Reply to “Comment on “Development Effects of Electrification: Evidence from the Topographic Placement of Hydropower Plants in Brazil”

Molly Lipscomb, Ahmed Mushfiq Mobarak, Tania Barham, Dimitri Szerman

August 2024

1 Introduction

Lipscomb et al. [2013] developed an instrumental variables strategy that took advantage of topographic requirements for hydropower generation in order to estimate the development effects of electrification. We update the results of that paper based on three considerations: (1) The engineering model used to derive the instrument is based on simulated annealing. This means that there is some instability in the estimated constructed grid based on the randomized allocation from alternative grids proposed by the model. We show that this can be eliminated, and the instrument improved, through averaging the predictions of the model. (2) There are different possible definitions of the Amazon region, depending on whether one uses official jurisdictional boundaries or the ecological definition of the Amazon biome. The paper used the legal Amazon definition in the “engineering model” for construction of the instrument and the Amazon together with the Pantanal (similar to the Amazon biome) in the controls applied in the 2SLS regression specification. We use the Amazon biome definition throughout and change the controls in the 2SLS to control for a quadratic in hydrographic cost factors interacted with decade. (3) Lee et al. [2022] is a recent applied econometrics paper that shows that an F-statistic of over 104 is necessary in order to avoid unbiased estimates of standard errors in the second stage of IV estimates. We correct the standard errors according to the tF procedure suggested by Lee et al. [2022]. As explained by Lee et al. [2023], this is most likely overly conservative. The revised standard errors are larger, and several of the results from the original paper are no longer statistically significant. A small error in variable cleaning is also updated.[1]

[1] We also discuss why an additional change suggested by Ankel-Peters et al. does not make sense.

The two stage least squares estimates of the impact of electrification on housing values and HDI are approximately half as large, but still statistically significant (housing is not statistically significant using the Lee correction). The finding that OLS estimates of the impact of electricity are biased downward relative to the two stage least squares estimates remains. Many of the mechanisms variables are no longer statistically significant following these updates, though it appears that the impacts occur primarily through employment, particularly in rural areas.

The Instrument To estimate the development effects of electricity, Lipscomb et al. [2013] construct an instrument using variation in the cost of producing hydropower based on local geographic characteristics, together with variation in the budget for hydropower across decades. We predict the placement of electricity in an "engineering model" based only on cost factors and the national budget of Brazil for electricity generation, leaving out all demand considerations. The process starts with a probit regression to estimate the relative importance of each geographic factor (Table 1 in Lipscomb et al. [2013]). This creates a "suitability index" that ranks every location in Brazil in terms of how sensible it is to place a hydropower dam at that location, if one were solely focusing on geographic factors. The underlying parameters for this index were water flow accumulation, average and maximum slope in the river, and an indicator for whether the point was in the Amazon. The Amazon coefficient was large and negative in this probit regression because building infrastructure in the rainforest is costly. Generation plants were therefore relatively unlikely to be placed there by the model. The time variation in the instrument comes from the fact that less and less ideal regions are available as non-electrified potential sites for new hydropower plants in later decades as the lowest cost locations are selected first. The most "suitable" locations that had not yet been electrified receive those plants together with simulated transmission lines and a distribution network which are allocated through simulated annealing.

Instrument Stability The "engineering model" which forms the basis of the instrument relies on simulated annealing which proposes alternative grids for the transmission lines in every decade, and compares the cost of the grid to the cost of alternative randomly generated grids. Cheaper grids are iteratively selected and the process repeats for 80,000 replications for each decade. This was by no means meant to be the best possible estimate of a model of Brazil's electricity grid, and the statistical advances in the field since 2013 have been significant–now machine learning tools offer straightforward alternatives for predicting placement of the grid. However, the simulated annealing algorithm provides us with a prediction of where the grid would be placed in each decade. Because simulated annealing relies on a randomized selection of alternative grids, there is natural variation in the final version of the grid that it achieves after 80,000 replications. There are

Table 1: Stability over different levels of averaging

	Instrument averaged over .... runs of the simulated annealing				
	5	25	50	100	500
Averaged Instrument	0.352***	0.386***	0.398***	0.396***	0.394***
	[0.070]	[0.073]	[0.074]	[0.073]	[0.074]

Notes: Shown are alternative levels of averaging of the instrument in runs of the preferred specification of the first stage (including the quadratic in hydrographic cost factors interacted with decade as controls as well as county (AMC) fixed effects, and decade fixed effects). All regressions have county size weights. Standard errors are clustered at the AMC level.

two main reasons for differences in the final grid between replication seeds: path-dependency–grid points categorized as electrified in 1960 can not be un-electrified in future decades, and natural random variation. Ankel-Peters et al. notes that this leads to instability in the first and second stage estimates. We run the model 500 times, and there is, as expected, some variation in the first stage coefficient estimates: the 1st percentile coefficient in our preferred specification with fixed effects by AMC and hydrographic cost factors interacted with decade controls is 0.12 while the 99th percentile is 0.32 with the F-stat varying from 14.1 in the first percentile to 23.9 in the 99th percentile.

The natural response to instability from simulated annealing is to average runs of the instrument. There are several benefits to averaging the estimated instruments: first, it improves the "signal" to "noise" ratio: the "signal" is guided by the hydrographic and budget factors from the data and is persistent between runs while the "noise" varies randomly. When the runs are averaged, the "signal" is retained across runs while the "noise" is canceled out. This effect can be seen as the first stage coefficient grows in the number of replications averaged for the creation of the instrument in table 1. In addition, the instrument created binary indicators for the placement of electricity for each grid point. Instead, the averaged instrument provides a continuous probability of the placement of electricity in each decade, thereby improving precision. As a result, the first stage is stronger: in our preferred specification the averaged instrument has a coefficient of 0.39 with an F-Statistic of 28.5. The instrument is stable once you have averaged at least 50 replications of the code as shown in table 1. Additional runs of the instrument have little impact on the estimates past 50. In this comment (and in a follow-on paper Szerman et al.), we use the averaged instrument across 500 runs of the Matlab code. We use this averaged version of the instrument in all specifications which follow.

The Specification Lipscomb et al. [2013] include county fixed-effects in their estimation in order to avoid relying on cross-sectional variation for identification which would otherwise introduce bias from differences between the counties. This isolates the identifying variation in the instrument to the cutoff between the last locations that received plants until the budget was exhausted for that decade, and the next most suitable areas that barely missed the threshold and had to wait for another decade for a plant. To mitigate concerns of common, spurious trends between the predicted grid expansion and development indicators, Lipscomb et al. [2013] directly control for any differential trends by adding geographic-specific time fixed-effects in their 2SLS estimation. Many potential specifications are possible with different geographic characteristics. In their preferred specification, Lipscomb et al. [2013] include interactions between the Amazon dummy and time-fixed effects in the first and second-stage regressions, which accounts for differential dynamics in Amazon counties.

Changes to the Specification regarding the Amazon Lipscomb et al. [2013] uses Brazil's Northern Amazon region defined by Brazil's official statistical agency, IBGE as a control in the estimation of the instrument. This area is primarily the area in which electricity transmission was managed by Eletronorte, and followed a separate planning process with more "isolated" generation plants than the rest of the Eletrobras network. Therefore, in Lipscomb et al. [2013] we control separately for the Legal Amazon in the creation of the instrument–both in the Probit regression and in the cost factors for transmission. In the 2SLS specification, the levels of development and trends in development indicators were different in the densely forested Amazon together with the Pantanal than the rest of Brazil, so we felt that this region merited a separate control; we added an Amazon and Pantanal indicator interacted with decade fixed effects to the 2SLS regressions. Both definitions of the Amazon are reasonable, but the ideal specification arguably would be consistent between the construction of the instrument and the 2SLS regressions despite the slightly different reasons for controls. As described by Peters, there are several definitions that could be used for the Amazon, and the results are similar across all of the specifications.

Table 2 reports the first-stage results from Lipscomb et al. [2013] together with the first stage results using the consistent Amazon biome definition (this is close to the Amazon and Pantanal) together with the averaged instrument. Each row of Table 2 corresponds to a different specification controlling for geographic-specific time trends, in the spirit of Table 10 in Lipscomb et al. [2013]. The Amazon biome definition is, conceptually, the closest in spirit to what Lipscomb et al. [2013] had in mind: a proxy for a geographic factor—dense forests—that increase construction costs of hydroelectric power plants.

The instrument becomes weak and therefore invalid, in Lipscomb et al. [2013]'s preferred specification (row 9). In fact, the table shows that the instrument becomes weak whenever the specification directly controls for Amazon-specific time fixed-effects (see rows 3,4,5,7,9, 11,12,14), and retains first-stage power when it does not (rows 1,2,6,8,10,13,15). The Amazon-specific trend controls simply absorb too much variation from the instrument.

Based on these considerations, we re-estimate the 2SLS specification using the averaged instrument and

Table 2: First Stage Results: Various Specifications Using the Amazon Biome Definition

Amazon Definition	AEJ Printed		Biome	
	(1)	(2)	(3)	(4)
Specification	First-Stage	F-Stat	First-Stage	F-Stat
1. Water Flow × budget	0.32*** (0.05)	45.8	0.47*** [0.07]	51.91
2. River Gradient × budget	0.32*** (0.05)	47.8	0.46*** [0.07]	49.68
3. Amazon dummy × budget	0.22*** (0.04)	25.4	0.13** [0.06]	5.33
4. Water Flow × budget, Amazon dummy × budget	0.22*** (0.04)	24.8	0.13** [0.06]	5.32
5. River Gradient × budget, Amazon dummy × budget	0.22*** (0.04)	25.3	0.13** [0.06]	5.22
6. River Gradient × budget, Water Flow × budget	0.32*** (0.05)	48.6	0.46*** [0.07]	49.79
7. River Gradient × budget, Water Flow × budget, Amazon dummy × budget	0.22*** (0.05)	24.4	0.13** [0.06]	5.09
8. Water Flow × year dummies	0.32*** (0.05)	45.6	0.47*** [0.07]	51.86
9. Amazon dummy × year dummies	0.22*** (0.04)	24.6	0.13** [0.06]	5.17
10. River Gradient × year dummies	0.22*** (0.05)	47.8	0.46*** [0.07]	49.23
11. Water Flow × year dummies, Amazon dummy × year dummies	0.22*** (0.05)	23.9	0.13** [0.06]	5.19
12. River Gradient × year dummies, Amazon dummy × year dummies	0.22*** (0.04)	24.9	0.13** [0.06]	4.92
13. Water Flow × year dummies, river Gradient × year dummies	0.32*** (0.05)	48.3	0.46*** [0.07]	49.36
14. River Gradient × year dummies, water Flow × year dummies, Amazon dummy × year dummies	0.22*** (0.05)	23.9	0.13** [0.06]	5.19
15. Quartic suitability rank × year dummies	0.24*** (0.04)	29.1	0.39*** [0.07]	28.50

Notes: All specifications include AMC60 fixed effects and decade fixed effects in addition to the listed controls. All regressions include county size weights. Standard errors are clustered at the AMC60 level.

with the following changes. First, we use the biome definition of the Amazon in all regressions and apply it consistently across all stage of analysis. The biome provides a control for both the Amazon and most of the Pantanal, and therefore is the sparsely inhabited region of Brazil where electrification would be prohibitively expensive and fail to follow the same cost considerations as the rest of Brazil. Second, we replace Lipscomb et al. [2013]'s preferred specification (shown in row 9 in Table 2) since the first-stage is now weak (coefficient is 0.13 with F-stat of 5.2), with a specification that controls directly for the levels and differential trends in the index of cost factors used in construction of the instrument, reported in row 15. This alternative specification controls directly for time variation related to the estimated geographic suitability of a location for electricity generation. The first stage F-statistic in our new preferred specification is 28.5, higher than the 24.6 reported in Lipscomb et al. [2013], and still above the commonly used cutoff of 10 suggested by Staiger and Stock [1997]. Because the F-statistic is below the 104 suggested by Lee et al. [2022] we also provide standard errors that have been corrected at the 95% level using their tF procedure (as shown in Lee et al. [2023], this is most likely overly conservative).

Data Cleaning Notes The UNDP's calculation of HDI changed between 1990 and 2000, and therefore the HDI variables needed to be adjusted to maintain consistency across the sample. 1990 values were provided for both the old and new methods of calculation of the HDI variables, so we adjusted the HDI variables by creating a ratio of the HDI for 1990 calculated under the old method to the HDI for 1990 calculated under the new method: $\frac{HDIold}{HDInew}$ and multiplying the new 2000 variable by the adjustment ratio. In cases in which the old variable was missing (as is the case for newly created municipios), we use the median value of the adjustment ratio for the AMC. Ankel-Peters et al. noted that while we did adjust the HDI component variables, we neglected to adjust the aggregate HDI variable in Lipscomb et al. [2013]. That is rectified in all regressions in this note.

Ankel-Peters et al. also suggests that there is an aggregation error in our data cleaning code. After checking our code, we find that Ankel-Peters et al. is in error. The comment suggests that in municipios created post-1990 the HDI and other outcome variables are missing in 2000 while the population data is not missing. This is incorrect. Population data is missing when outcome variables are missing, so there is no mismatch as suggested by Ankel-Peters et al.. The example municipios (Satubinha, Porto Maua and Novo Machado) provided in their appendix table A3 have original raw data available for the outcome variables from IPEA. They most likely suppressed the data for these municipios themselves inadvertently when correcting for the variable calculation change described above.[2]

[2]The adjustment ratio is not missing for new municipios in our code because we replaced missing adjustment ratios with the AMC median ratio under the assumption that the best approximation of the adjustment factor for new muncipios would be the adjustment factor for the area that they were carved out of.

Results Compared to Lipscomb et al, 2013 In table 3 we compare the results from the new preferred specification using the quartic of the suitability rank interacted with decade fixed effects as controls to those published in Lipscomb et al. [2013]. We find the magnitude of the coefficients is smaller and statistical precision is lost in several cases, but the direction of impacts remains the same in nearly all cases. Using the Lee et al. [2022] tF standard errors correction procedure, further statistical precision is lost.

Lipscomb et al. [2013] focused on two central dependent variables: housing values and the human development index (HDI). Rows 1 and 2 of Table 3 shows that the estimated effect of a zero to full increase in electricity infrastructure in a county leads to housing values increasing by $+378$ reais (the previous estimate was 881 reais). For the human development index, we estimate a 6 percentage point increase for new access to the electricity network (from 11 percentage points estimated in Lipscomb et al. [2013]).

The impact of electricity access on HDI components and other poverty measures was shown in table 11 of Lipscomb et al. [2013], they are now shown in rows 3-8 of table 3. The estimated impact of electrification on HDI components is positive but smaller in magnitude and not statistically significant in the new specification. Also shown are effects on infant mortality, income per capita and poverty; we find large effects on infant mortality, but effects on income per capita become very small and imprecisely measured negative. Effects on poverty become smaller and more imprecise (though the point estimate continues to be negative on poverty head-count).

The employment measures of economically active and formal employment remain positive, but the point estimates are smaller and economically active is no longer statistically significant. Effects on urban employment are also smaller, positive, and not statistically significant, while effects on rural employment are about half the size originally estimated, but still statistically significant. Effects on less than 4 years in school and number of years in school are no longer statistically significant and they have the opposite of the expected sign. Human capital accumulation has a smaller coefficient and is still not statistically significant.

Rows 14-16 of Table 3 show the effects on a few outcomes that Lipscomb et al. [2013] was trying to rule out as mechanisms. Effects on life expectancy and population density were statistically insignificant in both the original paper and here. Effects on urbanization also becomes statistically insignificant. Lipscomb et al. [2013] ran these regressions to examine whether the documented effects of electrification on development outcomes were the result of simple population movements, and we can now say with a bit more confidence that they were not.

Key Takeaways Estimating the causal impact of infrastructure services is notoriously difficult because access naturally expands to areas where demand is highest–the central contribution of Lipscomb et al. [2013]

Table 3: Estimates based on new specification

	AEJ Printed Results		New Specification Results	
	(1) OLS	(2) IV	(3) OLS	(4) IV
Results from Table 6 in LMB				
Housing Values	0.80***	8.81***	0.67**	3.78**
	[0.27]	[3.03]	[0.30]	[1.73]
				(2.11)
Results from Table 7 in LMB				
Human Development Index	0.01	0.11***	0.00	0.06***
	[0.01]	[0.04]	[0.01]	[0.02]
				(0.02)**
Results from Table 11 in LMB				
HDI: Education	0.03***	0.19***	0.02**	0.04
	[0.01]	[0.06]	[0.01]	[0.04]
				(0.05)
HDI: Longevity	0	−0.01	−0.00	0.00
	[0.01]	[0.05]	[0.01]	[0.03]
				(0.04)
HDI: Income	−0.03*	0.45**	−0.08***	0.02
	[0.02]	[0.15]	[0.02]	[0.13]
				(0.16)
Infant Mortality	−7.99***	−11.97	−10.27***	−71.30***
	[2.42]	[18.08]	[2.32]	[15.03]
				(18.34)**
Gross Income PC	−0.01	0.11**	−0.02***	−0.04
	[0.01]	[0.05]	[0.01]	[0.04]
				(0.05)
Poverty	−0.76	−42.17***	2.92*	−11.05
	[1.39]	[13.84]	[1.77]	[10.72]
				(13.08)
Results from Table 12 in LMB				
Economically Active	0.011*	0.173***	0.00	0.05
	[0.01]	[0.05]	[0.01]	[0.03]
				(0.04)
Formal Employment	0.010*	0.184***	0.00	0.07**
	[0.01]	[0.05]	[0.01]	[0.03]
				(0.04)
Urban Employment	0.01*	0.18***	0.00	0.04
	[0.00]	[0.05]	[0.01]	[0.04]
				(0.05)
Rural Employment	0.01	0.17***	0.01	0.08**
	[0.01]	[0.05]	[0.01]	[0.03]
				(0.4)**
Results from Table 13 in LMB				
Less than four years education	−0.36	−21.25***	2.37**	2.37
	[0.90]	[7.75]	[1.11]	[5.88]
				(7.17)
Years in School	0.06	2.02***	−0.16*	−0.11
	[0.08]	[0.67]	[0.10]	[0.44]
				(0.54)
Human Capital	2.09	11.54	−1.05	0.21
	[0.41]	[7.30]	[0.68]	[3.07]
				(3.75)
Results from Table 14 in LMB				
Life Expectancy	−0.44	−1.03	−1.01***	−2.69
	[0.32]	[2.39]	[0.31]	[2.07]
				(2.52)
Population Density	−1.11	−23.62	5.95***	8.42
	[3.74]	[19.20]	[2.19]	[9.47]
				(11.64)
Urban percent of pop	0.01	0.24**	−0.02	−0.01
	[0.01]	[0.11]	[0.02]	[0.11]
				(0.13)

Notes: Standard errors clustered by county reported in brackets, 0.05 tF adjusted standard errors corrected using the Lee et al. [2022] tF procedure reported in parentheses (adjusted by 1.22 based on their table 3A). All regressions have county size weights, county fixed effects decade fixed effects and hydrohat quadratic*decade fixed effects. The changes to the specification are (1) both the instrument and the 2SLS use Amazon Biome as the definition of the Amazon in the control, and (2) the instrument is based on averaging the results of 500 runs of the Matlab simulated annealing. The first stage F-Statistic is 28.5.

was to show that one can harness the differences in costs together with the budget for electricity over time to create an instrument for electricity access with panel variation. This provides two key improvements in estimation of the impact of electrification over OLS estimates: (1) the panel variation in the instrument allows us to include fixed effects to net out time-invariant factors unrelated to electricity that make certain areas more likely to both receive electricity and to have higher income and electricity demand. (2) the remaining identifying variation is from the difference between municipios which receive electricity in a given decade and those who are marginally higher cost than the decade cutoff based on geographic factors. It is unlikely that the outcome variables are correlated with this difference in municipios around the cutoff for electricity provision in a decade for reasons *other* than increased access to electricity, so the exclusion restriction is satisfied. Even with several key changes, we show that the instrument based on the evolution of optimal areas for building additional hydro-capacity is strong. The impact of electricity access is positive and important in magnitude though smaller than originally estimated in Lipscomb et al. [2013]. The finding that over the period 1960-2000 in Brazil IV estimates of the impact of electricity are more positive than OLS estimates, suggesting a downward bias in the OLS estimates, is retained.

References

Jörg Ankel-Peters, Gunther Bensch, and Colin Vance. Comment on "Development Effects of Electrification: Evidence from the Topographic Placement of Hydropower Plants in Brazil. *American Economic Journal: Applied Economics*.

David S. Lee, Justin McCrary, Marcelo J. Moreira, and Jack Porter. Valid t-ratio inference for iv. *American Economic Review*, 112(10):3260–90, October 2022. doi: 10.1257/aer.20211063. URL https://www.aeaweb.org/articles?id=10.1257/aer.20211063.

David S Lee, Justin McCrary, Marcelo J Moreira, Jack R Porter, and Luther Yap. What to do when you can't use'1.96'confidence intervals for iv. Technical report, National Bureau of Economic Research, 2023.

Molly Lipscomb, A. Mushfiq Mobarak, and Tania Barham. Development Effects of Electrification: Evidence from the Topographic Placement of Hydropower Plants in Brazil. *American Economic Journal: Applied Economics*, 5(2):200–231, 2013. doi: 10.1257/app.5.2.200.

Douglas Staiger and James H. Stock. Instrumental Variables Regression with Weak Instruments. *Econometrica*, 65(3):557–586, 1997. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/2171753.

Dimitri Szerman, Juliano Assunçao, Molly Lipscomb, and Mushfiq Mobarak. Electricity, agricultural productivity, and deforestation: Evidence from brazil.