idea behind the changes introduced in IDA18 is that countries that are moving towards middle-income country status and on towards graduation would prefer to receive higher volumes of financing. The big innovation in IDA18 – known as IDA plus – is to allow the IDA to borrow on capital markets to finance this less-concessional lending, making the IDA more like International Bank for Reconstruction and Development (IBRD) in that regard (although, with the exception of blend countries, the distinction remains that only the IBRD lends to creditworthy countries).

IDA18 will see the grant element of allocations (the real economic transfer from WBG to member countries) rise sharply. The grant element of IDA17 was 48% of \$52.1 billion, and under the IDA18 base scenario it is 49% of \$75 billion, yet proposed replenishment contributions into the fund by donors are flat (in national currency terms). The IDA18 Financing Framework paper is not yet public, and some non-concessional flows will be financed by repayments, but one implication is that IDA plus, the funds raised by leveraging the IDA balance sheet, are not being earmarked solely for non-concessional uses. Another implication is that the pot (equity) will shrink as grants disbursed exceed replenishments.²¹ This reveals that shareholders believe that demand for IDA resources will also shrink over time.

Some IDA members prefer to receive higher volumes on harder terms. One explanation for this is the current global low interest rate environment, meaning that the cost of debt service is less of a concern. Hence the interest rate discount available on IDA concessional terms, relative to its non-concessional terms, is perceived as less valuable by borrowers than the Bank's ability to underwrite large volumes of lending.²²

The WBG emphasises that IDA18 increases allocations to FCV countries in a variety of ways, including the increased base allocation, ending the grant discount, and the changed PBA formula, all of which tend to disproportionately benefit FCV countries. Although the Bank describes the change to the PBA formula, reducing the weight on CPR from 4 in IDA17 to 3 in IDA18, as 'increased poverty orientation' this description is questionable. Figure 3 shows annual per capita allocations generated using the PBA rule when reducing the weight from 4 to 3 based on data from 2014.²³ Figure 3 is not a comparison of IDA17 and IDA18 allocations – other than the reduced weight on CPR, all other changes introduced in IDA18 have been ignored and the sum being allocated has been held constant. The countries are ranked from highest to lowest GNI per capita. The allocations are not particularly pro-poor – the countries at the bottom do not tend to receive much more than those at the top. Reducing the weight on CPR does not greatly shift the allocation towards the poorer countries. Most of the countries with a relatively

²¹ Or perhaps IDA equity will merely stop growing. On the basis of publicly available information it is not possible to say.

²² This raises the issue of debt sustainability constraints, which is beyond the scope of the present paper. WBG (2016) contains analysis of this issue and concludes that IDA18 allocations are safely within those constraints.

²³ The figure was constructed as follows. Under the base scenario the core concessional IDA is \$52.3 billion, 92% of which is allocated by PBA. Annualised over three years that amounts to \$16 billion per year. There are 50 countries (excluding India, which only receives non-concessional funds, and including Eritrea and Somalia, which are currently inactive due to non-accrual). For some countries, GNI per capita data using the Atlas method was not available, so a prediction was made by regressing Atlas-method GNI on dollar GNI at market rates and the share of resource rents in GDP was generated and, where necessary, extrapolated to 2014; CPR scores were also not available for some countries so the CPIA scores used by Lea and Dercon (2016) were used. The PBA was then used to compute shares of the \$16 billion budget, each of which was then divided by population.

high CPR score, who see their allocations reduced under the change, are in the middle of the income distribution.²⁴

The WBG has considered increasing the magnitude of the weight on GNI per capita, but has rejected the idea. The reasoning is hard to understand: ‘The analysis concluded that reducing the CPR exponent provides a better balance between increasing financial support to FCS, maintaining the performance orientation in the system, and reducing volatility and limiting the impact on the affected countries’ (WBG, 2016: footnote 86). Compared with the alternative of putting more weight on income and keeping the CPR exponent unchanged, it is hard to see how reducing the CPR exponent maintains performance orientation. The phrase ‘limiting the impact on the affected countries’ reveals a reluctance to significantly reduce any allocations, relative to the status quo.

However, while the pattern of relative allocation across countries matters from a global efficiency perspective, and has not shifted under IDA18, from the perspective of each country what matters more is the absolute quantity of aid, not how their allocations compare to richer countries. Absolute quantities will change markedly under IDA18. Core concessional IDA will increase from \$39.1 billion under IDA17 to \$52.3 billion under the base scenario for IDA18, or \$50.7 billion after subtracting the new risk mitigation facility for Guinea, Nepal, Niger and Tajikistan, and \$1 billion of possible support for Syria. Running those numbers through the PBA formula shows there will be meaningful increases in the allocation per capita to poorer countries under IDA18, in absolute terms (see Table 1).²⁵

Table 2: Mean allocations across country income quartiles

Income quartile	IDA17 (\$ per capita, mean) CPR = 3	IDA17 (\$ per capita, mean) CPR = 4	IDA18 (\$ per capita, mean)
1	22.1	22.4	29.1
2	23.6	23.4	30.4
3	28.1	25.9	33.6
4	24.2	24.6	31.8

Source: author’s calculations

Despite the meaningful increase in the absolute quantity of resources allocated to poorer countries, the obvious critique of IDA18 is that rather than making minor changes to the pattern of allocation across countries, the WBG should have put a larger (more negative) weight on income and really shifted the distribution of resources towards poorer countries. As Table 1 shows, on average poorer countries still receive somewhat less per capita than richer countries. But, as argued above, it is impossible to take a stance on that issue without also taking a stance on absorption constraints. The next section compares IDA18


²⁴ In WBG (2016) the Bank calculates that the change does benefit poorer countries on average, based on fitting a line through the change in per capita allocations when plotted against GNI per capita. Based on Figure 3, one suspects the fit is not very good.

²⁵ Using 2014 data. The first column uses the IDA17 weight on CPR = 3, the second uses the IDA18 weight on CPR = 4. The GNI per capita quartile limits, including blend countries, are \$770, \$1,370, \$2,720 and \$8,080. The per capita dollar allocation figures are not annualised, they are for the three-year replenishment period.

against a benchmark created by DFID economists that takes a particular stance on absorption constraints.

From a political point of view, reallocating resources towards poorer countries may be more palatable if, as suggested in Section 3, the PBA rule is used to set the distribution of the grant element (real economic transfer) of WBG flows, leaving terms and volumes flexible, so that countries that see their grant element reduced can potentially be compensated by higher volumes. At present the PBA rule first allocates a gross sum, and the terms on which that volume is delivered (the grant element) is then decided by country characteristics (see Box 2: IDA member categories) so shifting the allocation towards poorer countries (holding the total budget constant) would mean reductions elsewhere. WBG (2016) stresses that nobody is worse off under IDA18, because any reallocation across countries is more than compensated for by increased volumes.

Figure 3: Allocations under IDA17 and IDA18 (annual \$ per capita)



Source: author's calculations

3 The DFID aid allocation benchmark

This section looks at IDA18 resource allocation in the light of an aid allocation benchmark produced by the UK's Department for International Development (DFID).

To support the department's recent Multilateral Aid Review, DFID economists Nick Lea and Stefan Dercon created an aid allocation benchmark model (Lea and Dercon, 2016).²⁶ The authors stress the model is not a prescription for donor allocations, merely a useful tool for thinking about how well multilateral allocations are aligned with DFID's objectives, and that it necessarily rests on a number of arbitrary simplifying assumptions.

The model is based on the idea of allocating aid to maximise the reduction of extreme poverty by stimulating economic growth, taking into account variation in the impact of aid on growth and of growth on poverty across countries, and also accounting for diminishing returns to aid within countries. This is the same approach as taken by Collier and Dollar (2002); it emerged from the empirical research of Burnside and Dollar (2000), who found that on average aid has a positive impact on growth in countries with good institutions, and represents the mainstream way of thinking about aid allocation.²⁷ The World Bank's PBA rule for IDA resources emerged from the same intellectual lineage (Wood, 2008).

In contrast to the WBG PBA rule, the Lea–Dercon benchmark is a fully optimal allocation, within the confines of the model. The authors define the objective in terms of an outcome, and write down a model that determines outcomes. The allocation that maximises that objective can then be derived. Every country receives exactly the right quantity of aid. Lea and Dercon define the objective of maximising the reduction of person poverty years: the number of people in extreme poverty multiplied by the number of years they are expected to remain poor. The PBA rule does not fine-tune allocations to individual countries to the same extent and is not directly derived from a formal model of how Bank resources affect outcomes.

The focus on extreme poverty is justified by the fact that DFID and other organisations share a focus on extreme poverty, and because a large number of the SDGs are also focused on dimensions of extreme poverty. Hence the Lea–Dercon model only reflects one of the WBG's twin goals. It does not reflect the goal of boosting the incomes of the poorest 40% of the population in any given country.

3.1 Two innovations

The Lea–Dercon model introduces two main innovations.²⁸ The first is that allocations are forward-looking, so that (when all else is equal) a country in which extreme poverty is

²⁶ Lea and Dercon (2016) is in the public domain but is not easy to find. A copy is available at <https://sites.google.com/site/paddycarter/>

²⁷ The Burnside and Dollar result is contested and some researchers find no role for institutional quality in explaining the relationship between aid and growth.

²⁸ The Lea–Dercon model also makes a number of assumptions that could be called into question. The importance of some of these assumptions is explored in robustness testing conducted by the authors. Any allocation model will

expected to persist for longer will receive more aid today. The second is that the time horizon for this forward-looking calculation ends when a country is judged able to self-finance the eradication of extreme poverty from its own resources.

Lea and Dercon judge a country able to self-finance poverty reduction once the sum of money required to fill the poverty gap falls beneath 1% of GDP. Lea and Dercon are not trying to estimate the costs of a realistic cash transfer programme to eliminate extreme poverty, hence they make no adjustments for administrative costs or leakage; they are merely trying to come up with a sensible yardstick for the point where a country can be said to have sufficient resources to tackle extreme poverty, if it chose to. So the objective is to maximise the reduction in person poverty years across countries, subject to the self-financing cut-off.²⁹ Appendix 1 discusses the mechanics of how aid is assumed to affect poverty in the Lea–Dercon model.

Critics have long argued that aid allocation decisions should be forward-looking (Wood, 2004). That approach is also in accordance with the mainstream economic way of thinking about dynamic welfare maximisation (Carter et al., 2015). In essence a dollar of aid today that lifts somebody out of poverty for a long time should be regarded as more effective than a dollar of aid that lifts somebody out of poverty for only a short while. The period of time over which aid lifts somebody out of extreme poverty is judged relative to what would have happened in the absence of aid. In the Lea–Dercon model, the impact of aid on poverty is permanent (it shifts the trend upwards), but in some countries people would have soon crossed the extreme poverty threshold in the absence of aid, hence aid will only lift them out of extreme poverty for a short while.

3.2 Ability to self-finance

The introduction of domestic resources (the ability to self-finance) to allocation decisions amounts to taking a position in the age-old debate about the allocation of aid to middle-income countries. As is well known, a large share of people living in extreme poverty live in middle-income countries (Kanbur and Sumner, 2012; Kharas and Rogerson, 2012), and some critics argue it is therefore a mistake to reduce the quantity of official development assistance (ODA) these countries receive. At the same time, many other observers argue that poorer countries currently receive too little aid in comparison to middle-income countries.³⁰

This debate raises some problems without technocratic solutions: there are moral and political judgements that must be made. One position is that, if all else is equal, countries that are unable to tackle extreme poverty should take priority over countries with domestic resources sufficient to tackle extreme poverty, but which choose not to. This is the position Lea and Dercon take, and it is also evident in the decisions of some bilateral donors to withdraw from some middle-income countries. Sending aid to relatively wealthy countries – even if they are home to many people living in extreme poverty – may also run counter to the wishes of the citizens in donor countries, which democratic governments and their

inevitable make simplifying assumptions that can be questioned, and the choices made by Lea and Dercon can be defended as reasonable on grounds of tractability.

²⁹ In the Lea–Dercon model, aid today delivers a permanent shift down the path of poverty; the authors argue that, empirically, improvements in headcount poverty tend to persist, and so reductions in the current poverty level translate into future reductions in the number of poor. The assumption that aid today affects the level of poverty in the distant future is probably not something to be taken literally – in the standard neoclassical growth model aid accelerates convergence towards the steady state but does not affect the level of that steady state, and the existing empirical literature is not capable of identifying the very long-run impacts of aid (existing long-run estimates rest on the assumed model and do not test alternative models). The presence of discounting in the model means that what happens in the distant future matters less than the near-term, so this is not a crucial assumption. The authors' robustness tests show that the assumed duration of the impact of aid upon poverty does affect the resulting allocation, but they study the extreme case of no persistent impact at all – there are also intermediate possibilities where the impact of aid is persistent but fades over time, such as when aid accelerates convergence.

³⁰ ODI (2015) analyses the allocation of ODA across countries in comparison to the estimated financing gap for social sector spending.

ministries of foreign affairs or aid agencies are supposed to represent. Another position is that it is unjust to punish the poor for the failings of their government, and many critics argue that middle-income countries need concessional finance to tackle pockets of extreme poverty within their borders.

There are also some important technocratic or pragmatic questions about aid to middle-income countries. If there is a moral hazard problem, wherein the presence of donors reduces a government's incentives to allocate its own resources to poverty alleviation, that would strengthen the case for withdrawing aid as countries get richer. If extreme poverty exists because the government does not care to address it, the extent to which it is possible for donors to reach the poor in middle-income countries without the assistance of a genuinely pro-poor government is also not clear, and circumventing uncooperative governments via NGOs and other channels may risk introducing other problems.

Without answers to these questions – which are beyond the scope of this paper – it is impossible to say how the allocation of WBG resources should respond to the availability of domestic resources, such as taxes and natural resource revenues, in partner countries.

The way in which Lea and Dercon introduce the ability to self-finance into their allocation model has stark consequences: many countries, such as Ghana and India, are already beyond that threshold and receive zero aid, while other currently poor countries who are expected to reduce poverty rapidly over coming years, such as Ethiopia, receive relatively low allocations.³¹

3.3 Absorption constraints

Lea and Dercon use the WBG's Country Policy and Institutional Assessment index (CPIA) squared to capture variation in the effectiveness of aid across countries. The distribution of CPIA across countries is somewhat different to that of CPR, as used by WBG (because CPR re-weights components of CPIA and includes a performance rating for past projects). Lea and Dercon's choice of weight (2) implies that, all else being equal, a country at the top of the third quartile of the CPIA distribution in their dataset, with a score of 3.65, would receive roughly 1.5 times as much aid as a country at the top of the first quartile, with a score of 3.01. They also assume diminishing returns within countries. Specifically, they assume that the change in GDP per capita is proportional to the log of aid per capita, so too much aid never becomes actively harmful.

As discussed above, that assumption would be absolutely crucial when thinking about PBA rules, which may give countries too much aid or, if calibrated to avoid doing so, will mean giving some poor countries too little. It matters less in the context of a fully optimal rule, like in the Lea–Dercon model. In that case, although allocations might be pushed towards the point where the marginal impact is zero, no country is ever given too much aid in the sense of pushing aid into negative returns territory.

But diminishing returns set in quite gradually in the Lea–Dercon model, and if poverty is high and persistent it makes sense to allocate lots of aid.³² So their baseline allocation gives large amounts of aid to some small economies (such as Malawi and Niger) far in excess of the 'turning point' of roughly 30% of GDP, where aid is often assumed to have no further impact on growth, based on the most widely cited estimates (Clemens et al.,

³¹ Ravallion (2010) proposed an alternative metric for capturing the 'capacity for redistribution' by estimating the marginal tax rate a country would need to impose on those who are not poor by rich-country standards to raise enough revenue to cover the domestic poverty gap. When he performed this exercise – a few years ago – he found that in the case of India not even a tax rate of 100% would be sufficient.

³² It is not possible to compare the severity of absorption constraints employed by Lea and Dercon with IDA18, because the PBA does not explicitly model absorption constraints.

2012). Lea and Dercon generate an alternative allocation, where aid is capped at 30% of GDP, but we use their main benchmark, where allocations are not capped.

Because Lea and Dercon are dealing with the global aid budget, their model can adjust the total quantity of aid in light of aggregate absorption constraints. Because the WBG is only allocating its own resources, which will be a fraction of the total amount of aid a country receives, it cannot treat absorption constraints in the same fashion. It might be possible to incorporate the predicted allocations of other donors, especially because other donors tend to follow the WBG lead (Knack et al., 2014), but in practice the WBG sees absorption constraints as pertaining to its own projects without considering the overall quantity of external finance a country receives.³³ However if a country is close to the limits of its aggregate absorptive capacity, that might be evident in its WBG portfolio performance rating.

3.4 Simplifying allocation rules

The WBG has a preference for simple and transparent formulae, such as the PBA, whereas the Lea–Dercon model is quite complex. It may be possible to achieve a reasonable approximation of the Lea–Dercon allocation by adjusting the weights in the PBA.

It may be possible to ‘shrink’ allocation models without much loss of information when the different inputs into the model are correlated with each other. For example, Lea and Dercon allocate more aid to countries where poverty is expected to be persistent. Expected poverty persistence may be correlated with having a lower CPIA score, which is already a part of the model. The Lea–Dercon model gives less aid to low-CPIA countries but more aid to persistent poverty countries – to some extent those two aspects of the model may cancel each other out. Hence the end result of giving more aid to countries where poverty is expected to be persistent may be achieved by using a model that does not explicitly include poverty forecasts but which places a smaller weight on CPIA. What matters is the strength of the correlation between expected poverty persistence and CPIA scores.

Lea and Dercon also base allocations on poverty, but they also want to account for the ability of countries to self-finance. Basing allocations on GNI per capita could be seen as achieving much the same result by less complicated means. Consider two countries with the same prevalence of poverty, but with differing levels of average GNI per capita: the richer of the two is closer to being able to self-finance anti-poverty measures. So using GNI per capita can be seen as a ‘reduced form’ means of taking a poverty orientation, adjusted for the ability to self-finance.

Of course it is not possible to precisely replicate a complex allocation, as used in the Lea–Dercon model, simply by adjusting the weights placed on just two variables in the PBA formula. Figure 4 shows PBA allocations per capita determined by adjusting the weights in the PBA formula to get as close as possible to the Lea–Dercon allocation, and for the sake of comparison also includes a PBA allocation using the IDA18 weights.³⁴ The weights that come closest to the Lea–Dercon allocation are –1.26 on GNI per capita and 0.65 on CPR (compared with –0.125 and 3 under IDA18). So to achieve a similar degree

³³ Finance from other sources enters into the equation in the sense that one of the CPIA components that underlie the CPR takes account of the sustainability of the country’s finances. However, that would imply countries with more non-WBG sources of finance will get a higher score and hence receive a higher allocation, all else being equal, whereas if the WBG was worried about countries breaching aggregate aid absorption constraints it would give less to countries that already receive large flows of aid from other sources.

³⁴ The set of countries shown are IDA members minus India. The set of countries used in the Lea–Dercon dataset also includes Angola, India, Swaziland and Syria. The PBA weights that come closest to the Lea–Dercon allocation are found by a numerical routine that minimised the Euclidian distance between the two allocation vectors, expressed as shares of the aid budget. Including a large country such as India, which receives zero under the Lea–Dercon allocation but which, because of its very large population, will always receive a reasonably large share using a PBA rule, has a significant impact on the weights. When India is included in the set of countries the weights that come closest to the Lea–Dercon allocation are –1.38 and 0.42.

of poverty orientation as the Lea–Dercon model, the PBA weights would require substantial revision.


The PBA with adjusted weights allocates substantially too little aid, in per capita terms, relative to the Lea–Dercon benchmark, to countries such as Liberia, Mali, Niger, Senegal and Somalia, where poverty is expected to be persistent. These countries are relatively small in population terms and do not account for much of the overall budget. The adjusted-PBA gives too much aid to two countries with large populations that account for a large share of the overall budget, but where Lea and Dercon expect poverty to fall rapidly: Bangladesh and Ethiopia. In the Lea–Dercon model, by far the largest share of the aid budget is allocated to the Democratic Republic of Congo, where the size of the population expected to be living in extreme poverty for a long time is very large; the adjusted PBA matches that allocation very well.

Compared with the IDA18 parameters, both the Lea–Dercon and the adjusted PBA allocations are far more pro-poor. The adjusted PBA allocation is much smoother – the only deviation from allocation in proportion to GNI is introduced by the relatively low weight on CPR – whereas the Lea–Dercon rule is more erratic, reflecting the variation in expected poverty persistence across countries. Evidently, GNI per capita and CPR scores do not correlate very closely with the predicted poverty rates of change used by Lea and Dercon.³⁵

Whether predicted rates of poverty persistence are reliable enough for the Lea–Dercon benchmark to represent a significant improvement on the adjusted-PBA allocation derived here is a question worthy of research. One thing is clear: the gap between Lea–Dercon and the PBA calibration used in IDA18 is large. This may be partially accounted for by the WBG also having the goal of promoting inclusive growth, as opposed to just eradicating extreme poverty. The PBA weight on CPR may also be chosen to provide countries with an incentive to improve their performance, a mechanism Lea and Dercon do not incorporate.

³⁵ It would be interesting to compare the predictive power of historic rates of poverty reduction and the level of CPR conditional on GNI, to see which does a better job of predicting future rates of poverty reduction.

Figure 4: Comparisons between PBA and IDA18 allocations and the Lea–Dercon benchmark



Source: author's calculations

4 The way ahead: an ideal allocation of World Bank Group resources

What would WBG allocation procedures look like, if we were to start with a blank sheet of paper? In a sense the ideal resource allocation process is bottom-up, starting from an assessment of the potential uses of those resources in each country, and equalising the marginal impact on WBG goals across countries on that basis. But in practice trying to make diverse potential uses of WBG funding comparable across countries would be extremely challenging, and it would be hard to make the comparison objective and transparent, free from the suspicion of political motivations or other sources of bias. The WBG has long seen the value in using a simple allocation formula based on publicly available data, and that is the approach we will pursue here.

4.1 Start with transfers

The starting point is to see the WBG as being in the business of transferring real resources to countries that face resource constraints, with the intention of accelerating development and achieving the WBG's twin goals.³⁶ Transfers (grants) are the scarce resource that the WBG must allocate to maximise the impact on its twin goals; this is the resource that it relies on shareholders to replenish every three years.³⁷ Although the WBG may face a liquidity constraint, implying that gross flows should also be regarded as a scarce resource, that constraint is less important.

This approach starts from the assumption that the WBG has a fixed budget of transfers (grant elements) to allocate, the magnitude of which is decided as part of the replenishment processes. Panizza (2015) has suggested that the WBG should allocate gross flows and the degree of concessionality (how much the country has to repay) should be chosen *ex post*, in response to events in borrower countries, implying that the WBG would not know in advance what demands on its equity there will be. That idea has merits, but we do not pursue it here and assume that the WBG wishes to define the magnitude of the grant element of its allocations *ex ante*.

In the case of a pure grant, the magnitude of the transfer is simply the face value of the grant, but when the WBG makes a concessional loan the grant element can be computed. The same grant-equivalent transfer from donor to recipient can be conferred via a menu of different flows, starting with a pure grant and then progressing through various combinations of volumes and terms (interest rates and repayment schedules). So, for

³⁶ Many observers stress that the real contribution WBG makes to development is – or should be – knowledge transfer (Clemens and Kremer, 2016; Ravallion, 2016). We ignore all such aspects of WBG activities here and focus entirely on the allocation of raw resources. We also ignore the IFC and MIGA.

³⁷ Within the WBG in the context of IDA the word 'transfer' refers to transfers from IBRD. Here we used the word to mean a transfer of resources from donor to recipient, or the grant-equivalent value of a financial flow. Another way of framing the scarce resource to be allocated is equity. Equity is depleted when financial flows do not generate sufficient returns.

example, a \$10 million transfer could take the form of a \$10 million grant, a \$30 million loan on highly concessional terms, or a \$100 million loan on close-to-market terms. In theory, any of these can be constructed by blending a \$10 million grant with a ‘hard’ loan, which confers zero transfer (and hence resembles a loan in an ideal competitive capital market, where real economic profits are zero in the absence of monopoly power or other sources of rent).

This observation suggests a two-stage procedure: first, the WBG decides the allocation of its available budget for transfers across countries; and second, the terms of that transfer are determined. In the second stage not every item on the menu would be available: the choice of volumes and terms would be made to maximise the impact of the transfer to that country on the WBG’s twin goals, which implicitly means staying within debt sustainability constraints (because having a public finance crisis does not help a country eliminate extreme poverty or deliver shared prosperity). A very poor country that merits a large transfer would not be able to take that transfer in the form of a huge loan on hard terms. The first stage allocation would be the product of a formulaic allocation rule like the PBA, while the second stage would be the product of collaboration between the WBG and partner countries, informed by the Bank’s assessment of the potential uses of those resources and so closer to the ideal bottom-up process, working within debt sustainability constraints. For simplicity’s sake we may think of a debt sustainability constraint for each country being defined in terms of the gap between existing outstanding debts and a maximum debt level (based on both capacity and willingness to pay), defined in present value terms.

Although this two-stage procedure could be implemented by a single organisation, for the sake of organisational continuity the grant element could be provided by the IDA, from a fund that is periodically replenished by donors, and a zero grant element loan provided by the IBRD, which is able to raise funds on capital markets. In theory there would be no need to divide countries into IBRD or IDA (or blend). In the case of a country that is currently classified as IDA-only and receives its allocation in the form of 50% grant and 50% standard IDA credits, that same combination could be delivered blending a 100% grant from the IDA with a zero grant element loan from the IBRD. Because the overall package is deemed sustainable – based both on capacity and willingness to pay – the risk exposure of the IBRD should still be acceptable. That is to say, preserving the IBRD credit rating would be a constraint on the menu available to the country.

Such an arrangement would require changes in the IBRD’s Articles of Agreement, which currently forbid lending to IDA-only countries (without a counter-guarantee provided by a creditworthy entity).

4.2 Absorption constraints, again

This two-stage procedure raises the question of when to apply the concept of absorption constraints: during the first stage or the second? At present, gross IDA allocations are adjusted for performance, but in the proposed two-stage procedure, gross flows would be chosen in the second stage.

The line of argument that the scarce resource to be allocated is the grant element of flows suggests that the concept of absorption constraints should be applied at the first stage of the procedure. But can the marginal impact of the grant element on WBG’s goals be captured without knowing what terms and volumes decisions are going to be taken in the second stage?³⁸

³⁸ This problem also arises in the current IDA allocation procedure, but reversed. Currently the gross IDA allocation is first decided by the PBA allocation formula and the degree of concessionality then decided by country characteristics. If the weight on CPR is there to capture absorption constraints it takes no account of the degree of concessionality that will be later chosen, and treats grants and non-concessional flows alike. In this section we suggest reversing the order and first

In the context of a loan, the concept of absorption constraints must be distinguished from that of debt sustainability. A country with fiscal space, perhaps with a fast-growing economy, can afford to repay a loan even if that loan is utterly wasted. A country with no fiscal space can afford to repay a loan if it is used to finance expenditure that yields sufficiently high and timely returns. We may think of financial absorption constraints as capturing the ability of a country to identify and execute positive net return investments, and monetise those returns through tax revenues. But some positive-return investments may have little impact upon the Bank's twin goals, so the concept of absorption constraints in a purely financial sense should not guide the allocation of WBG non-concessional flows. Rather, expected financial returns should inform which choices from the menu are available in the second stage.

Although a simplification, it might be helpful to think of absorption constraints as consisting of two elements: the potential impact of projects on the WBG's twin goals, and the efficiency with which those projects would be executed or the probability of successfully delivering those projects. In a low-capacity country, the provision of basic infrastructure or public service delivery could have the potential to make a big difference to the lives of people living in extreme poverty, but the probability of achieving that potential might be low. The expected returns on WBG's goals is the product of the expected cost, the potential returns and probabilities of achieving them. It is not obvious what distribution of expected returns across different levels of development looks like. There should be no presumption that expected returns in terms of the WBG's goals are higher in wealthier countries with higher CPR scores, because these countries may have higher existing capital stocks (human and physical) and fewer people living in extreme poverty. As argued above, the estimated relationship between aid and GDP growth across countries may not be a good indicator of the pattern of marginal returns on WBG goals. Aggregate economic growth is not the only channel via which WBG resources can reduce extreme poverty and promote inclusive growth, nor is it necessarily a good proxy for other channels.

Suppose that PBA is used for the grant element allocation in the first stage, based on two variables: GNI per capita and the CPR score.³⁹ If absorption constraints are to be accounted for entirely in the first stage, that amounts to saying that either countries with identical incomes and CPR scores will make identical terms and volumes choices at the second stage or, if they make different choices, the marginal impacts of those choices on WBG goals will be identical. Strongly stated, that may not seem very likely. But it is no more unlikely than the marginal impact of IDA resources on WBG goals being identical for two countries with the same levels of GNI and CPR under the current procedure.

Unless a truly bottom-up procedure is used, under any allocation procedure it will be impossible to identify absorption constraints in any detail *ex ante*: the ultimate impact of WBG resources on its twin goals will be determined by choices made after the allocation is determined. It is hard to see how to incorporate information about what volumes and terms will be chosen in the second stage when determining the size of the transfer in the first stage, while preserving simplicity and transparency. So it seems sensible to use a PBA rule in the first stage to account for absorption constraints (or variation across countries in the ability to transform WBG resources into progress towards WBG goals) without worrying about how choices made in the second stage interact with absorption

determining the size of the transfer and then choosing the gross volume associated with that transfer, and treating all transfers alike, whether conferred by grant or loan.

³⁹ GNI per capita is usually regarded as capturing need and CPR score as capturing absorption constraints, but actually GNI should also be seen as having a role in absorption constraints. A PBA rule based on GNI and CPR does not explicitly account for variation in potential marginal returns to investment. In a simple model of economic growth, with an aggregate production function, this omission could be justified because GNI per capita maps onto the capital-to-labour ratio, which in turn pins down the marginal return to capital. The idea that GNI per capita captures both the need for aid, in a welfare sense, and the scope for productive investment could be justified as a simplifying assumption, assuming that the marginal impact of projects on WBG goals are higher in poor economies; it is hard to see how it could be improved upon whilst preserving tractability and transparency.

constraints. In the second stage, decisions on volumes and terms are made to maximise the expected impact on the WBG's goals, in view of potential projects and within debt sustainability constraints.

4.3 Using CPR to guide investment in capacity-building, not volumes

The current approach to resource allocation – as used in both the WBG's PBA and the Lea–Dercon model – takes absorption constraints as a given. But in theory WBG resources can also be spent to relax absorption constraints and raise the impact of resources in the future.

Collier and Dollar (2004) argue that the scope for donors to either improve the policy and institutional environment or fruitfully work around governments is limited. More recent analysis has found a small positive association between aid and the quality of political institutions (Jones and Tarp, 2016). But rather than hope to improve the political and economic environment within a country, the WBG may also invest in the capacity to produce better results given the existing environment. Collier refers to 'investing in investing' (Collier, 2011). There are models of economic growth that incorporate this idea; for example, the model by Rappaport (2006) includes two complementary forms of capital, one of which bears a high installation cost and can be thought of as a capacity 'bottleneck'. In such a model, low capacity might not imply that a lower allocation is optimal, but rather that a higher proportion is spent on raising capacity. This possibility is not really reflected in the empirical analysis that formed the basis of the PBA rule.

It is often observed that WBG project ratings are not much worse in low-CPR countries, which the PBA rule heavily penalises.⁴⁰ But a potentially more relevant factor is how the cost of achieving results, in terms of staff costs and other overheads, varies across countries. Carter (2014a) models absorption constraints as determining the gap between gross aid and net *effective* aid (the gap is the amount of wasted aid). The gap between gross and effective flows, or waste, could be interpreted as including overheads. Carter demonstrates that when donors have the objective of maximising welfare across countries, the optimal allocation will tolerate high levels of waste in particularly poor countries because it is worth spending a lot of money to get a little money through to people in great need.

At present, IDA shareholders wish to say that all of their contributions go directly to poor countries, and that WBG overheads are funded through transfers from the IBRD, with the level of overhead allowable in-country being no higher in low-CPR countries. In the absence of such political considerations, it would make sense to use CPR scores to guide the level of 'investment in investing' funded by WBG resources, so that in low-capacity countries the WBG builds up over time the capacity to deliver projects that have a large impact on its twin goals. A formal analysis of aid allocation in a model that allows for a dynamic impact on absorption constraints through investment in local capacity might suggest a PBA with a lower weight on CPR, and more capacity investment in low-CPR countries. This could be a fruitful area for future research.

⁴⁰ See Geli et al. (2014) for example. World Bank project ratings are not necessarily a reliable guide to how absorptive capacity varies across countries, because they record the extent to which project objectives were met, and the chosen objectives may be less ambitious in more difficult circumstance. For resource allocation purposes an absolute performance measure is needed, not a relative one.

Box 4: Doing Development Differently

This paper approaches resource allocation from a very high level, and does not get into the details of how IDA resources are spent. Current thinking in development effectiveness emphasises the importance of tackling locally defined problems, being experimental and gathering information on performance and adapting projects in the light of that information as they proceed (in contrast to implementing projects that are designed in detail upfront and then rolled-out, with deviations from the plan regarded as something to be avoided). This approach goes by various names: Doing Development Differently, adaptive programming or problem-driven iterative adaptation (Andrews et al., 2013; Bain, 2016; Booth et al., 2016).

An important element of this approach is the introduction of greater flexibility into the resource allocation process, lengthening or rolling budget cycles, allowing task team leaders to either kill failing projects or scale up successful projects without being constrained by disbursement pressures. It can also mean tolerating higher overheads for more intensive oversight of implementation. Introducing flexibility to country resource allocations should be an important element of rethinking the WBG resource allocation process, but it is beyond the scope of this paper.

4.4 GNI per capita or poverty?

The WBG has the goal of promoting growth in the lowest 40% of the income distribution in every country it operates in. This serves the sensible purpose of expanding the focus of the WBG beyond narrowly defined extreme poverty; the idea that people should count for nothing in development policy once their income surpasses the extreme poverty threshold, intended to capture absolute destitution, is self-evidently absurd (Pritchett and Kenny, 2013). The goal of ending extreme poverty is something to be taken seriously, not literally.⁴¹

But having the goal of promoting inclusive growth in every country does not imply that inclusive growth in every country is equally important to the WBG. Consider two countries, one poor, where the lowest 40% of the income distribution have incomes below \$5 per day, and another richer, where the lowest 40% of the income distribution have incomes below \$10 per day; in light of the SDGs it is natural to think that promoting inclusive growth in the poorer country is relatively more important to the WBG.

One potential approach to increasing the poverty orientation of Bank allocations without going as far as to replace average GNI per capita with the extreme poverty headcount, as done by Lea and Dercon, is to use a measure that gives a more complete picture of the distribution of income within a country, but which still gives more weight to people the poorer they are. Such measures can be constructed using household survey data and a Foster–Greer–Thorbecke (FGT) poverty index (Foster et al., 2010). The Bank could define a high international poverty line, such as the US poverty line of around \$30 per day, and then base allocations on the number of individuals in each country weighted by the gap between their incomes and this poverty line. When the weight is 1, poor people count for more in proportion to how far their incomes are beneath this line, but a weight greater than 1 can be used so that the weight of individuals in the allocation process increases exponentially with the depth of their poverty. Thus the population living in extreme poverty would count for much more than those with higher incomes within each country, but, unlike in the Lea–Dercon model, it would not be the only thing that counts. Rather than chose a weight on GNI per capita, the WBG’s PBA could choose the ‘poverty aversion’ weight in an FGT poverty index. The advantage of doing this is that it preserves

⁴¹ This phrase is borrowed from Richard Manning, former OECD DAC chair.

the idea that those living in extreme poverty matter the most, but would also place some weight on the bottom 40% of most middle-income countries, who live on less than \$30 per day, in line with the WBG's second goal, and would not create a strong cut-off around a low poverty threshold.

This approach would be more demanding in terms of data. The WBG has committed to funding more regular representative household surveys in developing countries, but at present the data required to construct FGT indices might not exist for some countries. Another drawback is that unless those surveys are conducted very frequently, contemporary individual country poverty distributions, used to allocate IDA resources, would need to be estimated otherwise allocations would be based on historical data.

Alternatively, the continued use of GNI per capita could be justified on the basis that it represents a rough compromise between basing allocations on the prevalence of poverty with some implicit adjustment made for a country's ability to self-finance its own development spending.

4.5 Computing the grant element and differentiated pricing

A technical challenge in the approach outlined here is the calculation of the grant element, or the transfer from WBG to partner country. The WBG considers IBRD lending to be non-concessional, although IBRD loans clearly confer a benefit to member countries by offering credit at rates that are often substantially below sovereign lending rates on private markets.⁴² Capitalising the IBRD, including committing to contingent capital, also has an opportunity cost for shareholders (see Clemens and Kremer, 2016), which exceeds the returns on lending (which can be proxied by the transfer from net income that the IBRD makes to the IDA each year).⁴³ So while the WBG does not see things this way, from an economic perspective the IBRD confers a transfer on member countries.

In the current low interest rate environment, the interest rates on IBRD loans would imply a positive grant element when using the WBG and International Monetary Fund (IMF) unified discount rate for low-income countries of 5%.⁴⁴ The decision, taken in 2013, to unify discount rates for both debt sustainability and grant element calculations made sense for a relatively homogenous group of countries, but it may be harder to justify if the same rate is extended into current IBRD member countries, some of which are able to borrow on private capital markets at reasonably low rates. Countries that are able to borrow at rates close to 5% may feel that the discount rate is too high, and so to confer a given grant element the WBG ought to make larger volumes available at lower rates, in comparison to countries that face higher rates on capital markets.⁴⁵ Some may even be able to borrow below 5%, in which case a loan that the WBG would consider to confer a grant element would confer no benefit to them.

⁴² IBRD loans are provided on terms that are not considered concessional in respect of IDA's non-concessional borrowing policy

⁴³ Bulow and Rogoff (2005) suggest development bank accounting returns are illusory because they are created by defensive lending and grants, but the fact that IBRD is able to remit profits to IDA suggests it is making some genuine returns, even if they may be insufficient to represent positive returns once accounting for the cost of capital and risk bearing.

⁴⁴ An alternative method for computing grant elements is against opportunity costs. When lending to countries with access to capital markets, opportunity costs could be captured by sovereign bond prices. But while investment opportunities available to private investors are a genuine forgone opportunity when making a private investment, acting like a profit-motivated private investor is not a genuine forgone opportunity for the public sector. In the context of public sector financial flows, the grant element should be calculated with reference to an estimated marginal cost of public funds, perhaps with an adjustment for bearing risk. Professor Miles Kimble gives an accessible discussion of these issues in the blog post '[Discounting government projects](#)'. Quite apart from conceptual considerations, achieving consistency across countries where sovereign bond prices are available and countries where a discount rate must be assumed (to avoid either penalising or rewarding some countries) would be difficult.

⁴⁵ There are important differences in tenor and flexibility too. A single interest rate is referred to here for the sake of simplicity.

This problem is related to the issue of differentiated pricing with the IBRD. At present the IBRD operates like a cooperative and extends lending on the same terms to all its members. Bank shareholders and member countries hold some opposing views about the current distribution of IBRD lending and terms. Some take the view that better-off countries receive too large a share of IBRD funds. Some take the view that less risky countries ought to be able to borrow at lower rates, others that better-off countries, which could afford to pay higher interest rates, ought to do so.

The approach being proposed here would introduce differentiated pricing in the sense that some countries would receive larger grant-element allocations in the first stage and hence the frontier that defines the combinations of price and volume available to them would shift out. Holding the discount rate constant across countries, the idea that countries capable of paying higher interest rates should do so would be implemented if those countries receive a smaller grant-element allocation in the first stage (which they would tend to do, if ability to pay is correlated with the variables in the PBA). But countries that are able to borrow at low rates on capital markets, and which are allocated a small grant element in the first stage, may feel that the resulting frontier, which defines available combinations of price and volume, does not offer them enough. From the WBG's perspective, the combination of a unified discount rate of 5% and the allocation of small grant elements to upper-middle-income countries could also be seen as forgoing the opportunity to lend larger volumes to those countries, and to generate income for redistribution to poorer countries. The OECD Development Assistance Committee (DAC) currently uses a set of differentiated discount rates to compute the grant element of loans, but these are all strictly above the WBG–IMF unified rate of 5% (the DAC adds a premium for riskier countries). There could be a case for the WBG adopting differentiated discount rates that fall beneath 5% for countries that are able to borrow at similar (or lower) rates on capital markets.

It might appear that having an IBRD-like organisation make loans with zero grant element would mean forgoing the opportunity to generate income that could then be transferred to the IDA. This is not the case. A zero-return loan implies returns sufficient to compensate shareholders for the opportunity cost of capital and risk bearing. Shareholders could still remit those returns to the IDA, as they do at present. However, there may be scope to make positive-return (negative grant element) loans to higher income countries. That idea could be incorporated into this two-stage process by, for example, expressing GNI in the PBA rule as a deviation from some threshold, above which countries are allocated negative grants.

4.6 Two regimes

A more straightforward alternative to introducing differentiated discount rates could be to preserve two regimes: one that offers blends of grants and loans; and, for richer countries able to borrow at low rates on capital markets, an arrangement much like the IBRD. Because many IBRD countries currently borrow at more than 5% on capital markets, and because a first-stage PBA allocation would still allocate them a meaningful grant element, even when a relatively large weight on poverty is used (as will be illustrated below), many of them would be able to access large volumes on harder terms under the unified blend regime. So there would be an opportunity to shift the GNI per capita 'graduation' threshold, after which countries graduate into IBRD-only, much further up into the range of upper-middle-income countries. Doing so could ease concerns about countries losing access to concessional finance too soon, while also hardening terms as countries grow and the PBA allocates an ever-smaller grant equivalent to them.

4.7 Setting the menu

Given a grant-equivalent allocation, in the second stage of the two-stage process countries would choose from a menu of prices and volumes, in consultation with the WBG, taking into account available potential projects and debt sustainability constraints. The chosen combination of price and volume could be seen as an aggregate, consisting of a number of grants and loans for different projects on different terms.

Graphically, combinations of price and volume consistent with a given grant-equivalence describe a financial flow frontier. A debt sustainability constraint, interpreted as the maximum present value of additional debt a country is considered able to bear, can also be plotted. Figure 5 shows two examples.⁴⁶ In both cases the financial flow frontier is consistent with a \$10 million grant-equivalence. Anything above and to the left of that line represents a flow with a grant-equivalence of less than \$10 million; anything below and to the right, above \$10 million. There are two illustrative debt sustainability constraints: (a) labelled ‘low debt capacity’, with a present value of \$10 million, perhaps representing a low-income IDA country; and (b) labelled ‘high debt capacity’, with a present value of \$50 million, perhaps representing a middle-income IBRD country. Anything below and to the left of the debt constraint line is within the constraint, anything above and right exceeds it. So the available menu to each country is the twice-hatched area at the left of the figure, below the debt constraint and above the flow frontier. Assuming the country will want to take its maximum entitlement, it will choose some point on the segment of the flow frontier beneath the debt constraint.

These graphs are for illustrative purposes only. Toward the right of each graph, it can be seen that small increases in interest rate are consistent with large increases in volume, so middle-income countries allocated small grant equivalents could borrow large sums, debt constraints permitting.

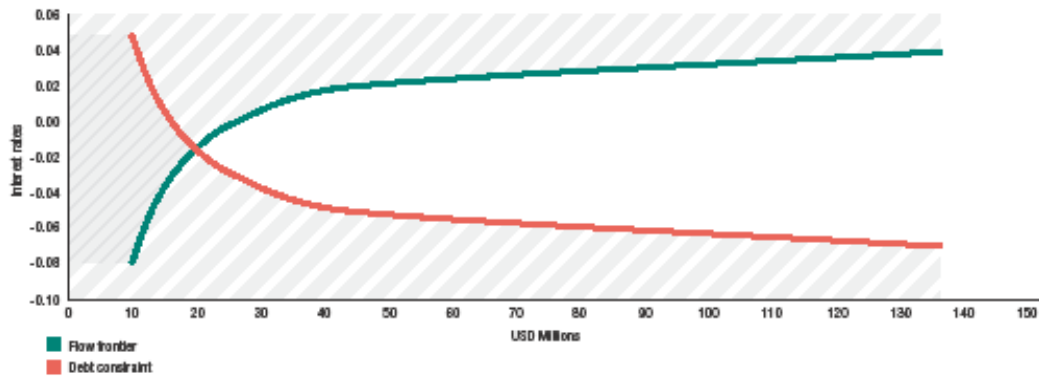
In reality the problem is rather more complicated than computing additional borrowing capacity quantities for countries in isolation, because the WBG credit rating also depends on its regional exposure.⁴⁷ But it should be no more complicated than under current arrangements.

⁴⁶ The financial flow frontier is calculated using a very simple 10-year maturity loan with interest paid once annually and a lump sum repayment of principal in year 10, with no grace period or fees. In reality WBG loans are more complicated than this, but a net present value can still be calculated and thus represented using a standardised loan format in price/volume space. For both the flow frontier and debt constraint, the WBG–IMF unified discount rate of 5% is used.

⁴⁷ The desire to maintain an AAA credit rating has been seen as a significant constraint on the Bank, but recently Standard and Poor’s revised its methodology, potentially giving the IBRD more headroom. See [‘Could multilateral banks be lending an extra \\$1 trillion?’](#).

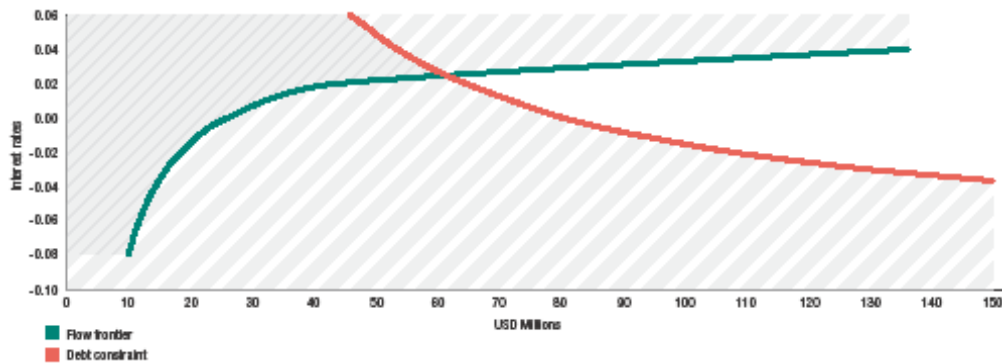
Figure 5: Financing menus

A: Low debt capacity



Source: author's calculations

B: High debt capacity



Source: author's calculations

4.8 Allocating the grant element

If a PBA-type rule were to be used to allocate the grant element of WBG flows across all countries, it would need to be recalibrated. Carter (2014a) shows how the optimal weights in a PBA rule change markedly when, for example, the ratio of the total sum being disbursed to the total incomes of recipients changes, or when the distribution of characteristics across recipients changes. In particular, adding IBRD member countries would introduce countries with higher scores on whatever metric, resembling CPR, is used to capture absorption constraints, which would again have implications for the optimal weights.

For the sake of illustration, the figures in Appendix 2 show allocations of grant resources across all IDA and IBRD countries using a PBA rule. Because CPR data are not available for IBRD countries, the Government Effectiveness rating of the World Governance Indicators (WGI) was used in its place, after transformation to resemble the distribution of the CPR data.⁴⁸ Two allocations are shown, one using the proposed weights in IDA18 and one using the weights that came closest to achieving the Lea–Dercon allocation. Again, the allocation ignores the minimum SDR floor and other adjustments.

⁴⁸ A positive constant was added to the WGI because negative numbers cannot be used in a PBA (unless one wishes to allocate negative aid flows), and then the resulting score was raised to the power of 0.75 to bring the variance across countries into line with the variance of CPR across countries. This was done because holding a weight constant but changing the variance of a metric would increase the importance of that metric in overall allocations.

To compute the budget to be allocated (the total transfer from WBG to its partner countries) the grant element of concessional IDA18 under management's base scenario of 49% was applied to the \$75 billion total and annualised over three years to give a figure of \$12.25 billion. The grant element inherent in IBRD lending is harder to estimate – the bank does not treat IBRD as explicitly concessional in the same way as it does IDA, although there is clearly a benefit conferred on member countries (and most IBRD loans have a large grant element when using the current 5% discount rate). To produce a rough estimate the figure of \$1.75 billion was taken from Clemens and Kremer (2016), which corresponds to their estimate of the opportunity cost to IBRD shareholders, including the risk of callable capital; but, to account for returns and avoid double-counting, \$0.65 billion was subtracted from this, representing the 2015 transfer made from the IBRD to the IDA.⁴⁹ This produces an estimated annual grant element of WBG operations, excluding IFC and MIGA, of \$13.35 billion.

This exercise shows that, if applied to the complete set of WBG member countries, the current weights in the PBA rule would produce a pro-rich allocation. The correlation between the allocation per capita and income per capita is 0.38. The largest allocations are to countries that score highly on indicators of government effectiveness, such as Chile and Malaysia. This reflects the combination of a large positive weight on WGI, a small negative weight on income, and a strong positive correlation between WGI and income (of 0.57). In contrast, the Lea–Dercon weights, with a large negative weight on income and a small positive weight on WGI, produce a strongly pro-poor allocation: the correlation between the allocation per capita and income per capita is –0.57. The distribution is much more strongly skewed using the Lea–Dercon weights: the maximum allocation (Benin) is \$22 per capita, whereas under the IDA18 weights the maximum is just \$3.7 (Mauritius).

However, even under the strongly pro-poor Lea–Dercon weights, richer countries still receive a small positive allocation. Chile, for example, receives \$0.19 per capita. That is not too far from the allocation for poor, low WGI countries when using IDA18 weights, where the minimum allocation is \$0.33 per capita (Somalia). In absolute dollar terms, the allocation of \$0.19 per capita for Chile using the Lea–Dercon weights translates to a grant equivalence of \$3.4 million, which would be consistent with a \$68 million loan with a 5% grant element, for example. Using the IDA18 weights, the per capita allocation for Chile of \$3.6 translates in absolute dollars to \$63 million, consistent with a \$315 million loan with 20% grant element, for example. According to the latest data, the IBRD original principal amount for Chile is \$330 million.⁵⁰

An interesting topic for future research would be to plot the financing flow frontiers for all WBG countries (combinations of price and volume that are consistent with the allocated grant equivalent), looking at both the different potential choices for weights in a PBA rule and the possibility of differentiated discount rates, and comparing these with status quo flows from the WBG.

⁴⁹ Source: '[The roles and resources of IBRD and IDA](#)'.

⁵⁰ Taken from World Bank Group Finances, '[Country summary: Chile](#)'.

5 Conclusion

The WBG has called the adoption of the SDGs a ‘trajectory shift’ in development policy, but only incremental changes to the allocation of its concessional resources are proposed in IDA18 (although a set of amendments targeted a countries afflicted by fragility, conflict and violence mean those countries will see rather more significant increases). In particular, the reduction of the weight placed on performance in the Bank’s PBA rule does not deliver a noticeably more pro-poor allocation, in relative terms. But IDA18 will see a meaningful increase in the size of the overall budgetary envelope, meaning that many poor countries will see significant increases in the absolute quantity of resources they receive, even if the relative pattern across countries has not changed markedly.

Any allocation can be rationalised on the basis that it gives poorer countries as much aid as they are able to use effectively, and without information on such ‘absorption constraints’ it is hard to say more about the efficiency of WBG resource allocation. This paper has compared proposed allocations under IDA18 against a benchmark created by DFID economists Nick Lea and Stefan Dercon, to inform DFID’s Multilateral Aid Review (forthcoming), which embodies a set of assumptions about absorption constraints motivated by mainstream economic research. A substantial revision to the weights in the WBG’s PBA would be needed to come close to that benchmark, which is much more pro-poor than the existing PBA. To an extent the allocation resulting from Lea and Dercon’s approach, arising from innovations such as being forward-looking and accounting for a country’s ability to self-finance anti-poverty efforts, can be achieved by appropriate adjustments to the PBA.

This paper has then sketched a potential integrated regime for allocating resources across the whole WBG, which would allow countries greater flexibility in choosing combinations of pricing and volume that suit their circumstances. The basic idea is to proceed from an initial allocation of the transfer (grant equivalent) conferred by the WBG and then permit any combination of pricing and volume consistent with that transfer and debt sustainability constraints. This could be implemented by blending 100% grants from the IDA with zero grant element loans from the IBRD. The attraction of this idea is that it avoids sharp discontinuities in the terms available to member countries as they develop, and builds in gradual diminution of the degree of concessionality offered by the WBG to richer countries. These things are much easier to sketch in principle than to implement in practice and there would be many technical and political considerations to think through before such a scheme could be adopted. Even if one holds that the current allocation of WBG resources is far from optimal, the political reality is that reforms will probably only be acceptable if they do not create too many losers, relative to the status quo. But many WBG member countries may prefer larger volumes on harder terms, so there may be scope to compensate some middle-income countries if the allocation of grant resources is shifted away from them.

This paper has also suggested that less weight could be placed on performance indicators when determining volumes and instead these indicators could be used to inform investments in WBG country capacity to deliver projects, with the intention of relaxing absorption constraints over time. Furthermore, empirical estimates of the relationship between aid and aggregate economic growth, which form part of the basis for performance-based allocation, may be a poor guide to the pattern of marginal returns to

WBG resources across countries, in terms of impact on its twin goals of eradicating extreme poverty and promoting inclusive growth. Higher costs and the lower probability of success in low-income countries with lower quality economic and institutional environments may be offset by the fact that improvements to basic infrastructure and public services could have a much greater impact on the lives of people living in extreme poverty than similar investments in better-off countries.

But the fact remains that we will probably never have good information regarding absorption constraints, and that resource allocation decisions will always be taken in the face of considerable uncertainty. Therefore, a fruitful avenue for future research should be formal analysis of decision-making under uncertainty. One well-established approach to decision-making under uncertainty is known as ‘min-max regret’, which entails trying to minimise the regret a policy-maker would feel in the event of a worst-case scenario. The application of that approach in this context is not straightforward, because if the worst-case scenario is that aid is positively harmful, the implication could be to shut down the WBG. That possibility aside, in this era of ambitious global development goals, one might have thought we would regret failing to try more than we would regret trying but failing. If so, the WBG should take a less conservative stance towards absorption constraints, invest more in capacity-building and introduce a more pro-poor pattern of resource allocation.

References

- Andrews, M., Pritchett, L., and Woolcock, M. (2013). Escaping capability traps through problem driven iterative adaptation (PDIA). *World Development*, 51, 234-244.
- Bain, Kathy (2016) The Implications for large aid organizations of Doing Development Differently – updating the plumbing to fit the architecture.
- Beynon, Jonathan (2003) Poverty Efficient Aid Allocations – Collier/Dollar Revisited. Economic and Statistics Analysis Unit Working Paper 2. ODI
- Booth, D., Harris, D., and Wild, L. (2016). From political economy analysis to doing development differently: A learning experience. London: ODI.
- Brock, W. A., Durlauf, S. N., and West, K. D. (2007). Model uncertainty and policy evaluation: Some theory and empirics. *Journal of Econometrics*, 136(2), 629-664.
- Bulman, D. J., Kolkma, W., and Kraay, A. (2015). Good countries or good projects? Comparing Macro and Micro Correlates of World Bank and Asian Development Bank Project Performance (April 21, 2015). World Bank Policy Research Working Paper, (7245).
- Bulow, J., and Rogoff, K. (2005). Grants versus loans for development banks. *The American economic review*, 95(2), 393-397.
- Burnside, C and Dollar, D (2000) “Aid, policies, and growth”, *American Economic Review*, vol 90(4).
- Carter, P. (2014). Aid allocation rules. *European Economic Review*, 71, 132-151.
- Carter, P. (2014b). Aid allocation rules: a non-technical summary. ODI.
- Carter, P., Postel-Vinay, F., and Temple, J. (2015). Dynamic aid allocation. *Journal of International Economics*, 95(2), 291-304.
- Clemens, M. A., and Kremer, M. (2016). The New Role for the World Bank. *The Journal of Economic Perspectives*, 30(1), 53-76.
- Clemens, M. A., Radelet, S., Bhavnani, R. R., and Bazzi, S. (2012). Counting chickens when they hatch: Timing and the effects of aid on growth. *The Economic Journal*, 122(561), 590-617.
- Collier, P (2011) IGC Growth Forum. Lecture by Professor Paul Collier [online] Available at: <http://www.theigc.org/wp-content/uploads/2014/08/collier-lecture.pdf> [Accessed 2 October 2016]
- Collier, P and Dollar, D (2002) “Aid allocation and poverty reduction”, *European Economic Review*, vol 45(1).

Collier, P., and Dollar, D. (2001). Can the world cut poverty in half? How policy reform and effective aid can meet international development goals. *World development*, 29(11), 1787-1802.

Collier, Paul, and David Dollar. "Development Effectiveness: What have we learnt?." *The economic journal* 114, no. 496 (2004): F244-F271.

Deaton, A. (2013). *The great escape: health, wealth, and the origins of inequality*. Princeton University Press.

Foster, J., Greer, J., & Thorbecke, E. (2010). The Foster–Greer–Thorbecke (FGT) poverty measures: 25 years later. *The Journal of Economic Inequality*, 8(4), 491-524.

Geli, Patricia and Kraay, Aart and Nobakht, Hoveida. (2014). *Predicting World Bank Project Outcome Ratings*. Policy Research Working Paper No. 7001. World Bank Group

IMF (2015) Eligibility to use the Fund's facilities for concessional financing, 2015 [Online] Available at: <http://www.imf.org/external/np/pp/eng/2015/062415.pdf> [Accessed 2 October 2016]

Jones, S., and Tarp, F. (2016). Does foreign aid harm political institutions? *Journal of Development Economics*, 118, 266-281.

Kanbur, R., and Sumner, A. (2012). Poor countries or poor people? Development assistance and the new geography of global poverty. *Journal of International Development*, 24(6), 686-695.

Kharas, H., and Rogerson, A. (2012). *Horizon 2025: Creative destruction in the aid industry*. London: Overseas Development Institute

Knack, S., Xu, L. C., and Zou, B. (2014). Interactions among donors' aid allocations: evidence from an exogenous World Bank income threshold. *World Bank Policy Research Working Paper*, (7039).

Lea, N and Dercon, S (2016) *Benchmarking aid allocation: Is the global and multilateral aid allocation poverty-focused?* Department for International Development, Chief Economist's Office.

Manski, C. F. (2013). *Public policy in an uncertain world: analysis and decisions*. Harvard University Press.

Panizza, U. (2015). *Debt Sustainability in Low-Income Countries-The Grants versus Loans Debate in a World without Crystal Balls* (No. P120).

Presbitero, A. F. (2016). Too much and too fast? Public investment scaling-up and absorptive capacity. *Journal of Development Economics*, 120, 17-31.

Pritchett, L., & Kenny, C. (2013). Promoting millennium development ideals: The risks of defining development down. *Center for Global Development Working Paper*, (338).

Rappaport, J. (2006). A bottleneck capital model of development. *Journal of Monetary Economics*, 53(8), 2113-2129.

Ravallion, M. (2010). Do poorer countries have less capacity for redistribution?. *Journal of Globalization and Development*, 1(2).

Ravallion, M. (2016). The World Bank: Why it is still needed and why it still disappoints. *The Journal of Economic Perspectives*, 30(1), 77-94.

Temple, J., & Van de Sijpe, N. (2015). Foreign Aid and Domestic Absorption (No. 15/658). Department of Economics, University of Bristol, UK.

WBG (2013) Updated IDA17 financing framework and key financial variables. IDA Resource Mobilization Department (DFiRM) World Bank Group.

WBG (2016) The Demand for IDA18 Resources and the Strategy for their Effective Use. IDA Resource Mobilization Department (DFiRM) World Bank Group.

WBG (2016) IDA Lending Terms [Online] Available at:
<http://ida.worldbank.org/sites/default/files/pdfs/ida-lending-terms-july-2016.pdf> [Accessed 2 October 2016]

Wood, A. (2008). Looking ahead optimally in allocating aid. *World Development*, 36(7), 1135-1151.

Appendix 1: Aid, growth and poverty in allocation models

Allocation models of the sort used by Lea and Dercon (2016), based on the work of Collier and Dollar (2002) and related research, conceive of economic growth as the only way to reduce poverty. These models offer no scope for aid to raise welfare by funding consumption. In Lea and Dercon, the effect of aid upon poverty is determined by the effect of aid upon growth, and the poverty elasticity of growth. The future expected path of poverty is then assumed to shift down by this amount.

The assumption that aid only affects poverty via growth has no real basis in mainstream economics. As Carter et al. (2015) show, under the mainstream approach of modelling economies as populated by forward-looking welfare-maximising households, when donors transfer resources – aid – to a developing country without placing restrictions on how it is used, households will split those resources between consumption and investment. The same would be true when aid is given to a recipient government that must also allocate revenues between consumption and investment on behalf of households. So growth is not the only way for aid to alleviate poverty: it can also fund current consumption. This point is also made forcefully by Clemens and Kremer (2016).

Funds from the World Bank’s concessional window, the IDA, are also often partially allocated to poverty alleviation expenditure that can be thought of as being more akin to consumption than investment, and there is also the potential for fungibility, wherein to some extent aid recipients are able to substitute aid earmarked for certain purposes for spending that would have taken place anyway and instead increase spending elsewhere.⁵¹ Consumption raises contemporaneous welfare; investment creates economic growth and raises future consumption. Under standard choices for model parameters, households will optimally spend the majority of aid on consumption. Because this is optimal behaviour, restricting aid to funding investment would not be in the recipients’ interest. The idea that recipients will use a large proportion of aid to fund consumption is not a mere artefact of economic theory; it is what we observe in the data (Temple and Van de Sijpe, 2015). Allowing for aid to fund consumption has the potential to significantly alter the conclusions of aid allocation models, when donors have the objective of maximising recipient welfare because it introduces the possibility of raising the welfare of households in stagnant economies, which would otherwise face a future of poverty, even when aid has no impact on growth.

Carter (2014a) shows that even when the impact of aid on growth is not pinned down by current income and institutional quality, and information on the impact of aid on growth

⁵¹ Here ‘consumption’ must be interpreted broadly to include anything that raises the contemporary quality of life, including food, shelter, healthcare and transportation. In reality many goods and services have both consumption and investment characteristics – for example, food is a consumption good, but malnourished people are also less productive; a clean cooking stove is a productive investment, but clean air also raises the contemporary quality of life. Abstract models in which households allocate resources to consumption and investment can be interpreted as capturing the mix of consumption and investment characteristics across the chosen bundle of goods and services.

is available and incorporated into an allocation rule, the impact of aid on growth would have little impact on the allocation of aid. That is because the use of aid to raise consumption dominates the use of aid to fund investment, and because in the neoclassical growth model aid only accelerates growth in countries that would be growing rapidly in the absence of aid. So the optimal allocation of aid barely responds at all to the expected impact on growth, in sharp contrast to Lea and Dercon.⁵²


The main consequence of incorporating aid-funded consumption into an allocation model is to allocate large sums of aid to countries with stagnant economies where aid will have little impact on growth. Although the Lea–Dercon model would not allocate any aid to a country where it is expected to have zero impact on growth, it does allocate more aid to stagnant economies where poverty is expected to persist for a long time. So in a sense it ends up with a similar result but via different means.

⁵² In Carter (2014a) the scope for aid to raise growth is determined by a country's distance beneath its steady-state level of capital per effective worker and hence the marginal product of capital. Some countries may be poor and close to steady-state, others poor and far beneath it, so initial poverty does not pin down the effect of aid on growth. In this model the quality of a country's institutions determines how efficiently aid is used to fund both investment and consumption. When the marginal product of capital is introduced to the allocation rule, simulations demonstrate that very little weight should be placed upon it. This is because aid recipients are modelled as neoclassical economies so that countries with high marginal products of capital would be growing rapidly in the absence of aid.


Appendix 2: The allocation of transfers from the WBG

This appendix shows two illustrative scenarios for allocating the grant element of WBG resource. Both use the PBA rule, but one uses the weights proposed in IDA18 and the other uses the weights that come closest to the allocation determined by Lea and Dercon (2016). The total sum being allocated is \$13.35 billion – see the main text for an explanation. The countries are sorted from richest to poorest, based on 2014 GNI per capita. The allocation shown is dollars per capita.


Figure 6: Grant element allocation scenarios across WBG



Grant elements 2/3



Grant elements 3/3



Source: author's calculations



ODI is the UK's leading independent think tank on international development and humanitarian issues.

Our mission is to inspire and inform policy and practice which lead to the reduction of poverty, the alleviation of suffering and the achievement of sustainable livelihoods.

We do this by locking together high-quality applied research, practical policy advice and policy-focused dissemination and debate.

We work with partners in the public and private sectors, in both developing and developed countries.

Readers are encouraged to reproduce material from ODI Reports for their own publications, as long as they are not being sold commercially. As copyright holder, ODI requests due acknowledgement and a copy of the publication. For online use, we ask readers to link to the original resource on the ODI website. The views presented in this paper are those of the author(s) and do not necessarily represent the views of ODI.

© Overseas Development Institute 2016. This work is licensed under a Creative Commons Attribution-NonCommercial Licence (CC BY-NC 3.0).

Overseas Development Institute
203 Blackfriars Road
London SE1 8NJ
Tel +44 (0)20 7922 0300
Fax +44 (0)20 7922 0399