

**The Distributional Impact of Large Dams:
Evidence from Cropland Productivity in Africa**

Abstract

We examine the distributional impact of large dams on cropland productivity in Africa. As our unit of analysis we use a scientifically based spatial breakdown of the continent that allows one to exactly define regions in terms of their upstream/downstream relationship at a highly disaggregated level. We then use satellite data to derive measures of cropland productivity within these areas. Our econometric analysis shows that while regions downstream benefit from large dams, cropland within the vicinity tends to suffer productivity losses during droughts. Overall our results suggest that because of rainfall shortages dams caused a net loss of 0.96 per cent in production in Africa over our sample period (1981-2000). However, further dam construction in appropriate areas could potentially lead to large increases in cropland production even if rainfall is not plenty.

JEL Classification: O20, Q19

Keywords: dams, agricultural productivity, Africa

“One day every last drop of water which drains into the whole valley of the Nile...shall be equally and amicably divided between the river people, and the Nile...shall perish gloriously and never reach the sea”

Winston Churchill, 1908

Section I: Introduction

It is well known that many African economies are heavily reliant on rainfed rather than irrigated agriculture as a source of income.¹ However, Africa’s rainfall regime is also notably erratic, making drought and desertification a regular phenomena in the region. As a matter of fact, Barrios et al (2008) have shown that a lack of rainfall may have lowered agricultural production in Africa in the latter half of the 20th century by as much as 32 per cent relative to the rest of the developing world. Additionally alarming in this regard is that some climatologists predict that water scarcity on the continent is in the future likely to further increase. It is important to note, however, that at least in terms of relative levels of overall freshwater resources Africa is actually not particularly bad off. As a matter of fact, its global share of freshwater is roughly equal to its global share of population (UNECA, 2005). What makes Africa different to most other continents, though, is that due to its great geographical and climatic diversity water as a scarce natural resource is spatially relatively unevenly distributed. For example, there are large portions of the continent, such as along the central forest rainforest belt, which are particularly rich in precipitation, whereas others receive little or only very sporadic rainfall.² Similarly, while Africa is home to some of the largest rivers and lakes in the world, these occupy only parts of Africa. It is not surprising then that the construction of large dams has had a long history of being promoted as a relatively affordable means to spatially re-distribute water as a resource across Africa from places of abundance to those in need.

¹ For instance, FAO data shows that the percentage of agricultural land irrigated in Africa is only about one eighth of that of the rest of the world.

² For instance, the Congo basin contains 30 per cent of total freshwater resources while harboring only 10 per cent of the continent’s population (UNECA, 2005).

Despite the ability of dams to reduce relative water scarcity, Africa as a continent is generally viewed as still being far from its potential.³ For instance, there is about one dam to every 683,000 persons in Africa, where the equivalent figure for the rest of the world is 168,000. Part of the reason for the relatively low level of dam infrastructure in Africa today is of course that the original enthusiasm for funding large dam projects by international donors had over time somewhat faded. More specifically, in the 1990s there emerged considerable concern about the fact that much of the actual distributional environmental and economic impacts of large dams had largely been ignored. The growing criticism eventually cumulated into the creation of the World Commission of Dams, which set up a committee to undertake a global review of their impact, and its judgment was in general rather damning, stating that "...a lack of equity in the distribution of benefits has called into question the value of many dams in meeting water and energy development needs when compared to alternatives" (WCD, 2000, p. XXVIII). Nevertheless, recently there appears to have been a return in the interest by international donors, such as the World Bank, in such large scale infrastructure projects.⁴

Arguably one of the most important reasons behind the unresolved controversy of how desirable large dams really are is that the evidence on their distributional impacts is scant at best. That is, while there are a number of case studies, the only comprehensive quantitative analysis is the innovative study by Duflo and Pander (2007). More specifically, Duflo and Pander (2007) set out to examine the distributional effect of large dams across districts in India. Importantly, their results suggest that while downstream districts can benefit in terms of increased agricultural production from dams located upstream, dams appear to have no net effect on agriculture in the districts they are constructed in, but do amplify the effects of rainfall shocks.⁵ The authors thus argue that neither markets nor institutions have alleviated the adverse distributional impacts of dam construction in India.

Given the differences in geography and policy contexts, it is of course not clear whether Duflo and Pander's (2007) findings can be generalized to African agriculture. In this paper we

³ UNECA (2005) argue that in Africa only about 5 per cent of its water management potential has been exploited.

⁴ For example, the World Bank launched its re-entry with its 2003 'high-reward/high-risk' water sector strategy; see Hathaway and Pottinger (2008).

⁵ They also find that dams generally increase poverty in the vicinity.

thus explicitly set out to provide a first study of the impact of large dams in Africa, focusing specifically on the distributional effect in terms of cropland productivity.⁶ However, our contribution to the literature is not just limited to this aspect. More specifically, we also make a number of methodological contributions. Firstly, Duflo and Pander (2007) examine the effect of dams only in terms of administratively defined regions. As we show here, these can at best only roughly provide a classification of upstream/downstream relationships of areas. We thus instead use a scientific spatial breakdown of the African continent into upstream/downstream areas based on topographic elevation and river network data that is independent of any ‘administrative’ classification. One problem of course with using a fairly disaggregated spatial unit of analysis is that particularly for developing countries consistent agricultural data over time tends to be sparse. Our second contribution hence lies in using spatially detailed satellite derived information to identify cropland and changes in its productivity over time. Overall this provides us with annual data of over 3,600 local areas covering a twenty year period (1981-2000) for which we are then able to evaluate the distributional impact of large dams on cropland productivity for the entire African continent.

The remainder of the paper is organized as follows. In the following section we outline the general framework and challenges involved in estimating the distributional impact of dams on cropland productivity. In Section III we introduce and describe our data sources. Our instrumental variable strategy is outlined in Section IV. Econometric results are provided in Section V. Section VI contains some simple back of the envelope calculations regarding the actual and potential aggregate impact of dams in Africa. Finally, the last section contains concluding remarks.

Section II: Empirical Framework

Our main goal is to investigate the distributional impact of dams on cropland productivity in Africa, a task which inherently entails a number of conceptual, data, and estimation

⁶ One should note that dams may have many other purposes, such as generating hydropower, that need to be taken into consideration when assessing their overall potential benefits. Our paper should thus be evaluated in terms of agricultural production.

challenges. For illustrative purposes of the underlying framework of the empirical approach we use to tackle these, we postulate the hypothetical case of a region consisting of agricultural plots and a river network, as shown in Figure 1. In this context we then consider the potential impact of the construction of a dam on the main river in the location as shown.

As can be seen, our hypothetical region consists of 6 different agricultural plots, for each of which the nature of dam engineering and features particular to African cropland production imply a priori expectations regarding the impact of the dam. More specifically, Plot #1, is assumed to be located entirely within what is to become the dam's catchment area. In this regard, Duflo and Pande (2007) argue that one should expect that it is likely to lose productivity due to the dam. This arises because water seepage from the reservoir and canal increases water logging and soil salinity. Moreover, there may be water restrictions in the catchment area to ensure maximum water storage.

In contrast, in the command area agricultural land, Plots #2 and #3 in our graph, are likely to be able to enjoy the productivity benefits from irrigation, since the fixed cost of accessing irrigation within the canal network of the command area is generally lower than for other forms of water harvesting. While we assume that downstream Plots #4 and #5 are too far away to avail of the benefits of an irrigation canal network, this does not necessarily mean that they are unaffected. As a matter of fact, Duflo and Pande (2007) note that potentially the beneficial effects of a dam may extend much further downstream from its command area if it is explicitly used to prevent floods and droughts by regulating the flow of water downstream, although there is a trade-off between these two goals. In the African context, there is another important aspect to consider with regard to agriculture further downstream from the command area. More specifically, about half of Africa's total wetland area is of the floodplain type⁷, and due to their extremely fertile soils these areas play an important role in local agriculture.⁸ In this regard, upstream dams can, by influencing the transfer of water and sediment downstream

⁷ Floodplains are relatively flat areas adjacent to rivers created by sedimentary deposits of meandering channels as well as periodic flooding.

⁸ Some cover several thousand squared kilometers in area, such as, for instance, the Inner Niger Delta in Mali and the Kafue Flats in Zambia, amongst many others.

drastically, affect the inundation area and soil fertility of a floodplain; see Barbier (2002).⁹ Thus in the context of our hypothetical example we postulate that the dam may very well also have a negative effect on plots #4 and #5. Finally, one should expect plot #6, which is not downstream to the dam, to experience no change in productivity.

Our simple example makes it clear that in order to completely disentangle the distributional effect of a dam on cropland productivity one would ideally like to know the location of dams, the spatial extent of its catchment and command areas, as well as what agricultural plots lie further downstream. Unfortunately, detailed comprehensive spatial data on dams' catchment and command areas do not exist in the African context.¹⁰ Nevertheless, one can still investigate distributional effects if one knows the location of a dam and the location of cropland in terms of its downstream relationship to the dam, and then makes some reasonable assumptions with regard to the likely spatial extent of the command and catchment areas. In terms of our hypothetical example we can illustrate this by assuming a spatial breakdown of our area in terms of a upstream/downstream relationship of sub-regions according to the river network, as shown by the four quadrants in Figure 1. Accordingly, Quadrant I lies immediately upstream to Quadrant III (and hence Quadrant III immediately downstream to Quadrant I). Similarly, Quadrant IV, although not neighboring on Quadrant I, lies (further) downstream from it. In contrast, while Quadrant II neighbors on Quadrant I it lies neither down- nor upstream from it. One should note that here we have also made the assumption that the command area lies completely within the quadrant that the dam is located in. We also are postulating that both the dam's own as well as a downstream region may contain some of its command area. In other words dams are not necessarily always located on the border of where one area can be classified downstream to another. We shall address the plausibility of these assumptions in the next section.

⁹ For example, Thompson and Polet (2000) have estimated that in part due to upstream water development the maximum extent of flooding in the Hadejia-Jama'are area in Nigeria has declined by about 60 per cent since the 1960s.

¹⁰ As a matter of fact, we know of no other country, let alone continent, where such comprehensive spatial data would be available for all dams.

What would one then expect the net impact of a dam constructed in Quadrant I on agriculture in all quadrants to be? Clearly Quadrant II, which although neighboring is neither up- nor downstream, should be unaffected. For the dam's own quadrant the net a priori effect would be ambiguous since we assume that it contains all of the catchment as well as part of the command area. In the immediately downstream Quadrant III there are again two opposing effects. First, there is the productivity boost for land located within the command area. However, whether the overall a priori expected effect is positive will depend on whether other agricultural land that is not part of the command area is negatively affected – as, for instance, when it is of the floodplain type. Thus even for the area immediately downstream, i.e., Quadrant III, the net effect is a priori ambiguous. Finally, even in the further downstream Quadrant IV, the effect will be ambiguous since dams may be used to prevent floods and droughts in these areas, but changes in the river flow may also negatively affect the productivity of floodplain cropland.¹¹

From the discussion thus far it becomes clear that our task, in the face of no exact spatial information on dams' command and catchment areas, reduces to disentangling the effect on cropland productivity of dams located both within its vicinity and upstream from it. This translates into the following econometric specification:

$$P_{it} = \alpha + \beta_1 DAMS_{it} + \beta_2 UDAMS_{it} + \beta_3 X_{it} + \varepsilon_{it} \quad (1)$$

where P is region i 's cropland productivity, $DAMS$ refer to the number of dams within the region, $UDAMS$ are the number of dams that are located upstream from i , X is a vector of other explanatory variables, and ε is the unexplained error term. As noted above, a priori the expected signs of the coefficients β_1 and β_2 are ambiguous and an empirical matter.

While at first look simple, estimation of (1) in the African context, however, poses a considerable number of challenges. Firstly, one requires geo-referenced data of African dam location. Secondly, it is necessary to be able to correctly classify Africa into regions with reference to their upstream/downstream relationship. Within this spatial breakdown on then

¹¹ One should note that Duflou and Pande (2007) argue that the net effect in the downstream area will be unambiguously positive, while they do not consider any impact further downstream from dams. The key difference is that they do not consider floodplain agriculture as would be important in Africa.

needs to identify areas of interest, namely those used for cropland production and find some proxy of productivity for these. Finally, from an estimation point of issue, it is unlikely that dams were randomly allocated across the African continent. Rather dams are likely to have been strategically placed according to cost-benefit considerations, as well as influenced by political factors. In this regard, a simple OLS regression of (1) is likely to produce biased and inconsistent estimates of β_1 and β_2 unless one has enough information to include all relevant factors in X . Our strategy in overcoming the data challenges is the subject of the next section, while we outline our instrumental variable strategy to deal with the endogeneity problem in Section IV.

Section III: Data

A. Dams

In order to identify and geographically place dams on the African continent we use FAO's African Dams Database, which is a geo-referenced data base of large dams, providing the exact latitude and longitude of location.¹² Dams are considered to be 'large' if they have a height of at least 15 meters or a reservoir capacity of at least 3 million m³. Of the 1138 dams in the data, for 145 the date of construction was missing. Where possible we completed the missing information from updated ICOLD data and internet searches, leaving us with a total of 1032 dams for our analysis. The database also contains information on the major purpose(s) for 75 % of our sample. In this regard, 65% have irrigation listed as at least one of their major purposes, while for 15% hydropower is one of the major purposes. One may want to note, however, that even for those dams that did not have irrigation listed as one of their major purposes or those that had no major purpose listed, additional information from the internet, if available, always indicated that irrigation was at least a side product. We will thus as a benchmark operate under the assumption that all dams are likely to reduce the fixed costs of an irrigation system in their command areas.

¹² In order to identify dams the FAO resorted to The World Register of Dams, national reports, information obtained from national experts, and the internet.

We depict the location of the 1032 dams in Figure 2. As can be seen, these are not equally spatially distributed, with most located in the western-southern, central-western or eastern-northern parts. One may want to also note that there are many areas containing larger rivers and lakes in which there are no dams in proximity.

B. Spatial Unit of Analysis

As outlined earlier, being able to properly quantify the distributional effect of dams crucially depends on being able to identify how areas are interlinked in terms of the flow of water, i.e., that is in terms of their upstream/downstream relationship to each other. One possibility is to simply use administrative areas and classify these accordingly. This is exactly what Duflo and Pande (2007) do for their Indian study under the justification that visual inspection of available dam command and catchment maps revealed that most catchment areas are contained within the districts of their located and that one can classify districts as upstream/downstream to each other. However, there is no a guarantee that one can apply a similar classification by administrative areas in the African case. For instance, visual inspection of administrative regions in Africa at the same level as Indian districts¹³ reveal that many borders of regions run at least partially along rivers – see, for example, the case of the southern African tip in Figure 3.¹⁴ If one considers such a case in terms of our hypothetical example of Figure 1, where the area would be broken up into two regions along the main river, then it becomes clear that it would be difficult to say anything about a priori expectations regarding distributional effects. More specifically, in our case both areas contain parts of both the catchment and the command area, as well as the potentially affected plots of floodplain agriculture downstream. Moreover, plot #6, which is neither up- nor downstream to the dam, would be included as a potentially affected in the analysis.

As an arguably superior alternative we here use a scientific spatial breakdown of Africa in terms of upstream/downstream relationships of regions that is based on actual hydrological data. More precisely, we use information from the U.S. Geological Survey's Data Center's HYDRO1K data set, which provides drainage basin boundaries data for the African continent as

¹³ Administrative Level 2 of the GAUL system.

¹⁴ This is true at both national and sub-national levels.

derived from river network and flow direction data.¹⁵ At its most disaggregated level this involves dividing the African continent into 7131 6-digit hydrological basins, with an average area of 4,200 km². Each basin is then assigned a Pfafstetter 6-digit code which allows one to determine whether it is upstream, downstream or not related to another basin in the data set.

We depict the spatial breakdown of the African continent according to our 6-digit basins in Figure 4. Accordingly, basins vary greatly in shape and size. Moreover, many basins cross national borders. As a comparison, we depict these also jointly with the outline of the administrative units in the Southern African tip, as shown in Figure 5. As can be seen, there is little correspondence between the two. This of course may just be a feature of Africa and we hence in Figure 6 also depict the comparison of a 6-digit Pfaffstetter code breakdown of Southern India in conjunction with the district spatial disaggregation used by Duflo and Pande (2007). As with Africa a similar picture emerges, where at least visually there appears to be no obvious correspondence between the two spatial classifications.

Finally, in terms of a priori expectations regarding the impact of dams within our empirical framework we made the assumptions that dams were generally not located exactly on the border of basins and that the catchment area was usually contained within a dam's own vicinity. With regard to the former, while not shown here, a closer look at our sub-basins in conjunction with the plots of dams revealed this almost always to be the case. For the latter assumption we examined specifically South Africa for which there is some spatial data of reservoirs of dams found that in most cases the catchment area was indeed within a single sub-basin.¹⁶

C. Cropland Productivity

Unsurprisingly, time varying agricultural measures of cropland productivity, such as crop yields, collected by statistical agencies or independent surveys to such a fine spatial extent as our hydrological regions do not exist. We thus proceed to construct a proxy for local cropland productivity by using two satellite data sources. The first is the Global Land Cover 2000 data set

¹⁵ These are derived from a 30-arc second digital elevation model.

¹⁶ Detailed plots of these checks for both assumptions are available from the authors upon request.

(GLC 2000) which allows us to identify cropland locally within Africa. More specifically, the data classifies land cover across the globe into 22 distinct land cover categories at the 1km resolution based on 14 months (1 Nov. 1999 - 31 Dec. 200) of daily 1-km resolution data acquired by the SPOT 4 satellite. We used land cover categories (i) cropland, (ii) mosaic of cropland / shrub or herbaceous cover, and (iii) mosaic of cropland / tree cover / other natural vegetation to identify cropland area within our river basins. One should note in this regard that 49% of the 6-digit code basins did not, according to the GLC 2000, contain any cropland.

It is of course important to point out the possible drawbacks of using this satellite data to identify cropland in our analysis. For one, given that the classification is at a 1km resolution, it is likely that one is only capturing larger cropland areas and not small farming. Moreover, the fact that this data is only available for 2000 means we cannot capture any changes in cropland area over our sample period. It would be difficult a priori to determine what kind of bias this might create in our estimated distributional impact. In some areas farmers may cease to plant crops if a dam had reduced cropland productivity enough, which would tend to underestimate the negative impact of cropland productivity captured using the time invariant 2000 classification. In other areas, in contrast, the access to irrigation may induce an expansion of area used for crop production. Moreover, whether this would tend to over- or underestimate the positive impact of a dam will depend on what the use of the expanded area was prior to its agricultural cropland classification in 2000.

We depict the cropland areas in Figure 7. It becomes apparent that much of the cropland lies within the semi-arid belt just outside the equator, where rainfall is known to be erratic. This is mostly due to the fact that in the more rainfall plenty zones humidity is high and hence would make agriculture extremely susceptible to diseases, while in the arid zones rainfall is too scarce; see Barrios et al (2009). One may also want to note that most dams are located near areas designated as cropland.

Given our identification of croplands within our hydrological regions we next needed data to measure its productivity over time. In this regard we resort to the concept of ‘net primary production’ (NPP). More specifically, ‘production’ refers to the creation of new organic matter. For example, when a crop of wheat grows, new organic matter is created by the process of

photosynthesis, where light energy is converted to energy stored in plants, in turn spurring plant growth. NPP then is the rate at which an ecosystem's producers convert solar energy into chemical energy as biomass minus a plant's respirations costs, measured in terms of kcal/m²/year. In essence NPP quantifies the conversion of atmospheric CO² into plant biomass and the resultant values can then serve as a proxy of cropland productivity.¹⁷ As noted by Hicke et al. (2004), one of the advantages for using NPP to proxy cropland productivity over large areas and over time is that it, unlike economic data, provides a common metric among different crop types, thereby facilitating comparisons and aggregation over all crop types.

Here we use NPP values derived from satellite images of spectral reflectance of terrestrial vegetation. The physical basis for the observed correlation between spectral reflectance and NPP is the existence of a relationship between spectral reflectance and the absorption of solar radiation by vegetation canopies, and in turn the link between the amount of absorbed photosynthetically active radiation and its utilization for NPP.¹⁸ Our source of data for NPP for the African continent are the MOD17A2 measures derived from observations of the Moderate Resolution Imaging Spectroradiometer (MODIS) on the NASA Earth Observing System (ES) Terra satellite, as computed by the Earth Observing System Data Information System (EOSDIS). The data are available at the 1 km spatial level on an annual basis for the period 1981 to 2000. As an example of the temporal variation in NPP we depict its values for cropland located within the Central Western African semi-aridity belt in Figure 8 for 1988 and 2000. Accordingly, the values of NPP were visibly lower during the substantially drier year of 1988. For our purposes we use the NPP data to calculate average annual productivity values for cropland within our spatial regions.

D. Other Explanatory Variables (X)

¹⁷ Examples using NPP to proxy cropland productivity include Heinsch et al (2005) and Hick et al (2000a, 2000b). One should note that there are also numerous studies have used actual crop yield data to derive NPP estimates; see, for instance, Monfreda et al (2008) and Veron (2002) to name a few.

¹⁸ It is of course of interest to know how well satellite derived measures of NPP are able to approximate productivity derived from ground data computed alternatives. In this regard, Lobell et al. (2002), for instance, have shown that there is good agreement between yield and satellite derived estimates of NPP for the US.

Given the lack of regional data and our particular spatial classification the only control variable that we could include in X in (1) was population density. The data for this is derived from the African Population Database (APD), which provides information on the spatial distribution of population of the region for 2.5 minute grid cells for 1960, 1970, 1980, 1990, and 2000. Using these cells we calculate population densities within each of our spatial 6-digit Pfaffstetter code units for 1980, 1990, and 2000, and then interpolate these linearly to arrive at annual figures.

E. Summary Statistics

The data restrictions given above imply that we have a total of 3619 spatial units for which we were able to identify at least some cropland and for which we could calculate annual cropland productivity over the 1981 to 2000 period. Summary statistics for the variables described above for this sample are given in Table 1. As can be seen the NPP of cropland has a mean of 1287, but is characterized by considerable variation. One also finds that on average 21% of regions with cropland have at least one dam within their area, while 7% have a dam upstream.

Section IV: Instrumental Variable Strategy

As noted earlier, dams are unlikely to be randomly allocated across regions. In their study of Indian districts Duflo And Pande (2007) address this potential endogeneity problem by using the share of dams located in a state (which is at a more aggregate level than districts) prior to their sample period and multiplying this by the total dams in India. They then regress the district level of dams on this state level variable interacted with district level river gradient indicators to predict arguably ‘exogenous’ district level dam construction proxies. One should note in this regard that they use the river gradient proxies on the grounds that certain river gradients are more appropriate for dam construction than others. In contrast, the use of state share of dams implicitly assumes a state level variation in the policy of dam allocation and that this can be proxied by the ex-ante state share interacted with time dummies.

A. Policy Context

In contrast to Duflo and Pande (2007), we in this paper take advantage of an explicit policy context driving local major dam construction on the African continent to build an instrument for dam location. More specifically, it is well known that water as a resource in Africa is of a transboundary nature. For example, every country on the African continent has territory in at least one transboundary river basin¹⁹ and these basins cover over 60 per cent of Africa's total land area. It also appears to be a stylized fact that "internal water is a resource whose characteristics tend to induce cooperation" (Wolf and Hammer, p. 66), and Africa, despite its well known high incidence of civil and international conflict, is no exception in this regard.²⁰

Focusing specifically on dams, arguably a substantial proportion of their construction was preceded by the establishment of transboundary river basin treaties and authorities.²¹ As such the earliest such water related transboundary agreements coincided with the period of European colonization of the continent.²² The use of treaties to deal with the management of water resources then carried on into and accelerated with the beginning of political independence for many African countries in the late 1950s. One may want to note that, in contrast to the treaties signed in the colonial period, where most focused on the division of water resources or encouraged construction of dams, those of the early independence period tended to emphasize more general joint management and water development by establishing river or basin authorities. However, even within these more general terms dam construction tended to an important component.²³

Of course one may ask oneself what the incentives were for using transboundary agreements to predate dam construction in Africa. As argued by Dinar (2008), there are a

¹⁹ Transboundary river basins are equivalent to one and two digit Pfaffstetter code classifications in the HYDRO1K data set.

²⁰ For example, if one uses data from the International Water Events Database, which has meticulously compiled all water related international 'events' and classified these on a 'good' and 'bad' scale since 1950, then in Africa those that are positive events clearly outnumber those that are negative in essentially every year.

²¹ See Lautze and Giordano (2005) for a more detailed account of the history of transboundary water agreements in Africa.

²² See Lautze and Giordano (2005)

²³ For example, in their review of the evolution of transboundary agreements Lautze and Giordano (2005) note that about three quarter of the treaties explicitly cited dam construction to facilitate hydropower development and/or expansions in the area of irrigated land as a goal.

number of reasons rationalizing cross-national cooperation through the use of transboundary agreements in terms of managing water resources. Firstly, there is often no clear division even within river basins of one country as upstream or downstream to another if they share more than one river. However, even if a state is more or less upstream to another it may have an incentive to cooperate with its downstream neighbor in dam construction. For instance, landlocked countries, as many are in Africa, may be able to gain navigational rights to offset the upstream advantage. Treaties and organization may also ensure that an upstream state can guarantee side payments for the benefits of dams that accrue to the downstream nation.²⁴ Furthermore, dam construction as a result of cooperation can resolve funding problems. For instance, downstream nations may be willing to share a large burden of the costs if they are the ones to benefit the most.²⁵ Perhaps even more importantly in this regard, transboundary cooperation generally served as a prerequisite for funding from multilateral organizations for dam construction. As a matter of fact, in early periods of dam construction, the big dams in Africa were primarily financed by the World Bank which viewed transboundary cooperation generally as a prerequisite for receiving funds; see Hathaway and Pottinger (2008) and Lautze et al (2005).^{26,27}

B. Construction of Instruments

In order to construct a quantitative measure of the policy context of transboundary water related agreements we avail of two databases. The first is the International Freshwaters Treaties Database, which provides a comprehensive collection of international freshwater related agreements since 1820, including summaries of these, as well as coding them according to the year signed and the river basin and countries involved.²⁸ We additionally use information from

²⁴ For instance, under the Lesotho Highland Water Project South Africa (LHWP) bears the full cost relating to water delivery while Lesotho pays of the cost of the hydropower component, which is about 5 per cent of the total cost of the LHWP.

²⁵ Again, the LHWP serves as a prime example of such a case.

²⁶ The first dam to be officially financed by the World Bank was the Kariba Dam in Zimbabwe in 1959. While the World Bank continues to play an important part in financing of dams in Africa, there are now also a number of other multilateral donor organization actively involved.

²⁷ For example, the World Bank (1995) has explicitly state as a policy objective that "...it attaches the utmost to having riparians enter into appropriate agreements" and that "...a country) proposing to execute any project which will regulate, abstract or otherwise change river flows must notify co-riparian states of its intentions so that each state may consider whether it wishes to lodge an objection".

²⁸ See <http://www.transboundarywaters.orst.edu/database/interfreshwaterdata.html>

the database on the historical formation of transboundary river basin organizations in Africa compiled by Bakker (2006). Combining these two data sources reveals a total of 98 treaties and basin organizations formed, involving 53 countries and 59 basins, since 1884. The 59 basins involved are shown in Figure 9. There are two noteworthy features in this regard. Firstly, as emphasized earlier, the river basins are clearly transnational, cutting generally across several countries. Secondly, their size, and hence their potential extent of coverage, is at a substantially larger scale than the individual regions that we use as our geographical unit analysis to identify up- and downstream relationships. With regard to the latter one may want to note that examination of the individual agreements shows, as mentioned earlier, that most only cover a sub-portion of countries that have at least some part of their territory within a basin. Moreover, the raw data revealed that the group of countries involved in any agreement pertaining to a particular river basin varies considerably over time. Within our constructed database, an agreement covers on average 245 Pfaffstetter 6-digit code regions, where the minimum and maximum number are 2 and 1054, respectively. We depict also the historical evolution of the number of agreements and dams constructed since 1920 in Figure 10. As can be seen, both the number of agreements signed and dams constructed have increased in a roughly similar pattern over this period. Moreover, it appears that at least in aggregate for many segments of time a rise in the number of dams constructed is preceded by the number of treaties signed.

To create an instrument variable for dams from the information on relevant agreements we calculated a five year moving average of the total number of agreements covering each spatial unit each year. Our choice of a moving average is a result of intensive internet searches that seemed to suggest that many dams were built within five years of commencement of the project. Similarly to Duflo and Pande (2007) we also consider the river gradient of an area as an important determinant of its suitability for dam construction. In this regard we follow Duflo and Pande (2007) and use gridded elevation data²⁹ of each spatial unit to isolate the cells within each hydrological basin through which rivers flow and calculate the fraction the fraction of these whose gradient is flat (1.5 per cent), moderate (1.5-3 per cent), steep (3-6 per cent), and very

²⁹ We use the elevation data from the HYDRO1K data set which provides elevation data in terms of mapping Africa into 1km*1km cells.

steep (above 6 per cent). The interaction terms of these three gradient variables with our treaties proxy then constitute our set of instruments for dam construction.

Section V: Econometric Results

A. First Stage Regressions

As a first step we use our set of instruments to create predicted values of the number of dams for our hydrological regions. One should note that we use the full sample of 6 digit level areas, rather than just those with cropland, since many dams are located in areas cropland free areas but still may affect regions downstream. The results of various specifications of regressing local dam construction on transboundary agreements are given in Table 2, where we include time dummies as well as control for region specific fixed effects.³⁰ Moreover, given the spatial nature of our data, we calculate standard errors which are robust to very general forms of spatial and temporal dependence using the nonparametric covariance matrix estimator proposed by Driscoll and Kray (1998).³¹

In the first column we first only include the five year moving average of the cumulative number of treaties which cover the region. Reassuringly, the coefficient on this variable is positive and significant, suggesting that treaties at the basin level increase the number of dams constructed in our spatial regions. In the second column we then interact the treaties coverage proxy with our three river gradient variables. Accordingly, all interaction terms are significant, indicating that treaty coverage has a greater effect on regions with gradients that are not very flat. If one compares the coefficients across interaction terms, their size shows that this additional effect is more than twice than for the low and medium river gradient interaction terms. This latter result may not be surprising given that most of the major dams in Africa appear to be multi-purpose and that, as noted by Duflo and Pande (2007), high gradient areas are likely most

30

³¹ We allow for an autocorrelation structure for up to 4 lags, although altering this made no qualitative difference in our results.

suited for purely hydropower dams.³² In the third column we also experimented with using a 10 year moving average of our treaties proxy. As can be seen, although the resultant coefficients are qualitatively similar, their size is noticeably smaller.

B. Second Stage Results

Results for estimating our main equation (1), where the dependent variable is are logged values of NPP, are given in Table 3.³³ We first use the actual own area dam and upstream dam numbers under OLS, as shown in the first column. Accordingly, both own area dams and upstream dams have a positive and significant effect on cropland productivity, although the coefficient on the latter is nearly six times larger than the former. Introducing fixed effects, as shown in the second column renders the effect of own area dams insignificant, however. Moreover, the coefficient on UDAM is reduced by nearly 40 per cent. This suggests that time invariant unobservables that determine both dam location as well as crop productivity are likely to be important. We next proceeded in using our predicted values of DAM and UDAM as obtained from the second column of Table 3. While using the predicted own area and upstream dam numbers under a fixed effects estimation, shown in Table 3, produces qualitatively similar results to the simple fixed effects estimations, notably the coefficient on UDAM is more than twelve times larger. This suggests that not controlling for other time varying factors that are correlated with both dam location and cropland productivity will tend to produce a substantial downward bias in the effect of upstream dams. One may want to note that our results are in line to those of Duflo and Panda (2007) in that we find that a net zero own area effect, but that dams located upstream can aid local agriculture.

C. Robustness Checks

We next proceed to undertake a number of robustness checks as shown in Table 4. First we experiment with using values for DAM and UDAM predicted from the 10 year moving average of cumulative treaties and gradient interaction terms taken from the third column of Table 3.

³² For example, when one re-runs this specification only for dams for which hydropower was listed as one of the major purposes, one finds that the interaction term with the high category gradient is more than twice that of the low gradient category. Detailed results are available on request.

³³ As for the first stage regressions we use the Driscoll and Kray (1998) adjusted standard errors.

Accordingly, the results are very similar, although with a slightly lower coefficient for UDAM. In the second column of Table 4 we also included the predicted number of dams of neighboring non-downstream, non-upstream regions into our base specification.³⁴ Accordingly, and as expected, it is only dams in upstream areas, that have an effect on local cropland productivity.³⁵

Thus far we have restricted our analysis to the effect of dams located immediately upstream. As argued earlier, a dam may feasibly have much more far-reaching impact. To investigate this we additionally include in the subsequent column the number of (predicted) dams in areas that are upstream to the immediate upstream areas. However, as can be seen, unless cropland is immediately downstream, it is not affected by upstream dams.³⁶

Finally, we also re-estimated our base specification for the mean NPP value of all other non-cropland areas within a spatial unit. Assuming that our cropland designation is relatively accurate, then a positive effect of upstream dams on local non-cropland areas would suggest that upstream dams may have a more general effect on NPP regardless of whether the land is used for cropland or not. In contrast, a negative coefficient on UDAM might imply that upstream dams disrupt the natural water flow system enough to have detrimental environmental effects on non-cropland areas. A similar line of reasoning could be employed for the interpretation of a significant coefficient on DAM. However, as shown by the last column in Table 4, using values NPP of non-cropland areas as the dependent variable for regions shows no significant effect of dams, either in their own region or downstream.

D. Droughts

An important argument for constructing dams in Africa is that Africa as a continent is particularly prone to spells of drought. In this regard dams may ensure the supply of water to those areas of need through their water storage capacity. To identify drought events within our hydrological regions we use the Inter-Governmental Panel on Climate Change (IPCC) data

³⁴ In terms of our hypothetical example shown in Figure this would constitute examining the effect of the dam on agricultural productivity in Quadrant II

³⁵ Including dams from downstream areas did not change the insignificance of the corresponding coefficient on this variable.

³⁶ Proceeding to include dams located even further the upstream ‘chain’ of areas similarly only produced insignificant effects – detailed results are available upon request.

climatic set, which provides monthly precipitation measures at the 0.5 degree level over the 20th century. In order to define periods of drought in our hydrological regions we first calculated the local standardized precipitation index (SPI), which has been argued to be particularly good at capturing the cumulative effect of reduced rainfall over time in a chosen locality³⁷, from the mean monthly precipitation values within our regions.³⁸ Following McKee et al. (1993) we define a drought event as starting when the SPI reaches an intensity of -1.0 or less and as ending once the index become positive again. As the summary statistics in Table 1 show, on average over our sample period the hydrological regions with cropland were in a period of drought 68 per cent of the time.³⁹

In the first column of Table 5 we include a dummy taking a value of one when a region is in drought and zero otherwise into our base specification. As would be expected, the negative and significant coefficient suggests that droughts reduce NPP. We next interacted our predicted dam variables with the drought dummy in order to see if dams can help alleviate the negative effects of water shortage during such periods. As can be seen from the second column, while there is no additional positive effect of upstream dams during times of drought, own area dams act to decrease cropland productivity when a drought occurs. This latter effect might potentially be due to the possibility that dams place particularly severe water restrictions in areas near the dam during a drought.

One problem of course in trying to disentangle the potential role of upstream dams in alleviating the negative effects of drought on cropland is that although precipitation varies regionally, there still tends to be a strong regional correlation of precipitation patterns.⁴⁰ In other words, if there is a simultaneous water shortage in areas upstream, water stored may already be

³⁷ One may want to note that the SPI is now used widely to identify droughts in Africa as, for instance, by the South African Weather Service.

³⁸ The calculation of the SPI is based on modeling the probability distribution of precipitation as derived from long term records by fitting these to a gamma distribution via maximum likelihood. An important component in this regard is the chosen time scale. Since we are interested cropland productivity and soil moisture conditions are known to respond to precipitation anomalies over a relatively short time period we use a 12 month scale. See <http://www.drought.unl.edu/whatis/indices.htm>.

³⁹ One should note that particularly the 1980s constituted a period in which continental wide rainfall was at its all time low compared to the rest of the 20th century.

⁴⁰ See, for instance, Barrios et al (2009).

low or needed for alternative uses, so there may be less water left for irrigation downstream. To investigate this we summed the number predicted dams in regions upstream only if these were not in a period of drought as defined above. The results of including this variable, UDAM(NDR), are given in the third column of Table 5. Accordingly, upstream dams located in regions that are not experiencing a drought induce a significant additional (to the average) 47 per cent effect on cropland productivity. In the last column of the table we then interacted UDAM(NDR) with our drought dummy of the region. As can be seen, while the coefficient on UDAM and the interaction of DAM and DROUGHT remain significant, the coefficient on the interaction between DROUGHT and UDAM(NDR) is also significant. This suggest that when a region experiences a drought, upstream dams can only help buffer the negative shock when they are themselves in a region that is not experiencing rainfall shortage.

E. Dam Purpose

Thus far we have considered dams as a homogenous group, assuming that all have at least some potential irrigation benefits. Nevertheless one might expect dams whose explicit major purpose (although possibly among several) is for irrigation to have a greater impact on cropland productivity. To explore this we divided all dams in our data into those which have irrigation listed as a major purpose and those that do not, and then re-ran the first stage separately for both groups.⁴¹ The results in the fourth column of Table 2 show that, as argued by Duflo and Pande (2007), high gradient rivers are not as suitable for irrigation dams in that the interaction term between our treaties variable and a the high river gradient category is insignificant, while the other two interaction terms remain significantly positive as in the full sample. One may also want to note that for this subsample treaties only have an effect on dam construction through the two gradient variables rather than on their own. For the remaining group, denoted as OTHER, one discovers that all three interaction terms, including the one with HIGH, remain statistically significant, denoting perhaps the variety of major purposes represented.

To generate predicted variables for our second stage regressions we then re-ran the two specifications for the IRRIG and OTHER samples with only the significant variables and time

⁴¹ This includes those that had no major purpose listed.

dummies. The results of including these in our base specification are given in columns (1) and (2) of Table 6, respectively. For both groups we find that dams located in the region itself have no net impact on cropland productivity. While both types of dams have a positive impact if they are located upstream dams for which at least one major purpose listed was for irrigation, have a greater quantitative effect.

In the final two columns of Table 6 we ran the specification in the last column of Table 6, but separately for the IRRIG and the OTHER groups. As can be seen, for dams which are listed as being intended for irrigation, there is a positive effect of upstream dams, where this effect is amplified when there is a drought but no such water shortage in the upstream area. In contrast, while similar effects exist for the OTHER group of dams, the average effect of upstream dams is considerably smaller than for the IRRIG sample. Moreover, the negative effect of being located in a region itself during a drought found for the total sample seems to be entirely driven by the OTHER group.

Section VI: Economic Significance

Thus far we have only addressed the statistical significance of dams in terms of affecting cropland productivity. Ideally, however, one would like to know what role they may have actually played in terms of aiding aggregate cropland production on the African continent over our sample period. To this end we first calculated mean cropland NPP at the country level over our sample period and then merged these data with actual annual cropland production figures taken from FAO's FAOSTAT database. We then regressed the log of cropland production as taken from FAOSTAT on country level log cropland NPP values, controlling additionally for country level fixed effects and time dummies. This produced a statistically significant positive coefficient of 0.197 on log cropland NPP⁴², confirming that NPP indeed plays an important role in determining the level of cropland production. Using this coefficient, the coefficient on UDAM from the fourth column in Table 3, and the mean number of actual upstream dams for all regions

⁴² Detailed results available from the authors.

within a country⁴³ then suggests that major dams increased annual cropland production by 2.88 per cent. If one then incorporates the effect of dams in times of drought by using the coefficients of the significant variables from the last column of Table 5 and their country level averages, one finds that the overall effect on cropland production was an average annual drop of 0.96 per cent. This is due to the fact that although upstream dams, particularly those located in regions not experiencing a drought, increase cropland productivity, in times of water shortage dams in a region's own area have a negative effect.

We also performed similar calculations across the two dam group types. Accordingly, using the estimated coefficients reveals that 'non-irrigation' dams have increased annual cropland production by 5.04 per cent, but that this effect changes to a negative 5.68 per cent effect once one takes into account periods of drought. In contrast, dams with a listed major purpose of irrigation increased cropland production by an average annual 6.11 per cent, where this effect rises in times of water shortage periods rises marginally to 6.12 per cent.

As a final exercise we use our results to roughly estimate what the potential increase in cropland production could be if further dams were constructed. To this end we calculated the potential increase in number of upstream dams by classifying our regions as low, medium, or high river gradient regions, depending on what gradient types was most prominently represented. For each category we then set the number of 'potential' dams as either the mean number of dams found in regions of that river gradient type or the actual number located there if the former was smaller than the latter. The potential increase is then just due the different between the potential number of dams upstream to a region relative to the actual observed number. Using our coefficients from above, our simple calculation reveals that annual cropland production could increase by a further 20 percentage points under a no-drought scenario and 17.8 percentage points otherwise.⁴⁴ This could be achieved by building a further 1382 in appropriate areas. One should note that the large potential increase is mostly due to the fact that dams can be upstream to several regions at the same time and hence can overcompensate for the negative effect they have in their own vicinity. If one considers only irrigation dams, and assume that only low and

⁴³ Whenever a region crossed national borders we assigned it to the country which contained most of its area.

⁴⁴ We assumed that the drought was evenly probable across all regions and equal to the mean over our 1981-2000 sample period.

medium gradient areas are suitable for these, then our simple rule suggests room for construction of a further 1143 dams. These would then increase, using the estimated coefficients, cropland production by a further 49 percentage points in times of drought, and 38 percentage points otherwise.

Section VII: Concluding Remarks

In this paper we provide the first comprehensive study of the distributional impact of large dams on agricultural productivity in Africa. Using a scientifically based disaggregation of the continent into local hydrological regions and satellite derived measures of cropland productivity we find that while dams benefit regions downstream, they have a negative impact within their own vicinity during periods of drought. Given that a large part of our sample period constituted relatively severe rainfall shortages, our results imply that overall dams had a small net negative impact of annual productivity losses of 0.96 per cent. Some simple back of the envelope calculations however indicated that further dam construction in Africa could provide considerable benefits.

Finally, one should note that our study only refers to one of a number of distributional aspects of large dams that have been highlighted, for instance, by the World Commission of Dams (2000). For example, we have not addressed such important issues as the displacement of people, the impact on fishing, energy provision, and river line erosion, amongst other negative impacts. Similar dams may provide additional benefits in terms of energy production and urban and industrial water supply. A comprehensive assessment of the total benefits and costs of dams in Africa clearly necessitates a credible quantitative analysis of all of these factors.

Table 1: Summary Statistics

Variable	Mean	St. Dev.
NPP_{AG}	1287	456
DAMS	0.21	0.99
UDAMS	0.07	0.38
POPDENS	21.1	82.9
DROUGHT	0.69	0.46

Notes: Summary statistics refer to all level 6 areas with at least some cropland.

Table 2: First Stage Regression Results

	(1)	(2)	(3)	(4)	(5)
T(5)	0.132** (0.047)	0.041** (0.014)		-0.005 (0.006)	0.046* (0.018)
T(5)*LOW		1.213* (0.564)		0.486* (0.284)	0.727** (0.281)
T(5)*MEDIUM		1.359** (0.512)		0.770** (0.257)	0.589* (0.314)
T(5)*HIGH		0.567* (0.235)		0.179 (0.165)	0.388* (0.184)
T(10)			0.022*** (0.006)		
T(10)*LOW			0.731** (0.337)		
T(10)*MEDIUM			0.751** (0.316)		
T(10)*HIGH			0.413*** (0.137)		
DAMS:	ALL	ALL	ALL	IRRIG.	OTHER
FE	Yes	Yes	Yes	Yes	Yes
# Obs.	142620	142620	142620	142620	142620
# Areas	7131	7131	7131	7131	7131
F-Test	8.05**	6.37**	30.97**	3.73**	44.20**
R²	0.014	0.017	0.017	0.010	0.011

Notes: (1) Time dummies included; (2) Driscoll and Kraay (1998) standard errors in parentheses; (3) ** and * are 1 and 5 per cent significance levels. (4) T(5) is given in 100s.

Table 3: Second Stage Regression Results

	(1)	(2)	(3)	(4)
DAM	0.006** (0.001)	0.004 (0.005)	-0.027 (0.183)	-0.040 (0.183)
UDAM	0.033** (0.002)	0.019** (0.006)	0.243** (0.040)	0.236** (0.041)
log(POPDENS_{t-1})				-0.010 (0.007)
DAMS:	Actual	Actual	Pred. T(5)	Pred. T(5)
FE:	No	Yes	Yes	Yes
# Obs.	72380	72380	72380	72380
# Areas	3619	3619	3619	3619
F-Test	186.21**	5.69**	22.75**	15.73**
R²	0.02	0.16	0.16	0.16

Notes: (1)Time dummies included; (2) Driscoll and Kray (1998) standard errors in parentheses; (3) ** and * are 1 and 5 per cent significance levels.

Table 4: Second Stage Regression Robustness Checks

	(1)	(2)	(3)	(4)
DAM	-0.123 (0.217)	0.028 (0.146)	-0.044 (0.182)	-0.068 (0.270)
UDAM	0.209** (0.044)	0.251** (0.047)	0.232** (0.039)	0.126 (0.077)
NDAM		-0.026 (0.017)		
U2DAM			0.035 (0.039)	
log(POPDENS_{t-1})	-0.010 (0.007)	-0.010 (0.007)	-0.010 (0.007)	-0.008 (0.004)
DAMS:	Pred. T(10)	Pred. T(5)	Pred. T(5)	Pred. T(5)
AREA:	Cropland	Cropland	Cropland	Non-Cropland
FE:	Yes	Yes	Yes	Yes
# Obs.	72380	72380	72380	68060
# Areas	3619	3619	3619	3403
F-Test	7.78**	13.59**	16.50**	11055.4**
R²	0.16	0.16	0.16	0.24

Notes: (1)Time dummies included; (2) Driscoll and Kraay (1998) standard errors in parentheses; (3) ** and * are 1 and 5 per cent significance levels.

Table 5: Second Stage Regression – Drought Effects

	(1)	(2)	(3)	(4)
DAM	0.006 (0.167)	0.191 (0.198)	-0.056 (0.175)	0.186 (0.194)
UDAM	0.218** (0.034)	0.212** (0.066)	0.248** (0.044)	0.205* (0.084)
DROUGHT	-0.033** (0.008)	-0.003 (0.015)		-0.005 (0.013)
DROUGHT*DAM		-0.228** (0.071)		-0.214** (0.068)
DROUGHT*UDAM		0.013 (0.041)		-0.021 (0.064)
UDAM(NDR)			0.117** (0.037)	-0.003 (0.058)
DROUGHT*UDAM(NDR)				0.123* (0.056)
log(POP DENS_{t-1})	-0.011 (0.007)	-0.011 (0.007)	-0.010 (0.007)	-0.011 (0.007)
DAMS:	Pred. T(5)	Pred. T(5)	Pred. T(5)	Pred. T(5)
AREA:	Cropland	Cropland	Cropland	Cropland
FE:	Yes	Yes	Yes	Yes
# Obs.	72380	72380	72380	72380
# Areas	3619	3619	3619	3619
F-Test	15.95**	30.27**	23.26**	37.47**
R²	0.16	0.16	0.16	0.24

Notes: (1) Time dummies included; (2) Driscoll and Kraay (1998) standard errors in parentheses; (3) ** and * are 1 and 5 per cent significance levels.

Table 6: Second Stage Regression – IRRIGATION and OTHER Dams

	(1)	(2)	(3)	(4)
DAM	-0.095 (0.393)	-0.074 (0.328)	0.278 (0.440)	0.403 (0.338)
UDAM	0.514** (0.085)	0.425** (0.082)	0.497** (0.162)	0.340* (0.161)
DROUGHT			-0.011 (0.014)	-0.003 (0.012)
DROUGHT*DAM			-0.345 (0.191)	-0.456** (0.100)
DROUGHT*UDAM			-0.061 (0.133)	-0.026 (0.121)
UDAM(NDR)			-0.015 (0.123)	0.001 (0.112)
DROUGHT*UDAM(NDR)			0.263* (0.117)	0.231* (0.108)
log(POPDENS_{t-1})	-0.010 (0.007)	-0.010 (0.007)	-0.011 (0.007)	-0.011 (0.007)
DAMS:	Pred. IRRIG.	Pred. OTHER	Pred. IRRIG.	Pred. OTHER
AREA:	Cropland	Cropland	Cropland	Cropland
FE:	Yes	Yes	Yes	Yes
# Obs.	72380	72380	72380	72380
# Areas	3619	3619	3619	3619
F-Test	17.38**	11.15**	35.83**	45.34**
R²	0.16	0.16	0.16	0.24

Notes: (1)Time dummies included; (2) Driscoll and Kraay (1998) standard errors in parentheses; (3) ** and * are 1 and 5 per cent significance levels.

Figure 1: Hypothetical Example – Hydrological Regions




Figure 2: Large Dams



Figure 3: Administrative Areas and Rivers: Southern Africa




Figure 4: Disaggregated Hydrological Regions

