UNITED NATIONS UNIVERSITY
UNU-WIDER

WIDER Working Paper 2022/23

Impact of teacher content knowledge on student achievement in a low-income country

Anna Holvio*

March 2022

United Nations University World Institute for Development Economics Research
wider.unu.edu

Abstract: This paper estimates the causal impact of teacher content knowledge on student achievement in Mozambique, a low-income country where a large share of fourth-graders fail to meet the minimum requirements of literacy and numeracy. I use nationally representative data from the Service Delivery Indicator survey, and exploit within-student across-subject variation in a sample of students taught by the same teacher in maths and Portuguese, thus circumventing bias caused by unobserved student and teacher heterogeneity. I find that, on average, teacher content knowledge does not have an impact on student achievement. However, the impact varies significantly by student's first language, urban or rural location of the school, and the match of students and teachers in their knowledge.

Key words: teacher content knowledge, student achievement, teacher quality, Mozambique

JEL classification: C23, I21, O15

Acknowledgements: This paper is part of a research project titled 'Schooling, Learning and the Political Economy of Education' at Aalto University School of Business and Helsinki Graduate School of Economics, directed by Professor of Practice Ritva Reinikka. Financial support from the Ministry for Foreign Affairs of Finland is gratefully acknowledged. I would like to thank Ritva Reinikka and Jukka Pirttilä for their guidance, and Ezequiel Molina and Pekka Pere for their helpful comments. The paper is based on my Master's thesis, 'Teaching and Learning in Mozambican Primary Schools: How Does Teacher Content Knowledge Affect Student Achievement?' at University of Helsinki, available at http://urn.fi/URN:NBN:fi:hulib-202111023967.

* Aalto University School of Business, Espoo, Finland; corresponding author: anna.holvio@aalto.fi

This study has been prepared within the UNU-WIDER project Academic excellence.

Information and requests: publications@wider.unu.edu

ISSN 1798-7237 ISBN 978-92-9267-154-9

https://doi.org/10.35188/UNU-WIDER/2022/154-9

Typescript prepared by Gary Smith.

United Nations University World Institute for Development Economics Research provides economic analysis and policy advice with the aim of promoting sustainable and equitable development. The Institute began operations in 1985 in Helsinki, Finland, as the first research and training centre of the United Nations University. Today it is a unique blend of think tank, research institute, and UN agency—providing a range of services from policy advice to governments as well as freely available original research.

The Institute is funded through income from an endowment fund with additional contributions to its work programme from Finland, Sweden, and the United Kingdom as well as earmarked contributions for specific projects from a variety of donors.

Katajanokanlaituri 6 B, 00160 Helsinki, Finland

1 Introduction

Primary school enrolment has grown to be nearly universal on a global scale, even in many low-income countries such as Mozambique. By contrast, learning results have not kept up with the rapidly expanding school systems. In many cases, these results have stagnated at a level that is not high enough to meet even the minimum requirements of literacy or numeracy for a large share of the students. Children progressing into higher grades without having a solid foundation is further hampering learning outcomes. A situation in which more children are in school than ever, but are learning very little, has been dubbed the 'learning crisis' (UNESCO 2013). In Mozambique, fewer than one in every eight fourth-grade students in my data could correctly read a single sentence. Likewise, only half could perform single-digit additions. Solely scaling up what has already been done will not be enough, but rather calls for actions that put learning at the centre.

A considerable amount of literature has been published on the impact of teacher quality on learning, and it is well established that good teachers can lead to significant gains. This is not true only for short-term learning results, but also for longer-term labour market outcomes (Chetty et al. 2014b). However, there is little consensus on which factors actually constitute teacher quality. Among all observable characteristics, earlier research has most consistently found teacher content knowledge to have a positive impact on student achievement. From this perspective, it is possible that raising teachers' knowledge by improving their training, for instance, could lead to better learning results. Still, existing studies have provided mixed evidence and found substantial variation across countries. This seems natural, because the contextual differences in education systems can be large.

Because of these contextual differences, what has been found to improve learning in one country cannot be copied as such to another. Pritchett et al. (2013) argued that such policies often fail because of 'isomorphic mimicry'—they only imitate the successes of others to make their own institutions appear legitimate, without having similar functionality. Mozambique is among the poorest countries in the world, and its economic development has not been inclusive, but has disproportionately benefited certain populations, such as higher-income households, urban areas, and the southern regions (Baez Ramirez et al. 2018). The education system itself was severely damaged by the long civil war, and had to be rebuilt starting from the 1990s. In general, both its teacher knowledge and student achievement are at a particularly low level (Molina and Martin 2015).

The objective of this study is to determine the causal impact of teacher content knowledge on student achievement in Mozambican primary schools. I use nationally representative data from a Service Delivery Indicator (SDI) survey conducted in 2014, which included assessments of fourth-grade students and their teachers in maths and Portuguese. The empirical analysis exploits within-student across-subject variation. This allows introducing into the model not only student fixed effects, but also teacher fixed effects, because all students in the sample were taught by the same teacher in both subjects, therefore strengthening the causal identification.

I find that, on average, teacher content knowledge does not have an impact on student achievement. However, I find evidence that the impact depends on the first language of students, the location of the school (urban or rural), and whether the student and his/her teacher have an above-the-median knowledge. Raising teacher content knowledge by one standard deviation (SD) was associated with an increase of 0.14 SD for students whose first language was Portuguese, and with an increase of 0.13 SD for students in urban schools. In addition, raising the content knowledge of a teacher whose knowledge was above the median by 1 SD was associated with an increase of 0.12 SD in student achievement when the student was also above the median.

This paper has been organized in the following way. Section 2 reviews the existing literature on teacher quality and teacher knowledge, and their impact on student achievement. Section 3 gives an overview of the Mozambican education system. Section 4 presents the SDI data, with a focus on the student and teacher assessments. Section 5 presents the empirical strategy. Section 6 discusses the results. Finally, Section 7 concludes.

2 Teachers and learning: existing evidence

2.1 Teacher quality

Teacher quality matters a great deal for student achievement, and there is an extensive literature on the topic. Still, much of what actually constitutes teacher quality is accounted for in unobservables. Perhaps this is why teacher quality is often understood in terms of outputs rather than inputs. Hanushek (2002) simply considers good teachers as the ones who get large gains in student achievement, and bad teachers as those who do not. It is typical to use student achievement as the output because it can be measured through tests, and it gives a rather immediate impression on how the students are doing.

Much of the economic research on education is built upon the education production function, which relates various inputs in education to the obtained level of student achievement. Common inputs include characteristics of students, households, teachers, classrooms, and schools, while the outcome is most often a test score used as a measure of student achievement. Strøm and Falch (2020) use the following general formula for the education production function:

$$y_{it} = f\big(X_i(t), S_i(t), T_i(t), \alpha_{i0}, e_{it}\big) \tag{1}$$

where i refers to the student and t to time (or grade). This function models the outcome y_{it} as a function of different relevant inputs up to time t. $X_i(t)$ represents individual and family inputs, $S_i(t)$ represents school inputs, and $T_i(t)$ represents teacher quality. The innate ability of the student is expressed with α_{i0}, and the idiosyncratic error with e_{it}.

A key concern for research on teacher quality is the non-random assignment of students to teachers based on some background characteristics. Non-random assignment makes it difficult to derive causal estimates that are unaffected by the pre-existing differences between students. Rothstein (2009) describes possible reasons for the non-random assignment both between and within schools. Parents who consider education important may choose a particular school because it is known to have better teachers, or school directors might assign students to specific teachers either to gather students of similar ability in the same classrooms, or to spread students of different abilities through all classrooms.

A widely used approach for obtaining estimates on the impact of overall teacher quality is the value-added approach. This uses longitudinal data that matches the test scores of individual students to their teachers at different points in time, and decomposes these test scores into components that represent teacher quality and student heterogeneity. Preferably, students improve when moving to higher grades, and teacher value-added is understood as the part of this improvement that is due to the teacher. Generally, effective teachers have a high value-added, and are therefore good at improving their students' learning. A common critique is whether these models are able to identify the causal effect in the case of non-random assignment, but recent studies indicate that the value-added approach can yield unbiased estimates when controlling for students' prior test scores (Chetty et al. 2014a).

Findings from value-added studies confirm that there is large variation in teacher value-added, indicating that effective teachers can significantly improve their students' learning results. Bau and Das (2020) used data from Pakistani students in grades 3–5, and found that increasing teacher quality by 1 SD increased students' test scores by 0.15 SD. The impact was larger for maths (0.21 SD) and English (0.17 SD), and smaller for Urdu (0.06 SD). Hanushek and Rivkin (2010) summarize the results of ten studies from the United States, and found estimates that ranged from 0.08 to 0.26 SD in reading (0.14 on average), and from 0.11 to 0.36 SD in maths (0.19 on average).

2.2 Teacher knowledge

Among the observable characteristics of teachers, teacher knowledge has most consistently been found to have a positive association with student achievement (Coenen et al. 2018; Hanushek 1997; Wayne and Youngs 2003). It is not obvious, though, how best to conceptualize teacher knowledge, as teachers need a vast array of knowledge in their profession. A widely cited categorization of teacher knowledge is that of Shulman (1987), who argued that there are, at minimum, seven categories for teacher knowledge that promote learning. The most commonly studied category is content knowledge, also referred to as subject knowledge, which Shulman describes as referring to the knowledge, understanding, and skills that students should learn.[1]

Studies often estimate the impact of teacher content knowledge directly by including it in the education production function, most often in the form of a test score. Because the data are mostly observational, the concern again is whether the studies are able to account for the non-random assignment of students to teachers. One way to overcome this problem is to identify a causal effect by exploiting within-student across-subject variation. When students are tested in multiple subjects, it is possible to rule out any characteristics of students that do not vary across subjects, such as socio-economic background.

While there is limited literature on teacher content knowledge and its impact on student achievement in developing countries, a number of reviews summarize research on different teacher characteristics, mostly from the United States. Coenen et al. (2018) came to the conclusion that both overall teacher test scores and test scores in maths affected student achievement in maths, but teacher language scores did not have a similar effect. The impact of increasing teacher content knowledge by 1 SD ranged from –0.033 to 0.087 SD change in student achievement, and was positive in 9 out of 11 cases. Also, Wayne and Youngs (2003) argued that improvements in teacher content knowledge lead to increases in student achievement, but pointed out that studies gave differing results based on the controls they used. Hanushek (1997) showed that out of 41 estimates, 47 per cent were statistically significant, out of which 37 per cent were positive and 10 per cent were negative. Recently, Hanushek et al. (2019) found that the impact of teacher content knowledge was 0.105 SD in a sample of 31 OECD countries.

More recently, research on the role of teacher content knowledge for student achievement in low- and middle-income countries has started to emerge. To provide a better reference point for this study, I review eight studies in detail, all concentrating on primary schools. Four of the studies were peer-reviewed (Bietenbeck et al. 2018; Marcenaro-Gutierrez and Lopez-Agudo 2020; Metzler and Woessmann 2012; Zakharov et al. 2016), while the others were working papers or background papers (Bold et al. 2019, 2017; Filmer et al. 2015; Shepherd 2015). Most of these studies utilized either data provided by the Southern and Eastern Africa Consortium for Monitoring Educational Quality (SACMEQ),[2] or SDI data.[3]. An exception was Metzler and Woessmann (2012), who used data from Peruvian national evaluations. Most of these studies aggregated data from various countries, while Shepherd (2015) fo-

[1] Shulman (1987) names general pedagogical knowledge, curriculum knowledge, pedagogical content knowledge, knowledge of learners and their characteristics, knowledge of educational contexts, and knowledge of educational ends, purposes, and values as the other six categories.

[2] Bietenbeck et al. (2018); Marcenaro-Gutierrez and Lopez-Agudo (2020); Shepherd (2015); Zakharov et al. (2016).

[3] Bold et al. (2019, 2017); Filmer et al. (2015).

cused solely on South Africa, and Metzler and Woessmann (2012) on Peru. The sample sizes of these studies were rather large, ranging from 6,996 to 74,708 students.

All of these papers exploited within-student across-subject variation to identify the causal effects. The exact methodological approach was divided so that the models of Bietenbeck et al. (2018), Marcenaro-Gutierrez and Lopez-Agudo (2020), Bold et al. (2019, 2017), and Zakharov et al. (2016) included student fixed effects, while Shepherd (2015), Metzler and Woessmann (2012), and Filmer et al. (2015) used correlated random effects. More specifically, Bold et al. (2019) divided the impact of teacher content knowledge between the current and all previous teachers, while the others only took the current teacher into account. Furthermore, some of these studies used a sample of students taught by the same teacher in both subjects to rule out the impact of unobserved teacher heterogeneity.

The results indicate that teacher content knowledge generally had a positive effect on student achievement in low- and middle-income countries. Yet, the estimates were rather small, and in some cases statistically insignificant. Out of the significant ones, Bietenbeck et al. (2018) estimated that a 1 SD increase in teacher content knowledge increased student achievement by 0.025 SD, while the estimate was 0.02 SD for Bold et al. (2017). Moreover, Metzler and Woessmann (2012) found an impact of 0.087 SD for maths, but a statistically insignificant 0.022 SD for reading. Filmer et al. (2015) had largely similar results, with an estimate of 0.05 SD for maths, and a statistically insignificant 0.02 SD for reading. For others, the coefficients ranged from –0.017 to 0.034 SD, but were statistically insignificant.

An interesting finding that many of the papers highlighted was that there was heterogeneity in the effect in terms of country, material resources and wealth of schools, students' socio-economic background, and different levels of student and teacher knowledge. Bietenbeck et al. (2018) related the heterogeneity between countries to their level of economic development, and argued that better (material) resources may be a driver for the larger impact experienced by countries with a higher per capita GDP. In fact, one can argue that the estimates generally tend to be larger the higher the income level in the country.

Another common finding seems to be that those who are already in a more advantageous position benefit more from teacher content knowledge. Shepherd (2015) found that the effect was larger for the wealthiest 20 per cent of schools, and for students from households with an above-average socio-economic status. Additionally, Metzler and Woessmann (2012) compared the estimates for different sub-samples where students and teachers from above and below the median were matched. The largest effect was found in maths when students who were above the median were taught by a teacher also above the median. This effect was twice as large as it was for students below the median when taught by a teacher above the median.

3 Education in Mozambique

3.1 Education system

Mozambique revised its National Education System Law in 2018. The law includes making education compulsory through the first nine grades, and reorganizing the structure of primary education from two levels (EP1, which covered grades 1–5, and EP2, which covered grades 6 and 7) into one that covers grades 1–6 (MINEDH 2020b; UNESCO and MINEDH 2019). These changes are being implemented gradually, and will be completed by 2023. This is not the only reform that has happened in the education system; the system itself could be characterized as undergoing constant change. Reinikka and Napaua (2019) interviewed local academics, who reported that a total of 23 different reforms had been introduced over the years since the end of the civil war in the early 1990s.

Most children start their educational path in primary school at age six, as preschool enrolment is still rare. As mentioned, primary school covers grades 1–6, after which students progress to secondary school for grades 7–9. After this compulsory education, students can continue into the final grades of secondary school, or enrol in vocational education. Higher education is available for those who have completed secondary school. The majority of children attend public schools, where there are no tuition fees, but families have to pay some expenses, such as for school uniforms. The role of private education at the primary level is small, and it is estimated that only around 3 per cent of primary school students are enrolled in private schools (MINEDH 2020a).

The high number of students has put pressure on schools to operate in shifts, but official student–teacher ratios remain high. Based on administrative records, there were 55.3 students per teacher in 2018, which is well above the sub-Saharan African average of 37.4 (World Bank 2021a). However, this figure is inflated when compared to actual observations because of high student absence rates. Molina and Martin (2015) reported a student–teacher ratio of 21.4 based on direct observation of classrooms. This is only half of the average ratio of 42 students per teacher in public schools in all African countries in the SDI data.

The language of instruction in primary schools is either Portuguese or Portuguese and a local language in the case of bilingual education. Bilingual education has gradually become more widely offered since the 1990s. Still, the share of students participating in these programmes remains limited. Based on stakeholder interviews, Reinikka and Napaua (2019) suggested that the main reasons for this have been the lack of funding, expertise, and leadership from the government, and the complexity of implementing bilingual education for dozens of local languages.

Teacher training in Mozambique has also experienced frequent reforms. Today, over 97 per cent of more than 120,000 primary school teachers in Mozambique have received some form of teacher training (World Bank 2021b,c). This training, however, can range from one to three or possibly more years, and the prior years of general education from 4 to 12, which makes the teaching workforce very heterogeneous. MINEDH (2017) lists a total of 17 different models of teacher training that have been applied since independence, while Beutel (2011) mentions that there have been 21 different models. Whatever the exact number, it is clear that the pace of these reforms is rapid. Even today, primary school teachers are trained in two programmes that differ in their durations.

In 2007, the country adopted a so-called 10+1 model, replacing the earlier 10+2 model. In the 10+1 model, teachers need to have completed ten years of general education, followed by one year of teacher training. The model was originally introduced as a short-term solution to teacher shortages, but since then became the standard. The 10+1 model has been widely criticized for its short duration that does not adequately prepare the prospective teachers for their profession (Beutel 2011; Reinikka and Napaua 2019; UNESCO and MINEDH 2019). It should be kept in mind that teacher candidates arrive from a system that produces poor learning results, and one year is not enough time to both account for these flaws and prepare the teacher candidates with the needed pedagogical knowledge and skills. In response to these criticisms, a longer, alternative 10+3 model was piloted in 2012.[4] Still, a majority of new primary teachers enter the profession with only one year of pre-service training (MINEDH 2020b).

In addition to its short duration, teacher training has also received criticism for its inefficient selection process, and a curriculum that does not sufficiently reflect the actual process of teaching. Reinikka and Napaua (2019) interviewed local actors during field visits, and pointed out that the current entry

[4] It seems that more recently there have been plans to both raise the requirements of general education and to increase the standard duration of teacher training. The latest strategic plan for education mentions that in the coming years teacher candidates would enter training after 12 years of general education, and the training would last three years (12+3), replacing the existing models (MINEDH 2020b).

examinations do not effectively detect the best applicants. This is unfortunate in a situation in which teacher knowledge is generally at a weak level, and the training itself is not enough to correct for this. Beutel (2011) interviewed both teacher trainers and students of teacher training institutions about their views on the training, and while the trainers called for both subject learning and practical skills, students stressed the role of practical skills that they were not learning in their training.

3.2 Learning results

Learning results from Mozambique show that learning is at a poor level compared to other countries, and particularly compared to what is expected by the curriculum. Bold et al. (2019) calculated that the knowledge of a median fourth-grade student was equivalent to 0 years of schooling—that is, if a student is expected to have completed the curriculum corresponding to each grade by its end. In general, there are few data available to make comparisons and put the results in context. Available sources that allow for comparisons over time and/or across countries include two rounds of SDI surveys in 2014 and 2018 for fourth-graders, and two studies conducted by the SACMEQ in 2000 and 2007 for sixth-graders.[5] What should be considered, though, is that the SACMEQ results are over a decade old, and many changes have since taken place.

Table 1 summarizes the student learning results of various countries in the SDI data. Looking at the average, it becomes evident that learning is at a low level in all of the surveyed countries, but especially in Mozambique. Mozambique had the lowest language score out of all the countries, and the second lowest maths score. Only Niger compared to Mozambique when it came to student test scores. Similarly, out of 15 countries the SACMEQ assessment from 2007 placed Mozambique 12th in reading and 10th in maths (Magaia et al. 2011). What is striking is that Mozambique was the only country whose reading results worsened, and the only country besides Uganda whose maths results worsened.

Table 1: Comparison of student learning results from SDI surveys across different countries: test scores (out of 100)

	Mozambique 2014	Kenya 2012	Madagascar 2016	Niger 2015	Nigeria 2013	Tanzania 2014	Togo 2013	Uganda 2013	SDI average
Language	18.7	75.4	44.5	21.7	31.4	36.5	45.5	47.1	49.5
Maths	25.1	59.0	56.8	11.5	31,9	58.2	44.6	43.4	47.3

Note: the results are for all schools, except for Mozambique and Tanzania, where only public schools were included. The sample for Nigeria consists of only four states (Anambara, Bauchi, Ekiti, and Niger).

Source: author's calculations based on Wane and Rakotoarivony (2017) and Marouma et al. (2017).

Data from the second SDI survey from 2018 are not yet publicly available, but a report by Bassi et al. (2019) provides a summary of how the indicators have evolved over the four years (Table 2). These results suggest that there has been slight improvement in learning, particularly in the Portuguese scores, which increased by 12.5 points, on average. However, the improvements were mostly driven by the southern region and urban areas.

Table 2: Comparison of student learning results for Mozambique from 2014 and 2018: test scores (out of 100)

	Total		South		Center		North		Urban		Rural	
	2014	2018	2014	2018	2014	2018	2014	2018	2014	2018	2014	2018
Portuguese	18.7	31.2	38.0	54.8	17.5	22.6	11.2	19.9	19.2	44.8	18.6	29.8
Maths	25.1	31.4	28.6	37.8	25.2	29.6	23.5	27.4	25.1	37.7	25.1	30.7

Source: author's calculations based on Bassi et al. (2019) and Molina and Martin (2015).

[5] The SACMEQ is an association of 16 Ministries of Education that carries out studies to assess the quality of education in its member countries.

Large disparities in learning are characteristic for Mozambique. Both SDI and SACMEQ results show differences in learning by region, urban and rural location, gender, and socio-economic background. The 2018 SDI data reveal that a student in the south scored more than double the number of points in Portuguese (54.8) compared to a student in the north (19.9), and urban students had significantly higher scores. The gender gap also showed geographic deviations; in the south it was practically non-existent, whereas in the centre and the north the gender gap favoured boys. The 2007 SACMEQ data, on the other hand, showed that students with a higher socio-economic status (top 25 per cent) had about 15 per cent higher reading scores than students with a lower socio-economic status (bottom 25 per cent).

4 Data and measurement

This study uses data from the first SDI survey in Mozambique from 2014,[6] which are anonymized and publicly available at the World Bank Microdata Library.[7] The SDI are a joint initiative between the World Bank and the African Economic Research Consortium to measure the performance and quality of service delivery in education and healthcare across Africa. The data are designed to give answers to three questions: *What do teachers know? What do teachers do? What do teachers have to work with*? Data from schools, teachers, and students were collected by local partners through questionnaires, classroom observations, and assessments of teachers and fourth-grade students that were carried out in-person by the enumerators.

The surveys were implemented using a multistage, cluster sampling approach that produces nationally representative estimates, taking into account the location and size of different schools. The sampling frame for the survey was all public primary schools with at least one fourth-grade classroom. In total, the data set contained test results for 1,761 students and 673 teachers from 203 schools across Mozambique, all of which were mixed-gender. As I am studying the effect of teacher subject knowledge on student achievement, the sample needed to be narrowed down to observations where both the student and the teacher were tested. This sample consisted of 1,505 students and 171 teachers from 171 schools. In this same-teacher (ST) sample, all the students were taught by the same teacher in maths and Portuguese. I also focused on a sub-sample of students in schools where there was only one fourth-grade classroom to rule out the possibility of non-random assignment within a school. This same-teacher one-classroom (STOC) sample consisted of 975 students and 113 teachers from 113 schools.

4.1 School, teacher, and student characteristics

Table 3 presents characteristics of schools, teachers, and students. Looking at the schools, most of them were rural. This is in line with the low urbanization levels of Mozambique (37 per cent; World Bank 2021d). The majority of schools were modest in their infrastructure. Only 5 per cent of classrooms had a working electricity connection. Teachers were predominantly female and young. A median teacher had six years of teaching experience, but one-fifth of teachers had only up to one year of previous experience. Over 95 per cent of the teachers had received some teacher training, and over half reported having completed 12 years of general education.

[6] A second survey was conducted in Mozambique in 2018, but its data are not yet publicly available.

[7] https://microdata.worldbank.org.

Table 3: School, teacher, and student characteristics

	All schools		ST sample		STOC sample	
	Mean/%	SD	Mean/%	SD	Mean/%	SD
Schools						
Location						
Urban	15.3%		15.2%		9.8%	
Rural	84.7%		84.8%		90.3%	
*Equipment index**						
1	76.4%		74.9%		68.1%	
0	23.6%		25.1%		31.9%	
Students with textbooks	0.67	0.02	0.66	0.02	0.65	0.03
Student–teacher ratio	19.7	1.0	20.2	1.0	16.7	1.1
Fourth-grade classrooms	1.6	1.6	1.5	1.0	1.0	0.0
Number of students	532.5	640.5	510.2	528.6	310.3	142.8
Teachers						
Gender						
Female	63.4%		71.2%		77.8%	
Male	36.6%		28.8%		22.2%	
Age	32.8	8.2	31.2	8.0	31.6	8.4
Experience (years)	8.9	8.5	7.6	8.2	7.6	8.6
Students						
Gender						
Female	47.1%		46.9%		45.1%	
Male	52.9%		53.1%		54.9%	
First language						
Portuguese	4.2%		4.1%		1.0%	
Other	95.8%		95.9%		99.0%	
Age	11.1	1.9	11.1	1.9	11.3	1.9

Note: * in the SDI data this was defined as having a functioning blackboard and chalk, and the share of students with exercise books and pens being over 90 per cent. *N* = 203 schools, 673 teachers, and 1,761 students for all schools. *N* = 171 schools, 171 teachers, and 1,505 students for the ST sample. *N* = 113 schools, 113 teachers, and 975 students for the STOC sample.

Source: author's calculations based on the SDI data.

There were some minor differences in school enrolment between girls and boys. In rural schools boys comprised 54 per cent of students, whereas in urban schools it was 48 per cent. Students' first language also displays large differences between urban and rural students. Nearly 17 per cent of urban students spoke Portuguese as their first language, while this was extremely rare among rural students, with less than 2 per cent speaking Portuguese as their first language.[8] Students' ages suggest that many of the students have either enrolled in a primary school later than at the intended age of six, or have had to repeat grades.

It is worth noting that the ST sample shows few deviations from the whole sample, which provides support for interpreting the results as nationally representative. Teachers were slightly less experienced, and were more often female. These differences were repeated in the STOC sample. In additions, schools in the STOC sample were more often rural, more modestly equipped, smaller in size, and had a smaller share of students with Portuguese as their first language.

4.2 Student and teacher assessments

Students' and teachers' knowledge in maths and Portuguese was measured using tests designed by professionals in pedagogy, and compared against the national curricula of 13 African countries. The questions were harmonized to allow for comparison across countries and over time. The student test was

[8] The share of Portuguese speakers by province ranged from 1 per cent in Sofala to 65 per cent in Maputo City.

administered to ten randomly chosen fourth-grade students from each school. The assessments were carried out one-on-one with the student and the enumerator, and the questions were read out loud, which enabled students to respond to questions even without adequate reading skills to understand the assignment. In addition to students, their current and previous teachers were tested to the extent possible, as were teachers from other grades. Values of Cronbach's alpha suggest that the assessments could be considered as reliable measures of subject knowledge. For the maths and Portuguese assessments of the Mozambican SDI data, the results were $\alpha_M = 0.72$ and $\alpha_P = 0.86$ for the student assessment, and $\alpha_M = 0.79$ and $\alpha_P = 0.81$ for the teacher assessment.

The student assessment covered topics from first to third grade, and measured students' knowledge in maths and Portuguese, along with a four-question non-verbal reasoning section to control for innate ability. The Portuguese section consisted of six tasks, while the maths section consisted of six tasks comprising 17 questions. The teacher assessment similarly covered topics from first to third grade, but also from higher grades as it is, for example, necessary to know Portuguese beyond the level of a fourth-grader to be able to teach it. In addition to the maths and Portuguese sections, the assessment also included a pedagogy section. The format of the assessment was different for teachers: instead of taking the same test, teachers were asked to mark mock student exams and provide the correct answer where necessary. This format better matches teachers' daily activities, but also acknowledges teachers as professionals.

The results of the assessments confirm that there is a learning crisis in Mozambique. Table 4 shows how students performed in selected tasks in the assessments. Many of the students lacked the basic skills required to read in Portuguese. However, these Portuguese skills were spread on a wide scale, where the gap between the best-performing and the worst-performing students was vast. A student on the 10th percentile scored only 3 points out of the total of 100, while a student on the 90th percentile scored 87 points. There were few students who placed in the middle. Students' knowledge in maths was at a slightly higher level, at least when looking at the share of students who could perform the most basic tasks. It seems, however, that many lacked deeper understanding of the topics they were learning, and were unable to apply what they had learned, as can be seen from the small percentage of students who were able to solve a word problem that included a multiplication task.

Table 4: Student knowledge

	All students	ST sample	STOC sample
Portuguese			
Read a letter	39.9%	41.2%	36.7%
Read a word	24.5%	26.2%	23.6%
Read a sentence	11.8%	12.5%	11.0%
Read a paragraph	3.7%	4.1%	3.1%
Total points (out of 100)	21.9	23.1	21.2
Mathematics			
Identify numbers	83.3%	83.9%	81.8%
Add single digits	52.7%	52.8%	53.6%
Add double digits	21.8%	22.3%	23.3%
Subtract single digits	38.9%	39.4%	39.8%
Subtract double digits	7.2%	7.4%	7.1%
Multiply single digits	7.0%	7.0%	6.9%
Divide single digits	17.8%	18.5%	16.8%
Solve a word problem	3.9%	3.4%	4.3%
Total points (out of 100)	25.4	25.7	25.1

Note: N = 1,758 for all students; N = 1,502 for the ST sample; N = 972 for the STOC sample. In tasks in which a student was asked to read a letter, read a word, and identify numbers, the task was repeated three times and the student is reported as being able to perform this task if he/she correctly did so on all three occasions.

Source: author's calculations based on the SDI data.

Similar to student knowledge, teacher knowledge was at an alarmingly low level, and did not cover the topics included in the curriculum (Table 5). The SDI data define a teacher as having the minimum knowledge if he/she receives 80 per cent or more of the points in the assessment. This threshold was practically never reached in Portuguese, and rarely in maths. While the total points of both assessments were roughly the same, teachers' knowledge in maths was more equally distributed, whereas over 90 per cent of teachers could not respond correctly to even half of the tasks in the Portuguese assessment. A teacher is not able to teach students anything past his/her own knowledge, and teacher knowledge below the minimum makes it practically impossible for students to learn the skills listed in the curriculum.

Table 5: Teacher knowledge

	All teachers	ST sample	STOC sample
Portuguese			
100%	0.0%	0.0%	0.0%
80%	0.2%	0.0%	0.0%
50%	9.8%	5.9%	3.5%
Total points (out of 100)	33.6	32.4	29.8
Mathematics			
100%	0.3%	0.6%	0.0%
80%	7.0%	7.6%	6.2%
50%	30.6%	29.8%	26.6%
Lower primary (out of 100)	37.4	38.4	36.2
Upper primary (out of 100)	20.5	22.4	22.5
Total points (out of 100)	31.5	32.8	31.4

Note: N = 673 for all teachers; N = 171 for the ST sample; N = 113 for the STOC sample.

Source: author's calculations based on the SDI data.

Item response theory (IRT) was used to derive measures of the underlying levels of knowledge for students and teachers in the empirical analysis. This is because using the sum of correct answers ignores the levels of difficulty of each task. Measures of teacher knowledge in maths and Portuguese were derived using a two-parameter logistic model, a standard model for data whose items (tasks) are scored as either correct or incorrect. The level of knowledge (θ) is determined by two item parameters, difficulty (β) and discrimination (α). The difficulty parameter, according to its name, reflects the difficulty of an item so that when $\beta = \theta$ the probability of a correct answer is 0.5. The discrimination parameter, on the other hand, describes how well an item differentiates between individuals with different levels of knowledge. In the two-parameter logistic model, the probability of a person with knowledge level θ giving a correct response to an item i is given by

$$P_i(\theta) = \frac{\exp\big(\alpha_i(\theta - \beta_i)\big)}{1 + \exp\big(\alpha_i(\theta - \beta_i)\big)} \tag{2}$$

Measures of student achievement in maths and Portuguese were derived using a generalized partial credit model, which is an extension of the two-parameter logistic model for polytomous data, where items are scored on an ordered scale of $x = 0, 1, 2, \ldots, m_i$. This is because students were asked to repeat some tasks three times, for instance, and the partial credit model also takes into account the partially correct answers. Instead of giving the probability of a correct response, the generalized partial credit model specifies the probability of scoring $x = j$ points from item i. This is modelled as

$$P_{ix}(\theta) = \frac{\exp \sum_{j=0}^{x} \alpha_i(\theta - \delta_{ij})}{\sum_{r=0}^{m_i} (\exp \sum_{j=0}^{r} \alpha_i(\theta - \delta_{ij})} \tag{3}$$

where $\sum_{j=0}^{0}(\theta - \delta_{ij}) \equiv 0$.

5 Estimation of the education production function

The question that this study tries to answer is: *how much does teacher content knowledge impact student achievement in a low-income country?* The starting point is the education production function, whose general formulation was given in Equation 1. I use a specification of the education production function that puts special emphasis on relating teacher content knowledge to student achievement, given by

$$y_{ijcs} = \beta_0 + \delta_0 P_j + \beta_1 T_{jcs} + \beta_2 X_{ics} + \beta_3 Z_{cs} + \beta_4 W_s + v_{ijcs} \tag{4}$$

The dependent variable is y_{ijcs}, the IRT score of student i in subject j in classroom c of school s. The content knowledge of his/her teacher is denoted with T_{jcs}, the teacher IRT score. P_j is a subject dummy that receives a value of 1 if the subject was Portuguese, and 0 otherwise. Other independent variables include vectors of subject-invariant characteristics of students, denoted with X_{ics}, teachers, denoted with Z_{cs}, and schools, denoted with W_s, while v_{ijcs} is the error term.[9] It should be noted that the data did not include any other subject-varying factors besides those from student and teacher assessments. Any other subject-varying characteristics were therefore included in the error term.

Estimating Equation 4 with ordinary least squares (OLS) is likely to yield biased estimates due to omitted variables that can be both determinants of student achievement and correlated with teacher content knowledge. In the context of education this is likely to arise from the non-random assignment of students to specific schools or teachers based on unobservable characteristics, as discussed earlier. If the assignment is non-random, OLS cannot distinguish the causal effect of teacher content knowledge from the pre-existing, unobservable differences of students who were taught by the same teacher. For my analysis, the key was to distinguish whether the assignment was based on subject-invariant or subject-varying characteristics.

I treated the data as a panel, and included the student fixed effect into the model. Panel data typically include observations of the same individuals over time, but here, the data included observations of the same students and their teachers in two subjects, maths and Portuguese. This is crucial for the identification, because it allowed removing any characteristics that stayed constant across the two subjects. An example is the innate ability of a student, which cannot be directly measured, but affects academic achievement, and can also be a reason for non-random assignment of students to teachers.

To distinguish between the subject-invariant and subject-varying unobserved student characteristics, the error term v_{ijcs} was broken down into a composite error term with two parts, a_{ics} and u_{ijcs}, where the former represents any subject-invariant unobserved characteristics of students, and the latter is the remaining idiosyncratic error. Oftentimes, a_{ics} is referred to as the student fixed effect, which emphasizes the fact that it stays constant across subjects. The education production functions for maths and Portuguese could therefore be written as

$$y_{ics,M} = \beta_0 + \beta_1 T_{cs,M} + \beta_2 X_{ics} + \beta_3 Z_{cs} + \beta_4 W_s + a_{ics} + u_{ics,M} \tag{5a}$$

$$y_{ics,P} = \beta_0 + \delta_0 + \beta_1 T_{cs,P} + \beta_2 X_{ics} + \beta_3 Z_{cs} + \beta_4 W_s + a_{ics} + u_{ics,P} \tag{5b}$$

To remove the problematic student fixed effect a_{ics} from the model, Equation 5a was subtracted from Equation 5b, leading to a first-differenced model:

$$y_{ics,P} - y_{ics,M} = \delta_0 + \beta_1 (T_{cs,P} - T_{cs,M}) + (u_{ics,P} - u_{ics,M}) \tag{6}$$

[9] Subject-invariant characteristics refer to characteristics that stay constant across subjects, such as gender or innate ability. Examples of subject-varying characteristics other than teacher content knowledge include pedagogical methods and teaching time, if they vary across subjects.

which can also be written as

$$\Delta y_i = \delta_0 + \beta_1 \Delta T_{cs} + \Delta u_{ics} \tag{7}$$

In this model, student achievement is only explained by teacher content knowledge. The coefficient β_1 is the causal impact of teacher content knowledge on student achievement. Because the IRT scores were standardized to have a SD of 1, increasing teacher content knowledge by 1 SD would lead to a change of size β_1 SD in student achievement. Although not the main interest, the intercept δ_0 is the change in intercept between maths and Portuguese.

Besides removing the student fixed effect from the equation, first-differencing also removed all subject-invariant characteristics of teachers and schools. This is because all students in my sample were taught by the same teacher in both maths and Portuguese, and a student naturally attended only one school. Using observations from the same teachers in both subjects corresponds to including the teacher fixed effect into Equations 5a and 5b (Bietenbeck et al. 2018). All variation is now only within the IRT scores of students and their teachers, and the results from estimating the first-differenced Equation 7 with OLS are not subject to omitted variable bias.

The causal identification of teacher content knowledge relies on two main assumptions:

1. The impact of teacher content knowledge, β_1, is the same for both maths and Portuguese.
2. There is no correlation between ΔT_{cs} and Δu_{ics} $(\mathbf{E}[\Delta T_{cs}' \Delta u_{ics}] = 0)$.

I approached the first assumption by comparing the OLS estimates separately for both subjects. These results should be interpreted with caution, however, because they are subject to the limitations of OLS discussed earlier. A more credible method to assess the assumption would be to follow Metzler and Woessmann (2012). These authors used a correlated random effects model, which allowed testing the similarity of the coefficients. The second assumption, on the other hand, could be violated if students were assigned to teachers based on some unobserved subject-varying heterogeneity. This could happen if students who were relatively stronger in maths than in Portuguese in the previous year, for example, were assigned to teachers who also have a stronger knowledge in maths. I addressed the issue by including a specification where the sample was narrowed down to students in schools with only one fourth-grade classroom (STOC sample). In this setting, it is not possible to assign students to specific teachers within schools.

5.1 Measurement error

Regardless of the assessments being carefully designed and implemented, they are still only indirect measures of the underlying knowledge. Therefore they include some amount of measurement error, which biases the causal estimate of teacher content knowledge. Following the classical measurement error model, teacher content knowledge can be written as

$$\widetilde{T_{jcs}} = T_{jcs} + e_{jcs} \tag{8}$$

where $\widetilde{T_{jcs}}$ is the observed teacher IRT score, T_{jcs} is the true value of the underlying teacher content knowledge, and e_{jcs} is the measurement error. It is necessary to assume that $\mathbf{E}(e_{jcs}) = 0$, $\text{Var}(e_{jcs}) > 0$, and that e_{jcs} is independent of T_{jcs} and v_{ijcs}. Substituting this into Equation 7 gives

$$\Delta y_i = \delta_0 + \beta_1 \Delta \widetilde{T_{cs}} + \beta_1 \Delta e_{cs} + \Delta u_{ics} \tag{9}$$

This creates an endogeneity issue, because the estimator $\hat{\beta}_1$ that describes the impact of teacher content knowledge not only depends on the teacher IRT scores, but also on the measurement error. More specifically, this biases the estimator towards zero, hence called the attenuation bias. This is because

$$\text{plim}\hat{\beta}_1 = \beta_1 \left(\frac{\text{Var}(\Delta T_{cs})}{\text{Var}(\Delta T_{cs}) + \text{Var}(\Delta e_{cs})} \right) = \beta_1 \lambda \tag{10}$$

where λ is called the reliability ratio and has a value between 0 and 1. Fixed effects estimates are particularly susceptible to attenuation bias, because differencing reduces the variance in the independent variable to a fraction. If there is not enough variation in the difference between the levels of teacher content knowledge in maths and Portuguese, much of the observed variation in the independent variable can in fact be due to the measurement error.

Equation 10 implies that the unbiased impact of teacher content knowledge is given by $\beta_1 = \frac{\hat{\beta}_1}{\lambda}$. It is not possible to observe the true value of λ, but following Metzler and Woessmann (2012), Cronbach's α can provide an estimate of it. Considering that the values for teacher assessments in both subjects were around 0.8, this suggests that the estimate of teacher content knowledge could be biased downward by around 20 per cent.

6 Results

6.1 Main results

Ordinary least squares

OLS was used first to analyse the relationship between teacher content knowledge and student achievement. Table 6 presents the results from estimating Equation 4 with OLS. These results show that teacher content knowledge was not associated with student achievement (column 1). There was a small, positive coefficient of size 0.047, but it was statistically insignificant.[10] This is a surprising result, because teacher content knowledge is expected to correlate strongly with overall teacher quality and, therefore, with student achievement. Adding a set of control variables did not change the result (column 2), but suggested that material resources were the most strongly associated with student achievement, as the coefficients for the share of students with textbooks and the equipment index show. On the contrary, there is no evidence that any of the observed teacher characteristics were associated with student achievement.

The final specifications compared the coefficients for maths and Portuguese separately. This gives a rough understanding of whether the role of teacher content knowledge varies between them. The results reveal that teacher content knowledge was not associated with student achievement for either of the subjects, although the coefficient for maths was slightly larger (column 3). After including a set of control variables, the relationship became even weaker (column 4). Again, the role of material resources was highlighted, and they were relatively more important for Portuguese than for maths. The results also show that there were gender gaps that disfavoured girls, and these were considerably larger in maths.

[10] If the relationship were statistically significant, this should be interpreted so that an increase of 1 SD in teacher content knowledge is associated with an increase of 0.047 SD in student achievement.

Table 6: Results from OLS regression

	Student achievement					
	(1)	(2)	(3)		(4)	
	Both	Both	Maths	Portuguese	Maths	Portuguese
Teacher content knowledge	0.047 (0.044)	0.026 (0.038)	0.041 (0.047)	0.055 (0.063)	0.031 (0.045)	–0.002 (0.057)
Student age		0.049*** (0.093)			0.050*** (0.014)	0.048*** (0.008)
Student is female		–0.146*** (0.048)			–0.219*** (0.053)	–0.071 (0.054)
Student's first language is Portuguese		0.248 (0.175)			0.195 (0.162)	0.304 (0.208)
School in urban location		–0.167 (0.108)			–0.158 (0.120)	–0.168 (0.119)
Student–teacher ratio		0.012*** (0.003)			0.009*** (0.003)	0.015*** (0.004)
Textbook availability		0.360*** (0.124)			0.306** (0.129)	0.413*** (0.141)
Equipment index = 1		0.207** (0.095)			0.152 (0.098)	0.268** (0.108)
Teacher is female		-0.056 (0.092)			–0.014 (0.095)	–0.096 (0.109)
Teacher experience		0.000 (0.006)			–0.002 (0.005)	0.003 (0.007)
Teacher has received training		0.073 (0.211)			0.067 (0.197)	0.086 (0.231)
Observations (students)	1,505	1,505	1,505	1,505	1,505	1,505
Observations (teachers)	171	171	171	171	171	171

Note: student achievement is indicated by student IRT scores, and teacher content knowledge by teacher IRT scores, both of which are standardized to have a mean of 0 and a standard deviation of 1. Standard errors are clustered at the classroom level, and presented in parentheses. Significance at *** 0.01, ** 0.05, and * 0.1 levels.

Source: author's calculations based on the SDI data.

Fixed effects

Table 7 presents the results of estimating the first-differenced Equation 7, which allows studying the causal relationship between teacher content knowledge and student achievement more convincingly. The results suggest that the null hypothesis of teacher content knowledge having no impact on student achievement cannot be rejected. The coefficient was small and positive, of size 0.041, but it was statistically insignificant.[11] It also seems that the results are not biased by subject-specific assignment within schools. Table A1 presents the results of estimating the first-differenced model for students taught by the same teacher in schools with only one fourth-grade classroom (STOC sample). In this specification, the estimate remained statistically insignificant, although became slightly smaller.[12]

The results are largely in line with those found in previous research. Bietenbeck et al. (2018) and Bold et al. (2017) obtained statistically significant effects of sizes 0.025 SD and 0.02 SD, respectively. Taking a slightly different methodological approach, Metzler and Woessmann (2012) found an effect of size 0.087 SD for maths, and Filmer et al. (2015) an effect of 0.05 SD for maths, while the results were insignificant for reading. Geographically closest, Shepherd (2015) estimated an effect of 0.013 SD for

[11] If the coefficient were statistically significant, it should be interpreted so that a 1 SD increase in teacher content knowledge leads to an increase of 0.041 SD in student achievement.

[12] It should be kept in mind that the STOC sample differs from the ST sample in two important aspects, namely the share of rural schools and non-Portuguese-speaking students, which goes against the generalizability of these results.

South Africa, which was statistically insignificant. Moreover, other studies such as Marcenaro-Gutierrez and Lopez-Agudo (2020) and Zakharov et al. (2016) obtained statistically insignificant results.

Table 7: Results from the first-differenced model

	Student achievement
Teacher content knowledge	0.041 (0.031)
t-statistic	1.31
Prob > t	0.19
Observations (students)	1,505
Classrooms (clusters)	171

Note: student achievement is indicated by student IRT scores, and teacher content knowledge by teacher IRT scores, both of which are standardized to have a mean of 0 and a standard deviation of 1. Standard errors are clustered at the classroom level and presented in parentheses. Significance at *** 0.01, ** 0.05, and * 0.1 levels.

Source: author's calculations based on the SDI data.

There are some possible causes or mechanisms behind not finding the impact. An obvious concern in the Mozambican context is that the majority of teachers do not master the curriculum well enough themselves. Teachers are not able to teach past their own level of knowledge, and this is perhaps reflected in the results, because the tests measured knowledge at a level that neither teachers nor students generally had achieved. It is possible that for teacher knowledge to have a larger effect there is a minimum level that teachers need to reach, such as mastering the curriculum themselves. Beyond that, existing research does not provide evidence on whether the impact of teacher content knowledge is linear.

A possibly important consideration is that the students in my sample were fourth-graders, and that the impact of teacher content knowledge may become more important when moving to higher grades. Much of the topics covered by the curriculum during the first grades could be considered basic cognitive skills, as opposed to specific content knowledge. As the topics become more difficult in the higher grades, it is possible that the role of teacher content knowledge is more relevant. Existing research does not directly touch upon this hypothesis, but Hanushek et al. (2019), for example, found an effect of 0.105 SD for 15-year-olds from 31 OECD countries. Similarly, my results only apply for maths and Portuguese. It is possible that assessments in other subjects might give different results, but existing research has primarily focused on maths and language/reading.

How to interpret these results? First, teacher content knowledge is currently poor overall. As teachers cannot teach past their own knowledge level, the majority of students do not even have an actual possibility to reach the requirements set by the curriculum. Second, measures to improve student achievement need to be found elsewhere. Considering that Mozambique is a poor country with very scarce resources, these should be targeted to what has been shown to improve learning more in the local context. This is particularly important because raising teacher content knowledge to a sufficient level would be an enormous task. An average teacher would need to more than double his/her points from the assessment to meet the minimum requirements.

Finally, there are some limitations related to data and methodology that may affect the results. Measurement error is likely to bias the estimate of the impact of teacher content knowledge downward, as already discussed. As I did not correct my results for measurement error, it is unclear whether the corrected estimate would still remain statistically insignificant. Moreover, my sample size of 1,505 students is considerably smaller than in comparable multi-country studies, ranging from around 7,000 to 75,000 students. Considering that earlier literature has found heterogeneity in their estimates across countries, it is nevertheless important to study individual countries, despite smaller sample sizes. This is especially relevant for policy, which is always national.

6.2 Robustness checks

Heterogeneity of effects

As was evident in the SDI data, there are large differences within Mozambique. Economic development has favoured urban areas, and Portuguese-speaking students and southern regions perform much better in student assessments. While the results showed no impact of teacher content knowledge on learning, it is possible that they hide in-country heterogeneity, where teacher content knowledge disproportionately benefits some groups of students. Heterogeneity in the effects across different groups may indicate that there are binding constraints to learning that need to be alleviated first (Glewwe and Muralidharan 2016). For instance, Glewwe et al. (2009) conducted a randomized experiment in Kenya and found that providing textbooks had little impact on an average student, but raised the scores of students who already had the best baseline scores. They concluded that the binding constraint were poor reading skills, and that those should be improved before providing textbooks in order for the books to have an impact.

To study the issue of heterogeneity, I introduce interaction terms into the model. These indicate whether teacher knowledge has a different impact on student achievement for different groups. This was done by adding an independent variable of form $\Delta T_{cs} \times$ *variable of interest* into the regression of Equation 7. Results from the analysis in Table 8 suggest that only students' first language and urban location of a school showed a larger, statistically significant interaction. An increase of 1 SD in teacher content knowledge increased student achievement by 0.14 SD for students whose first language was Portuguese, but only 0.03 SD for others. Likewise, an increase of 1 SD improved student achievement by 0.13 SD for students in urban schools, but only 0.02 SD for those in rural schools. It should be noted that the causal interpretation of these results is not without problems, because urban location, for instance, is related to several unobserved characteristics that this specification does not rule out, as opposed to my baseline model.

It seems plausible that students' poor language skills may be a binding constraint to learning. Portuguese-speaking students likely experienced a larger impact because they better understand the language of instruction, whereas the language skills of others are too weak to follow the teaching. Teacher content knowledge has very little importance if the message cannot be conveyed. In this sense, the results suggest that poor language skills may be a more fundamental obstacle for learning than poor teacher content knowledge. Improving language education at the earlier grades or a wider introduction of bilingual education that gradually introduces students to Portuguese might be fruitful approaches to improve learning. How best to do it remains an important policy question.

While urban schools had a larger share of students with Portuguese as their first language,[13] there were no large differences in their performance on the Portuguese assessment. Therefore language is unlikely to be the only explanation in this case. Urban schools, for example, are generally surrounded by a more prosperous environment, as opposed to rural schools. Considering that more recent economic development in the country has disproportionately benefited urban areas over rural ones, it is worrying that also teacher content knowledge affects those in urban but not in rural schools, particularly because there were no differences in levels of teacher knowledge between urban and rural schools.

[13] In urban schools, 16.9 per cent of students reported Portuguese as their first language. In rural schools the share was 1.6 per cent.

Table 8: Results from the interaction analysis

	Student achievement							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Teacher content knowledge	0.087 (0.066)	0.033 (0.032)	0.037 (0.042)	0.092 (0.126)	0.030 (0.030)	0.019 (0.036)	0.050 (0.070)	0.048 (0.057)
× student gender	–0.032 (0.038)							
× student first language		0.107* (0.064)						
× student non-verbal reasoning			0.008 (0.075)					
× teacher gender				–0.032 (0.071)				
× teacher pedagogy score					–0.046 (0.036)			
× urban location						0.110* (0.060)		
× equipment index							–0.012 (0.078)	
× textbook availability								–0.011 (0.084)
t-statistic	–0.83	1.68	0.11	–0.45	–1.26	1.84	–0.16	–0.14
Prob > t	0.405	0.095	0.916	0.654	0.209	0.068	0.876	0.893
Observations (students)	1,505	1,505	1,505	1,505	1,505	1,505	1,505	1,505
Classrooms (clusters)	171	171	171	171	171	171	171	171

Note: student achievement is indicated by student IRT scores, and teacher content knowledge by teacher IRT scores, both of which are standardized to have a mean of 0 and a standard deviation of 1. Standard errors are clustered at the classroom level and presented in parentheses. Significance at *** 0.01, ** 0.05, and * 0.1 levels.

Source: author's calculations based on the SDI data.

Even though teachers' pedagogical knowledge showed no interaction, the explanation can be similar as suggested for the main results—the level of pedagogical knowledge is possibly too low for it to have any effect. In general, most Mozambican teachers had really poor pedagogical knowledge. There may not only be a crisis of teacher content knowledge, but also of pedagogical knowledge. After all, one can expect that teachers with better, or at least minimal, pedagogical knowledge would be more effective at passing their content knowledge onto their students.

Interestingly, material resources in the form of textbook availability and classroom equipment index did not interact with teacher content knowledge. This is against the results of Bietenbeck et al. (2018),[14] who found that the impact of teacher content knowledge was twice as large for students who had textbooks, and also increased with an index of several school facilities. They also hypothesized that material resources may explain the heterogeneity in results between richer and poorer countries, for which my analysis does not provide support.

[14] From a methodological perspective, the equipment indices used by this study and Bietenbeck et al. (2018) are not fully comparable. My analysis focused on basic equipment, such as blackboards, whereas the index of Bietenbeck et al. (2018) also included more advanced resources, which are likely to be less relevant for the Mozambican context.

Student–teacher matching by levels of knowledge

Another possible source of heterogeneity that I addressed are the different levels of knowledge of students and teachers. Following Metzler and Woessmann (2012), I studied the impact for pairs of students and teachers, with their level of knowledge above or below the median. This should reveal whether the matching matters, and whether the impact varies by the level of teacher content knowledge. It should be noted that most of the teachers who were above the median still did not have adequate knowledge to master the curriculum, but were relatively better in this sample. As indicated by Table 5, narrowing down to teachers who actually had the minimum knowledge[15] would leave the sample far too small for analysis. Similarly, only a small group of students received high total points in the assessment, while the majority of students even above the median could not master the curriculum.

Table 9 presents the findings from the impact of teacher content knowledge for pairs of students and teachers below and above the median.[16] First, students who were below the median did not show differences in their results by teacher knowledge level, and overall the impact was statistically insignificant. Compared to students who were above the median and experienced larger impacts, this may indicate that teachers do not target their teaching to those who have fallen behind. As can be seen from the learning results, many students lack the foundation on which the topics from higher grades are based. Without this foundation, teacher content knowledge is rather irrelevant as the topics are too advanced in any case.

Table 9: Matching of students and teachers by levels of knowledge

	Student achievement			
Teacher below/above median	Below		Above	
Student below/above median	Below	Above	Below	Above
Teacher content knowledge	0.018 (0.040)	0.035 (0.072)	0.010 (0.048)	0.123* (0.068)
t-statistic	0.45	0.48	0.20	1.81
Prob > t	0.657	0.632	0.843	0.075
Observations (students)	384	369	370	382
Classrooms (clusters)	84	84	73	78

Note: student achievement is indicated by student IRT scores, and teacher content knowledge by teacher IRT scores, both of which are standardized to have a mean of 0 and a standard deviation of 1. Standard errors are clustered at the classroom level and presented in parentheses. Significance at *** 0.01, ** 0.05, and * 0.1 levels.

Source: author's calculations based on the SDI data.

The second finding to emerge from the student–teacher matching is that students who were above the median experienced a significantly larger effect from teacher content knowledge when taught by a teacher who was also above the median. For these students, an increase of 1 SD in teacher content knowledge improved student achievement by over 0.12 SD. Because these same teachers did not have an impact on students below the median, this further supports the claim that teaching is not targeted to those who have fallen behind, but the pace may be set by the more advanced students. As variation in student achievement can be large within a single classroom, it makes teaching all students at the right level difficult for the teacher.

[15] This was defined as receiving 80 per cent or more of the total points in the teacher assessments.

[16] These were calculated as the average of the IRT scores in maths and Portuguese for each student/teacher. This approach was chosen because it puts more emphasis on *general* academic ability instead of performing well in only one subject.

7 Conclusions

The aim of this study was to determine the causal impact of teacher content knowledge on student achievement in Mozambican primary schools. Learning results in Mozambique have stayed at a low level compared to other sub-Saharan African countries, calling for action to improve learning. Earlier research has established that teacher quality matters greatly for learning, and found that out of all observable characteristics, teacher content knowledge has most consistently had a positive impact on student achievement. Most of this research, however, comes from high-income countries.

I used data from the nationally representative SDI survey from 2014, which included assessments of students and their teachers in maths and Portuguese. To estimate the causal impact of teacher content knowledge, I exploited within-student across-subject variation. The non-random assignment of students based on unobservable characteristics is the leading cause of bias, but I was able to circumvent it through first-differencing, which eliminated the student fixed effect from the estimation. Furthermore, because all students were taught by the same teacher in both subjects, the teacher fixed effect was also eliminated.

My main results show that, on average, teacher content knowledge did not have an impact on student achievement, although measurement error is likely to have biased the estimate downward. The result is similar to the results of Marcenaro-Gutierrez and Lopez-Agudo (2020) from Botswana, Lesotho, and Zambia, and Shepherd (2015) from South Africa, for example. However, there was considerable heterogeneity in the effect. Raising teacher content knowledge by 1 SD was associated with increases in student achievement of size 0.14 SD for students whose first language was Portuguese, and 0.13 SD for students in urban schools. This is not unexpected, because Mozambique is a country with large contrasts between the urban and rural areas, different regions, and higher- and low-income households. Furthermore, raising the content knowledge of teachers who were above the median by 1 SD increased by 0.12 SD the achievement of students who were also above the median. But I found no interaction between teacher content knowledge and material resources of schools, whose importance Bietenbeck et al. (2018) highlighted.

The observed heterogeneity suggests that students' poor knowledge of Portuguese may be a binding constraint for learning, therefore lowering the effect of teacher content knowledge, and so something that should be prioritized in public policy. This could be achieved by improving language education at the earlier grades, or by extending bilingual education that only gradually introduces students to Portuguese. Because the content knowledge of teachers who were above the median only affected students who were also above the median, there is also reason to believe that the pace is set by the more advanced students. Targeting teaching to match the level of all students can be difficult in practice, because the distribution of student knowledge can be wide within a single classroom. In Portuguese, for example, a student on the 10th percentile only scored 3 out of 100 points, while a student on the 90th percentile scored 87 points. The results of this study encourage conducting further research on students' poor Portuguese skills as a possible binding constraint for learning.

In addition, there seems to be a gap in research on the impact of teacher content knowledge in low- and middle-income countries when it comes to students past primary school. The existing results may not generalize into secondary schools, where the role of teacher content knowledge can be more relevant as the subjects move from general literacy and numeracy to more subject-specific knowledge. Extending compulsory education in Mozambique up to ninth grade also underlines the need for high quality throughout this compulsory education.

References

Baez Ramirez, J.E., G.D. Caruso, C. Niu, and C.A. Myers (2018). 'Mozambique Poverty Assessment: Strong But Not Broadly Shared Growth'. Report 131218. Washington, DC: World Bank Group.

Bassi, M., O. Medina Pedreira, and L.J. Nhampossa (2019). 'Education Service Delivery in Mozambique: A Second Round of the Service Delivery Indicators Survey'. Report. Washington, DC: World Bank Group.

Bau, N., and J. Das (2020). 'Teacher Value Added in a Low-Income Country'. *American Economic Journal: Economic Policy*, 12(1): 62–96. https://doi.org/10.1257/pol.20170243

Beutel, M. (2011). 'Teachers Talking: Primary Teachers' Contributions to the Quality of Education in Mozambique'. Report. London: Voluntary Service Overseas (VSO).

Bietenbeck, J., M. Piopiunik, and S. Wiederhold (2018). 'Africa's Skill Tragedy: Does Teachers' Lack of Knowledge Lead to Low Student Performance?' *Journal of Human Resources*, 53(3): 553–78. https://doi.org/10.3368/jhr.53.3.0616-8002R1

Bold, T., D. Filmer, E. Molina, and J. Svensson (2019). 'The Lost Human Capital: Teacher Knowledge and Student Achievement in Africa'. Policy Research Working Paper 8849. Washington, DC: World Bank Group. https://doi.org/10.1596/1813-9450-8849

Bold, T., D. Filmer, G. Martin, E. Molina, C. Rockmore, B. Stacy, J. Svensson, and W. Wane (2017). 'What Do Teachers Know and Do? Does It Matter? Evidence from Primary Schools in Africa'. Policy Research Working Paper 7956. Washington, DC: World Bank Group. https://doi.org/10.1596/1813-9450-7956

Chetty, R., J.N. Friedman, and J.E. Rockoff (2014a). 'Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates'. *American Economic Review*, 104(9): 2593–632. https://doi.org/10.1257/aer.104.9.2593

Chetty, R., J.N. Friedman, and J.E. Rockoff (2014b). 'Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood'. *American Economic Review*, 104(9): 2633–79. https://doi.org/10.1257/aer.104.9.2633

Coenen, J., I. Cornelisz, W. Groot, H.M. van den Brink, and C.V. Klaveren (2018). 'Teacher Characteristics and Their Effects on Student Test Scores: A Systematic Review'. *Journal of Economic Surveys*, 32(3): 848–77. https://doi.org/10.1111/joes.12210

Filmer, D., E. Molina, and B. Stacy (2015). 'What Goes on Inside the Classroom in Africa? Assessing the Relationship Between What Teachers Know, What Happened in the Classroom, and Student Performance'. Policy Research Working Paper 7956. Washington, DC: World Bank Group.

Glewwe, P., and K. Muralidharan (2016). 'Improving Education Outcomes in Developing Countries: Evidence, Knowledge Gaps, and Policy Implications'. In E.A. Hanushek, S. Machin, and L. Woessmann (eds), *Handbook of the Economics of Education*, volume 5. Amsterdam: Elsevier. https://doi.org/10.1016/B978-0-444-63459-7.00010-5

Glewwe, P., M. Kremer, and S. Moulin (2009). 'Many Children Left Behind? Textbooks and Test Scores in Kenya'. *American Economic Journal: Applied Economics*, 1(1): 112–35. https://doi.org/10.1257/app.1.1.112

Hanushek, E.A. (1997). 'Assessing the Effects of School Resources on Student Performance: An Update'. *Educational Evaluation and Policy Analysis*, 19(2): 141–64. https://doi.org/10.3102/01623737019002141

Hanushek, E.A. (2002). 'Teacher Quality'. In L.T. Izumi and W.M. Evers (eds), *Teacher Quality*. Stanford, CA: Hoover Institution Press.

Hanushek, E.A., M. Piopiunik, and S. Wiederhold (2019). 'The Value of Smarter Teachers: International Evidence on Teacher Cognitive Skills and Student Performance'. *Journal of Human Resources*, 54(4): 857–99. https://doi.org/10.3368/jhr.54.4.0317.8619R1

Hanushek, E.A., and S.G. Rivkin (2010). 'Generalizations about Using Value-Added Measures of Teacher Quality'. *American Economic Review*, 100(2): 267–71. https://doi.org/10.1257/aer.100.2.267

Magaia, F., T. Nahara, and A. Passos (2011). 'Trends in Achievement Levels of Grade 6 Pupils in Mozambique'. Policy Brief 1. Gaborone: Southern and Eastern Africa Consortium for Monitoring Educational Quality.

Marcenaro-Gutierrez, O.D., and L.A. Lopez-Agudo (2020). 'Does Teacher Subject Knowledge Contribute to Student Academic Performance in Developing and Least Developed Countries'. *South African Journal of Economics*, 88(3): 267–97. https://doi.org/10.1111/saje.12259

Marouma, A.L., A. Erambel, F.-L. Ibrah, M. Djibo, and A. Boukari (2017). 'Niger Service Delivery Indicators: Education 2015'. Washington, DC: World Bank Group.

Metzler, J., and L. Woessmann (2012). 'The Impact of Teacher Subject Knowledge on Student Achievement: Evidence from Within-Teacher Within-Student Variation'. *Journal of Development Economics*, 99(2): 486–96. https://doi.org/10.1016/j.jdeveco.2012.06.002

MINEDH (2017). *Relatório do Estudo Holístico da Situação do Professor em Moçambique, 2015*. Maputo: Ministério da Educação e Desenvolvimento Humano.

MINEDH (2020a). *Levantamento Escolar 2020: Educação Geral, Alfabetização e Educação de Adultos e Formação de Professores*. Maputo: Ministério da Educação e Desenvolvimento Humano.

MINEDH (2020b). *Plano Estratégico da Educação 2020–2029*. Maputo: Ministério da Educação e Desenvolvimento Humano.

Molina, E., and G. Martin (2015). 'Education Service Delivery in Mozambique: Results of 2014 Service Delivery Indicator Survey'. Report AUS3420. Washington, DC: World Bank Group.

Pritchett, L., M. Woolcock, and M. Andrews (2013). 'Looking Like a State: Techniques of Persistent Failure in State Capability for Implementation'. *Journal of Development Studies*, 49(1): 1–18. https://doi.org/10.1080/00220388.2012.709614

Reinikka, R., and E.A. Napaua (2019). 'Stepping Up Finland's Support to Education in Mozambique'. Report. Helsinki: Ministry for Foreign Affairs of Finland.

Rothstein, J. (2009). 'Student Sorting and Bias in Value-Added Estimation: Selection on Observables and Unobservables'. *Education Finance and Policy*, 4(4): 537–71. https://doi.org/10.1162/edfp.2009.4.4.537

Shepherd, D. (2015). 'Learn to Teach, Teach to Learn: A Within-Pupil Across-Subject Approach to Estimating the Impact of Teacher Subject Knowledge on South African Grade 6 Performance'. Stellenbosch Economic Working Paper 01/15. Stellenbosch: University of Stellenbosch.

Shulman, L. (1987). 'Knowledge and Teaching: Foundations of the New Reform'. *Harvard Educational Review*, 57(1): 1–23. https://doi.org/10.17763/haer.57.1.j463w79r56455411

Strøm, B., and T. Falch (2020). 'The Role of Teacher Quality in Education Production'. In S. Bradley and C. Green (eds), *The Economics of Education*, 2nd edition. New York: Academic Press.

UNESCO (2013). *The Global Learning Crisis: Why Every Child Deserves a Quality Education*. Paris: UNESCO.

UNESCO and MINEDH (2019). *Revisão de Políticas Educacionais de Moçambique*. Paris: UNESCO.

Wane, W., and R.A. Rakotoarivony (2017). 'Education Service Delivery in Madagascar: Results of 2016 Service Delivery Indicator Survey'. Report AUS18887. Washington, DC: World Bank Group.

Wayne, A.J., and P. Youngs (2003). 'Teacher Characteristics and Student Achievement Gains: A Review'. *Review of Educational Research*, 73(1): 89–122. https://doi.org/10.3102/00346543073001089

World Bank (2021a). 'Pupil–Teacher Ratio, Primary: Mozambique, Sub-Saharan Africa'. Available at: https://data.worldbank.org/indicator/SE.PRM.ENRL.TC.ZS?locations=MZ-ZG (accessed 22 February 2022).

World Bank (2021b). 'Trained Teachers in Primary Education (% of Total Teachers): Mozambique'. Available at: https://data.worldbank.org/indicator/SE.PRM.TCAQ.ZS?locations=MZ (accessed 22 February 2022).

World Bank (2021c). 'Primary Education, Teachers: Mozambique'. Available at: https://data.worldbank.org/indicator/SE.PRM.TCHR?locations=MZ (accessed 22 February 2022).

World Bank (2021d). 'Urban Population (% of Total Population): Mozambique'. Available at: https://data.worldbank.org/indicator/SP.URB.TOTL.IN.ZS?locations=MZ (accessed 22 February 2022).

Zakharov, A., G. Tsheko, and M. Carnoy (2016). 'Do "Better" Teachers and Classroom Resources Improve Student Achievement? A Causal Comparative Approach in Kenya, South Africa, and Swaziland'. *International Journal of Educational Development*, 50: 108–24. https://doi.org/10.1016/j.ijedudev.2016.07.001

Appendix A

Table A1: Results from the first-differenced model for the STOC sample

	Student achievement
Teacher content knowledge	–0.001 (0.041)
t-statistic	–0.02
Prob > t	0.988
Observations (students)	975
Classrooms (clusters)	113

Note: student achievement is indicated by student IRT scores, and teacher content knowledge by teacher IRT scores. The IRT scores are standardized to have a mean of 0 and a standard deviation of 1. Regression is estimated by OLS. Standard errors are clustered at the classroom level and presented in parentheses. Significance at *** 0.01, ** 0.05, and * 0.1 levels.

Source: author's calculations based on the SDI data.