Aid Dispersion: Measurement in Principle and Practice

CATHRIN FLØGSTAD and RUNE JANSEN HAGEN[*]

University of Bergen, Norway

Summary. — Excessive dispersion of development assistance has been high on the Paris Agenda on aid effectiveness. However, there is no agreement in the existing literature on how aid dispersion should be measured and few studies of the extent of the problem. We argue for using the Theil Index for both recipients and donors. This relative inequality measure has a major advantage: it allows for a perfect decomposition into variation between and within entities. Exploiting this property, we can rank official donors and recipients not only in terms of the total spread, but also assess the contributions of geographic and sectoral dispersion. We provide a detailed picture of developments along various dimensions (globally as well as for countries, income groups, and regions, over 1998–13). We further distinguish between bilateral and multilateral donors. Consistent with other studies using more limited samples, we find little effect of the Paris Agenda overall. Aid is more fragmented in Sub-Saharan Africa and in the poorest countries. Globally as well as for most donor and recipient countries, between variation is the main driver of the spread, lending support to the geographic concentration policies many donor countries have adopted. Bilateral aid has been somewhat more dispersed than multilateral aid and in both cases the large number of donors controlling similar shares of total funds is a major driver of the total spread. The latter suggests that concentration could also be achieved through a reduction of the number of actors on the donor side of the aid industry, a perspective that previous studies using other measures have been unable to capture.
© 2017 Elsevier Ltd. All rights reserved.

Key words — foreign aid, aid dispersion, Transaction costs, Paris Agenda, Theil Index

1. INTRODUCTION

The dispersion of development assistance has been high on the so-called Paris Agenda on aid effectiveness.[1] In short, the argument has been that there are too many actors funding too many activities in too many countries. It is widely believed that this leads to excessive transaction costs, i.e., to spending on planning, monitoring, reporting, and evaluation being disproportionate relative to spending on activities actually generating valued goods and services.[2] It is also argued that the current situation creates incentive problems on both sides of the aid relationship. For example, recipients might suffer from the tragedy of the commons if aid agencies compete for resources such as host government personnel or funds.[3] However, while the usefulness of transaction costs of aid as an analytical concept is reflected in its widespread use in the literature it is not clear that they are measurable.[4] This implies that we cannot directly assess how changes in the structure of aid delivery affect these costs, neither in the aggregate nor for any single actor. Moreover, it is obvious that the optimal level is not zero.[5] A project that is better prepared has a higher chance of being a success. Monitoring progress may reveal that it is lacking, allowing adjustments that put projects back on track to be made or misconceived programs to be terminated before they consume even more resources. Evaluations can provide valuable lessons learned, improving aid effectiveness in the future. One should also bear in mind that the issue of aid effectiveness goes beyond transaction costs and that fragmentation could in principle have positive effects in other dimensions.[6]

Still, we have indications that aid is currently spread too thinly, imposing excessive costs on recipients on average. Annen and Kosempel (2009), Djankov, Montalvo, and Reynal-Querol (2009), and Kimura, Mori, and Sawada (2012) all draw the conclusion that aid dispersion is associated with lower economic growth in recipient countries.[7] Furthermore, in addition to the commitments made as part of the Paris Agenda, several donors have adopted their own policies

of concentration. An interesting question in its own right is then whether these declarations have resulted in lower spreads. Somewhat surprisingly given the attention the topic has received at the policy level, there are rather few academic studies of it. We know of just three that have this as the main focus, as opposed to looking at the consequences of dispersion.[8] Acharya et al. (2006) has a fairly broad coverage of donors (22 bilateral ones) and recipients (179), but only for three years (1999–01). Aldasoro, Nunnenkamp, and Thiele (2010) have a longer time frame (1995–06), but only data for 10 members of the OECD's Development Assistance Committee (DAC). They conclude that despite the Paris agenda, donors have made little progress in concentrating their aid. This conclusion is echoed by Nunnenkamp, Ohler, and Thiele (2013). They have 19 DAC-donors in their sample, which covers the period 1998–09. This makes the "before-after" comparison more credible. Yet, as they split the sample at the halfway point and it ends in 2009, they might not have picked up longer run effects of a process that arguably gathered speed until 2008 at least.

Another limitation of the extant literature is that there is little discussion of and no agreement on how dispersion should be measured.[9] For example, Acharya et al. (2006) use different measures to gauge dispersion for recipients and donors, with no convincing argument as to why this is the correct approach. We will apply their terminology and call the former *fragmentation* and the latter *proliferation*. However, we will use the same measure for both. Fragmentation is the major cause of concern in both policy circles and the academic literature. Still, what donors control is proliferation and the link to fragmen-

tation is not straightforward; focussing on a single recipient could worsen fragmentation there and even a donor taking care to avoid this could see the effort nullified by the actions f other donors. Hence, it is important to check whether reduced proliferation is detectabe on the other side of the relationship. In addition, the spread of donor funds is important in its own right as transaction cost savings could result in larger transfers to recipients for given overall aid budgets. Given that we do not know the transaction cost functions of donors and recipients, it is arguably more consistent to apply the same measure to both types of aid dispersion. To our knowledge, this is the first time this has been done.

Our second contribution is to fully exploit the properties of our preferred index, the Theil. This is a relative inequality measure informing us how far the actual distribution of aid is from the extremes of maximum spread and complete concentration and we argue that there is no reason why this is an inferior alternative to the more commonly used Herfindahl–Hirschman Index (HHI). Moreover, the Theil has a major advantage: it belongs to the only class of inequality measures that allow for a perfect decomposition into variation between and within entities (Shorrocks, 1980). Using this property, we can rank donors and recipients not only in terms of the total spread, but also pinpoint whether the lion's share of it is due to having many partners (between) or to thinly dispersed aid at the sector level (within). In contrast, the standard approach of looking at the HHI calculated at the country level cannot account for the latter and hence could miss an important part of the total variation. And this information has obvious policy relevance as donor countries like the Netherlands, Norway, and Sweden have in recent years adopted policies aimed at reducing the number of partner countries.

Thirdly, the flip-side of perfect decomposability is perfect aggregation. While previous studies have focussed on individual donor and recipient countries, we can group these consistently in various ways. On the recipient side, we look at differences in fragmentation across regions and income levels, as well as aggregating all the way up to show the global picture. This enables us to provide new perspectives, including whether fragmentation globally is driven mainly by a relatively equal distribution across recipients or by high dispersion within them. On the donor side, we are able to study bilaterals and multilaterals separately. While proliferation is limited for most multilaterals by mandates that are restrictive in terms of geography or sector, their aggregate contribution is of interest, particularly in light of the increasing number of such actors (c. f. Figure 2 below). We believe our study is the first to analyze the consequences of this trend for aid dispersion.

Our final contribution is to look at these issues in a longer time-frame (1998–13) than previous studies. Consistent with these, we find little effect of the Paris Agenda on either fragmentation or proliferation. In fact, dispersion has increased globally. There are also both more donors and more recipients recording higher spreads in the latter half of our time frame than those seeing reductions. Apparently, the various international declarations and individual aid policies have not had much bite in practice. Fragmentation is more severe in Sub-Saharan Africa and in the poorest countries. Both globally and for most donor and recipient countries, between variation is the main driver of the spread, lending support to the geographic concentration policies mentioned above. Bilateral aid has been somewhat more dispersed than multilateral aid. Proliferation by both donor types are in the aggregate mainly caused by there being many actors with quite similar shares of total bilateral and multilateral aid, respectively. This finding points to a neglected part of the picture, viz. that other things

being equal concentration could also be achieved through a fall in the number of donors. Since it is likely that there will be even more bilateral donors in the future as emerging economies initiate their own aid programs, this can probably only be achieved through a reduction in the number of multilaterals.

The rest of this paper is organized as follows. In the next section, we discuss principles for measuring aid dispersion and state formulas for different variants of the Theil. Our data are described in Section 3. Section 4 contains the aggregate results, while the topic of Section 5 is developments in individual donor and recipient countries. As a robustness check, the correlation of the Theil and the HHI is briefly analyzed in Section 6. Finally, we summarize our findings in Section 7.

2. MEASURING AID DISPERSION: PRINCIPLES

The extent of fragmentation (within a recipient country, across donors) or proliferation (by a donor, across recipients) concerns how a certain sum (total aid to a recipient country or total aid by a donor) is spread across entities, which could be projects, thematic sectors, or recipient countries. There are many different measures of dispersion that could be used, but little discussion and no consensus in the literature on which of these are preferable. Some of those actually applied are fairly ad-hoc and/or only capture part of the phenomenon. This can be said of expressing fragmentation in terms of the number of donors, for example. In this section, we dispute the view that concentration measures like the HHI are better than inequality measures like the Theil in terms of capturing the effects of dispersion and contend that the best we can do currently is to assess dispersion itself. In our opinion, the Theil Index does this as well as the alternatives. Moreover, its perfect (dis)aggregation property opens new and policy-relevant perspectives on the topic at hand.

Dreher (2010, p. 11) argue that

To be appropriate for the assessment of in-country aid fragmentation [an] index should ideally fulfill all of the following requirements. It should (1) reflect fragmentation in a *theoretically correct* way, (2) be *easily understandable and computable*, and (3) use a functional form appropriate to reflect the *problems* involved with in-country aid fragmentation. (Emphasis in original)

These principles are sensible. However, on further reflection they are not easily applicable. The main problems concern requirements (1) and (3). What we ideally would like to have is a measure that relates fragmentation to transaction costs. However, we have neither a theoretical model nor empirical estimates of this relationship. Country- or sector-specific factors might imply that a certain level of fragmentation is more or less harmful, but there is currently no way of picking up these in an applied analysis. Moreover, we lack the data to take fully into account whether donors use aid modalities such as sector-wide approaches or multi-donor trust funds that are often argued to entail lower transaction costs.[10] Finally, although Dreher (2010) claim that due to (1) concentration measures are preferable to inequality measures, their argument is not completely consistent.

The HHI is probably the most frequently used basis for quantifying the effects of fragmentation in the academic literature.[11] Dreher (2010) find that it is overly sensitive to an increase in the number of donors at low levels. For this reason, they prefer measures that capture the cumulative shares of the 3–5 largest donors. However, these are ad-hoc and there is no way of knowing which is the "correct" one. Moreover, both these and the HHI are based on shares, like the Theil, and