STOCKHOLM SCHOOL OF ECONOMICS
Department of Economics
5350 Master's thesis in economics
Academic year 2021–2022

Rural Road Improvements and Local Agricultural Intensification

A Remote Sensing Evaluation in Mozambique

Jeffrey Clark (23500) and Tamina Matti (23660)

Abstract

Transaction costs serve as an obstacle to competitive market exchanges in rural and remote areas around the world. Improvements to transportation infrastructure are hypothesized to lower these costs and help alleviate poverty among smallholder farmers. Yet, few empirical studies estimate the effect of improved rural infrastructure on agricultural output, especially in the sub-Saharan context. This thesis investigates whether rural road upgrades in northern Mozambique have any short-term effects on agricultural output; specifically, we evaluate the early effects of an ongoing World Bank project. By employing remote sensing and machine learning methods, we identify rural road upgrades that took place between 2018 and 2021. Using a differences-in-differences approach, we find that areas in immediate proximity to roads that received an upgrade did not experience changes in agricultural output, compared to areas that did not receive an upgrade. We restrict the sample and find a significant increase in agricultural output, although not robust. Future research should consider the medium- and long-term impact of rural road upgrades for the complete picture to emerge. While sole dependence on remote sensing data remains a challenge in economics, it is a promising avenue for future research, particularly in contexts where comprehensive survey data is lacking.

Keywords: remote sensing, rural road improvements, agricultural output, Mozambique, economic development

JEL: O13, O18, R12

Supervisor	Anders Olofsgård
Date submitted	May 16, 2022
Date examined	May 23, 2022
Discussants	Benoît Bourban & Amelie Klaus
Examiner	Karl Wärneryd

Acknowledgements

We acknowledge Anna Tompsett and Tillmann von Carnap-Bornheim at the Insitute for International Economic Studies at Stockholm University who introduced us to the field of geospatital economics, provided us with data from their colleagues at World Bank, and freely gave of their time and valuable guidance.

We are grateful to our supervisor Anders Olofsgård, who read each working draft and provided us with thoughtful feedback.

We also thank Emelie Fröberg for her feedback on our machine learning method.

Finally, we thank our friends and family for their continuous love and support.

Code Repository

All our code, as well as instructions on how to implement our road upgrade detection algorithm, is publicly available and found in the following Github repository:

https://github.com/jeffrey-clark/road_upgrades

Contents

1 Introduction

> *Good roads, canals, and navigable rivers, by diminishing the expense of carriage, put the remote parts of the country more nearly upon a level with those in the neighbourhood of the town. They are upon that account the greatest of all improvements.*
>
> Adam Smith, *The Wealth of Nations*

Investments in transport infrastructure are important for economic development. Historically, the introduction of large transportation networks in Western Europe, Japan and the United States, happened simultaneously as these regions and countries experienced periods of rapid economic development (Banerjee et al. 2020). Certainly, the direction of causality is difficult to discern, mainly due to endogeneity in infrastructure placement. Nevertheless, the topic warrants attention, as low-income countries across the globe contemplate strategies to alleviate poverty and increase economic output. While the idea of providing people with the opportunity to exit subsistence agriculture via increased agricultural commercialization has gained traction as a means to this end (see e.g. World Bank n.d.), rural residents of the world face obstacles in the form of transportation costs to participate in market exchanges (United Nations 2014; Agrilinks 2022). At the same time, traditional subsistence agriculture plays an important role in some communities, and can be sustainable and climate friendly (United Nations 2022b).

Rural residents in regions that are less economically developed face limited access to public services, goods and labor markets. In addition to increasing road safety, it is commonly understood that rural road upgrades bring about various access-related benefits to rural residents (Iimi et al. 2015, p. 5). Not only do rural road upgrades increase market access, but also access to healthcare, finance, education, and other important institutions (Yoshino et al. 2018, p. 191). Since 1996, Mozambique has been one of the fastest growing countries in sub-Saharan Africa (Utrikespolitiska institutet [The Swedish Institute of International Affairs] 2022). The discovery of gas and coal increased foreign attention in recent years, and foreign investments in mining and industries contributed to the country's rapid economic growth (ibid.). However, investments in transport infrastructure have been lagging, or at least they have not accommodated foreign investors' demands (ibid.). If sustained, such an infrastructure gap could become a substantial hindrance to the country's continued economic growth. Although the country has experienced economic growth, the distribution of wealth has become more unequal, job opportunities have not flourished, and other sectors of the economy have not seen the same development (ibid.). One such sector is the agricultural sector, employing 80% of the population and comprising 25% of GDP (ibid.; USAID 2022).

The Nampula and Zambézia provinces of the northern[1] Mozambique, jointly make up one of the poorest[2] regions in the world (the northern and rural populations comprise the majority of the poor; see Utrikespolitiska institutet [The Swedish Institute of International Affairs] 2022). Agricultural productivity in this region has proven to be highly vulnerable to reoccurring climate shocks, exacerbating rural poverty (Baez et al. 2019). The roads in this region are primarily dirt roads. As such, road upgrades take the form of leveling or paving. In 2018, the World Bank approved a grant of 150 million USD to the government of Mozambique to finance rural road upgrades (hereafter referred to as "the World Bank project") in the Nampula and Zambézia provinces. The combined population in the two provinces amounts to about 11.75 million people in 2020, according to Statista (2021). The investment was motivated with a theory that improved roads will increase market access, stimulate agricultural intensification, and ultimately raise incomes for subsistence farmers, lifting them out of poverty (World Bank, n.d.). There are a number of studies investigating rural market accessibility, market participation and spatial economic distribution in sub-Saharan Africa theoretically or empirically (see e.g. Aggarwal et al. 2018; Olwande et al. 2015 Mather et al. 2013), although few empirical studies consider the event of improved rural infrastructure and the economic implications (Jedwab and Storeygard 2021 do this, but not only for explicitly rural contexts; Casaburi et al. 2013 do this in Sierra Leone).

By leveraging remote sensing, our research paper aims to overcome the high costs of data collection in less accessible regions. We contribute to the field by furthering the usage of emerging techniques to collect data for quantitative assessment such that other researchers may learn from our experiences, and by constructing a road upgrade detection algorithm. Our models are tailored to publicly available data sources and can be applied to any region in the world, which is also a contribution on our side. Furthermore, we tap into an existing literature gap as we attempt to estimate the effect of rural road upgrades on local agricultural intensification in Mozambique, where agricultural intensification is defined as in increase in agricultural output. Do rural road upgrades affect smallholder agricultural output? We limit our study by asking the following research question: do rural road upgrades, within the frames of the World Bank project, lead to increased agricultural output in the short term? We measure agricultural output at the grid cell[3] level and only consider areas located in immediate proximity to the roads participating in the World Bank project. Rural road upgrades are identified on a yearly basis, and our years of study are 2018-2022.

In order to study the short term effects of rural road upgrades on agricultural output, we propose an approach that uses panel data from 2018 to 2022 on the changes in

[1]Sometimes, Zambézia is referred to as a central province, see e.g. section 6.4.

[2]Both multidimensional and monetary poverty is high in Mozambique (United Nations 2021).

[3]We construct our own grid cells by considering areas adjacent to roads, that are large enough to theoretically capture agricultural land.

the normalized difference vegetation index (NDVI, defined in section 5.3) as well as associated average precipitation in a difference-in-differences framework. More specifically, we create a measure of agricultural NDVI, defined as the difference in NDVI between the growing and planting season (see section 6.4). Exploiting the World Bank project, we utilize information on all the roads participating in the project across the provinces of Nampula and Zambézia. This allows us to construct three sets of (overlapping) treatment and control groups, where our leveraging of roads' individual prioritization score and status as "selected" or "non-selected" for an upgrade enables us to investigate observations from theoretically more comparable groups in two restricted samples. Our findings differ between the primary sample that encompasses all upgraded and non-upgraded roads, and the secondary samples that restrict the treatment and control groups according to "selected for upgrade" status and prioritization score. In essence, our findings from the main regressions using the primary sample suggest that rural road upgrades cause no short-term effects on agricultural NDVI, whereas our findings from the main regressions using the secondary samples display economically as well as statistically significant short-term effects on agricultural NDVI from rural road upgrades. The former finding is robust to our sensitivity analysis, while the latter is not – as such, the findings from our three different (sub)samples corroborate each other. All in all, our results suggest that there are no short-term effects on agricultural NDVI from rural road upgrades in Nampula and Zambézia. While it is possible that our study is set out too early for us to detect any agricultural output effects, we also discuss changing identification strategy such that the parallel trends assumption may be cast aside.

Our thesis is organized as follows: section 2 presents a review of previous research and literature, section 3 introduces some background information about smallholder farmers in Mozambique, section 4 presents the theoretical framework, section 5 describes our data collection process. Section 6 details our econometric specification and considerations, section 7 presents descriptive statistics and our findings, section 8 discusses the interpretation of our results, and section 9 concludes.

2 Previous Research

While the potential effects of rural road upgrades are vast, they are highly dependent on contextual factors. For example, in terms of economic outcomes the Iimi et al. (2015, p. 2) specifically points out that the effects of road upgrades on rural firms' profitability can depend on the level of motorization. Despite broad support for the claim that road infrastructure plays an important role in economic development, empirical evidence does not conclusively show this; just as economic benefits may manifest themselves in different ways depending on context, they may not materialize at all (Peng and Chen 2021). Arguably, economic effects brought about by (rural) road upgrades are heterogeneous with respect to the particular context at hand.

The following review of previous literature will first present findings of the role of transport infrastructure in economic development, and subsequently a deeper review of the rural accessibility literature and the specificity of the Mozambican context. Finally, it concludes by identifying an existing literature gap.

2.1 Transport infrastructure in economic development

As Glaeser and Kohlhase (2004) and Redding (2010) suggest, a "new regional economics" or "new economic geography" literature emerged with Krugman's seminal 1991 JPE paper (Krugman 1991). Maintaining traditional firm location literature and central place theory, new regional economic models are underpinned by the theoretical assumption of transportation costs (Krugman assume "iceberg costs", which build on previous work by Samuelson; see ibid., p. 489), such that transportation costs play a central role in understanding the benefits of proximity in economic exchanges (Glaeser and Kohlhase 2004). In this emerging literature, understanding both regional and urban development becomes a question of understanding shifting transportation costs for people and for goods.

Important empirical work has focused on studying the economic development effects of major transportation network introductions in countries such as the United States (Chandra and Thompson 2000; Donaldson and Hornbeck 2016; Michaels 2008), China (Banerjee et al. 2020; Baum-Snow et al. 2020; Faber 2014, India (Alder 2019; Donaldson 2018), and Sweden (Lindgren et al. 2021). This research has reached various conclusions. Chandra and Thompson (2000) operate from the assumption that investments in infrastructure (highways specifically) may encourage specific industries of the region to gain a competitive advantage over other regions and thereby increase these specific industries' performances (ibid., p. 461). A new highway is envisioned as bringing about reductions in transportation costs (ibid., p. 463). They do find that highway interstate construction in the United

States positively affects earnings in some industries in which transportation costs decline, and that the effect is ambiguous for other industries (Chandra and Thompson 2000, pp. 471–479). The authors also show that highways can change the spatial allocation of economic activity by increasing it in counties where they are built, and drawing it from adjacent counties where they do not pass through (ibid., p. 487). Michaels (2008) tests the implications of a Heckscher-Ohlin model of trade on the construction of interstate highways in the United States. The author finds that trade brought about by highways in rural counties increased the relative demand for the abundant factor of production, skilled workers, and reduced it where skill was more scarce – which is in line with predictions of the model (ibid.). Banerjee et al. (2020) study the impact of transportation infrastructure in regional economic outcomes in China, and depart from the assumption that distance influences the mobility of goods and factors of production. These assumptions are of particular importance in the Chinese context, as in the post-1976 era, mobility of unskilled labor was largely restricted in the country and mobility of rural capital was low (ibid., pp. 8–9). The authors find that transportation infrastructure increased GDP levels of sectors, but did not contribute to GDP growth, consistent with their model (ibid.). In his 2019 study, Alder shows that the improved highway infrastructure between major economic cities in India increased (albeit unequally among districts) aggregate economic output (as measured by night light density) via increased market access, where the Eaton and Kortum (2002) inspired Ricardian trade model prediction of the estimate is not rejected by the empirical analysis (Alder 2019, pp. 22–28). Donaldson and Hornbeck (2016) find that the provision of railroads in the United States increased agricultural land value, consistent with their Eaton and Kortum (2002) inspired model. Donaldson (2018) finds that the provision of railroads in India decreased transaction costs and enabled regions to exploit a Ricardian comparative advantage, increasing real income and also supporting the Eaton and Kortum model. This may be juxtaposed to Faber (2014) who finds that Chinese highway construction reduces industrial output in peripheral regions, which corroborates the notion of a core-peripheral integration of trade (for the model, see Krugman 1991). Baum-Snow et al. (2020) borrow from the New Economic Geography (NEG) literature and from Eaton and Kortum (2002) inspired models (such as the one used in Donaldson and Hornbeck 2016) alike in the formulation of market access index, the key difference in this case being that the former assumes limited population mobility whereas the latter assumes perfect population mobility (Baum-Snow et al. 2020). The authors suggest that increased accessibility by road infrastructure in China may increase prefectural specialization, as well as create a distinctive pattern of economic "winners and losers" in more urban versus hinterland prefectures (ibid.). They additionally find that increased domestic market access has a negative effect on prefecture GDP and wages, for which they cannot draw support from an Eaton and Kortum (2002) nor NEG model (Baum-Snow et al. 2020). Finally, a study by Lindgren et al. (2021) looks at the causal effect of transport infrastructure on economic development in Sweden using historical data. The authors find very strong effects on local economic development from the introduction of railways, with real income sharply increasing and population and agricultural land value also increasing but in much smaller magnitude (Lindgren et al. 2021, pp. 21–22).

The survey by Redding and Turner (2015) advances that the empirical body on the role of transportation costs in spatial economic development can be summarized quite readily. Perhaps unsurprisingly, they argue that the type of transportation mode matters for the particular outcome variable of interest, that different transportation modes are not interchangeable, and that institutions matter (ibid., p. 3).

A prominent feature in this field of study is the endogeneity issue of estimating effects of transportation infrastructure on economic development. The challenge posed by the fact that infrastructure placement is not random is oftentimes discussed (see e.g. Banerjee et al. 2020; Lindgren et al. 2021), meaning that much emphasis is put on identifying an appropriate identification strategy. In fact, Lindgren et al. argue that overarching conclusions about railway infrastructure and economic development are premature, as much research has used the two-way fixed effects estimator that can be subject to substantial bias in cases of heterogeneity in treatment (2021). Specifically, the studies by Donaldson (2018) and Donaldson and Hornbeck (2016) are not robust to heterogeneity in treatment, according to Lindgren et al. (2021).

2.2 Rural accessibility

Moving on to transport infrastructure in an explicitly rural context, there is a growing body of research that estimates effects on economic development for rural households following rural access improvements (see e.g. Asher and Novosad 2020; Shamdasani 2021; Yoshino et al. 2018). This section starts by introducing studies that have looked at rural accessibility, as well as the distribution of spatial economic activity. Subsequently, we introduce different ways that agricultural decisions is affected by increased rural accessibility. Firstly, there is empirical evidence of shifting agricultural decisions as a consequence of increased accessibility, where farming may become increasingly commercial, and/or modern-day input uptake may increase. Secondly, there is empirical evidence of a reallocation of labor out of agriculture and towards other sectors of employment for rural residents following increased rural access. The reallocation of labor out of agriculture may take place simultaneously as agricultural practices shift towards becoming more commercial, as shown in Shamdasani (2021). After this, the section presents a study arguing that reducing transaction costs alone may not be enough for smallholder farmers to transition out of poverty, which is supported by the studies presented in section 2.3. We additionally present two studies that consider (rural) road improvements and economic development. Finally, the section presents a study that brings yet another perspective on structural transformation in Sub-Saharan Africa, that diligently uses remote sensing techniques to arrive at its conclusion.

While the estimated effects from road upgrades differ across studies, the common denominator among these is to depart from the assumption that road upgrades reduce transportation costs, just as theorized in the "new regional economics" literature (for instance, see Aggarwal et al. 2018; Asher and Novosad 2020; Shamdasani 2021; Yoshino et al. 2018). Researchers have also shown effort to estimate this mechanism. Aggarwal et al. (2018) quantify transportation costs in rural Tanzania using a quantitative model and a combination of distance-to-market calculations, survey data collected from transportation operators, and enumerator documented road quality and travel times. Indeed, the quality of road is important for determining spatial access, and as the authors point out, little of the road network in Tanzania is paved (ibid., pp. 5–6). Their model, calibrated to northern Tanzania, suggests that there is an economically significant spatial heterogeneity in input and output prices faced by farmers, and that farmers with better market access are more likely to use fertilizer and sell maize than farmers with less favorable market access (ibid., pp. 2, 11–13). Transportation costs are central in the 2003 study by Fafchamps and Shilpi, where authors highlight the relationship between the extent of geographical isolation and spatial division of labor in Nepal using a von Thünen inspired model[4]. They find that households in proximity to cities (and markets) are increasingly likely to be engaged in non-farm work, whereas the most remote households practice subsistence agriculture for crops consumption and diversify income by selling livestock (Fafchamps and Shilpi 2003, pp. 24–25). Households in between these "extremes" engage in agricultural production for the market; perishable agricultural goods are produced closer to the market in question, whereas storable goods are produced further away (ibid., pp. 24–25). While their study does not look at the event of increased rural accessibility, it informs us on the role of households within the local economy, as a function of their proximity to the city. Nepal is an interesting case study, as rural accessibility can be highly restricted, leading to isolation from urban centers for some – residents of the Himalaya, for instance (ibid., p. 24).

Suri (2011) suggests that hybrid seed adoption is lowest among the Kenyan farmers who face the highest expected gross returns from acquiring hybrid seeds, because net returns are not positive. This is because there is heterogeneity in costs of purchasing the seeds, and high cost is correlated with high gross return (ibid.). High costs are down to poor infrastructure, although Suri focuses on access to inputs using various distance variables (as opposed to using travel time and road quality) (ibid., pp. 200–202). The introduction of spatial heterogeneity as an explanation for hybrid seed adoption stands in contrast to some of the leading preceding work on the topic, that

[4]von Thünen specifically argued that agricultural goods with high transportation costs would be produced close to the city, whereas agricultural goods with lower transportation costs would be produced further away. As such, economic activity would display concentric circles around cities (Fafchamps and Shilpi 2003).

had focused on treating the decision to adopt the input in a learning framework (Suri 2011, p. 163). Of course, learning about the usage of hybrid seeds may be facilitated by access to markets and information exchanges.

There is empirical evidence to suggest that road upgrades improve the economic conditions of rural households. In their study of feeder road improvement in Papua New Guinea edited by Yoshino et al. (2018), the authors employ correlated random effect estimates to a previous study by Gibson and Rozelle that used the year of entry by the national highway network into a region as an instrument for travel times on feeder roads to the nearest highway (ibid., p. 215: see the authors' presentation of Gibson and Rozelle, 2003). It is shown that road sealing increases average consumption of households by 0.55% (ibid., pp. 217–218). They furthermore suggest that road sealing has a positive effect on households' market participation and transformation towards commercial farming activities, although the estimates are not statistically significant (ibid., pp. 219–221). Jacoby (2000) finds that the provision of access to markets by rural roads in Nepal would bring about substantial benefits, with much of these going to poor households, although with persisting income inequality (ibid., pp. 735–736). The author views agricultural land as an asset: its net present value is then the discounted stream of maximal profits obtained from cultivation (ibid., pp. 714–716). Where the asset is located such that transportation costs to markets are lower (measured in hours of travel time), this is reflected in its value (ibid.). The data indicates a negative relationship between input use (chemical fertilizer) and distance to market, which means that there is a relationship between access and the intensive margin of cultivation (ibid., p. 724). Using ordinary least squares and a semiparametric instrumental variable set-up, Jacoby shows that travel-time to markets is negatively associated with agricultural land value (ibid., pp. 728–729). One caveat is that the model assumes prices across markets are fixed; indeed, this is a plausible assumption where the representative rural village is small compared to markets, but perhaps not if the general equilibrium is expected to shift following a rural road access program. The findings by Yoshino et al. (2018) and Jacoby (2000) suggest that road upgrades may mediate positive rural household market participation and increase agricultural productivity.

Some empirical evidence suggests that increasing rural access leads to a reallocation of labor out of agriculture, due to increased labor mobility: the possibility for rural residents to participate in labor markets located in more densely populated areas, such as towns (see e.g. Shamdasani 2021, p. 4). This phenomenon is recorded in Shamdasani (ibid., pp. 8–10) and Asher and Novosad (2020, pp. 807–818), who study the same road construction program in India, namely the Pradhan Mantri Gram Sadak Yojana (PMGSY). Shamdasani (2021) employs a differences-in-differences strategy and exploits the staggered implementation of the PMGSY to identify causal impacts of road construction on agricultural decisions. She finds evidence of diverging effects of treatment along the spatial dimension, where households in less remote areas and within commuting distance to towns see an increased exit from agriculture, compared to more remote households that do not see an increased exit from agriculture (Shamdasani 2021, pp. 8–10). Asher and Novosad (2020) use a regression discontinuity approach that takes advantage of threshold levels on village population for enrolling in the road construction program to study economic outcomes of the PMGSY. Interestingly, the authors use remote sensing data to proxy agricultural production (vegetation measures NDVI and EVI) on village level (Asher and Novosad 2020, p. 803). They find that rural road construction increases the provision of transportation services, and reallocation of labor out of agriculture (ibid., pp. 809–813). The authors suggest that the effect is strongest for rural residents holding no or little agricultural land, although this heterogeneity cannot be justified statistically (ibid., p. 803). Furthermore, the authors conclude that they find no substantial impact on agricultural yields (agricultural production) for villages that receive new roads – standard errors of the coefficient of interest are large compared to the estimate itself (ibid., p. 815).

Moneke (2020) studies the impact of "big push" infrastructure investments, in this case the extension of electrification and improvements in the road network in Ethiopia. The author estimates the separate effects of road improvements on sectoral employment and compares the results to the estimated effects from the coupled investments (ibid.). It is an approach that acknowledges that there might be interactions (or complementarities) between different types of infrastructure investments, and interactions' potential importance in economic development. To overcome the endogeneity issue, Moneke uses an instrumental variable design and finds that road infrastructure improvements only has a statistically and economically significant effect on manufacturing employment and service employment shares, albeit in different directions (positive for services, negative for manufacturing), while the coefficient on employment in agriculture is statistically insignificant on conventional levels but negative (ibid., pp. 30, 62). When estimating the effect of the interaction between road improvements and extension of electrification on sectoral employment, the author finds a large, negative impact on agricultural employment, statistically significant at the 10% level (ibid., p. 62).

Taken together, these studies point towards a reallocation of labor out of agriculture as an effect of improved access to labor markets, or a combination of improved access and other infrastructural developments. Still, Shamdasani (2021) finds that improved road connectivity in a statistically significant way increases households' adoption of modern agricultural inputs such as hybrid seeds, leads to increases in labor hiring at farms, and increased market participation among rural farmers (ibid., pp. 11–12). Indeed, if rural farmers are credit constrained, sending a family member to work in the city could very well facilitate the household's agricultural production. At the same time, Asher and Novosad (2020) recorded a small and statistically insignificant effect on agricultural output in their study of the PMGSY, which could suggest that agricultural production was not altered to a great extent even as agricultural decisions became more commercial. Still, less costly transportation of labor and goods may channel both the reallocation of labor out of agriculture and the commercialization of agricultural activities, in a simultaneous manner. Finally, we note that reallocation of labor out of agriculture coupled with increased uptake of modern-day inputs in the agricultural sector, are patterns that could be observed in structural transformation, where the agricultural sector increases its productivity but decreases in terms of importance for employment (see e.g. Campenhout 2015, p. 1), although none of the above studies claim to show such a transformation.

Olwande et al. (2015) use panel data to estimate output supply functions for smallholder farmers in Kenya, a country where, similarly to Mozambique, smallholder farmers make up a large share of the population and produce a large part of total output; 75% of agricultural output in Kenya is produced on 0.2.-3 hectares of farmland (ibid., p. 22). Farming is generally characterized by low productivity and input usage, and can be described as “semi-subsistence” (ibid., p. 22), which means that the features of Kenyan agriculture as described in Olwande shares traits with Mozambican smallholder agriculture (see section 3). Using panel data and Cragg’s 1971 double hurdle model, the authors study market participation and quantity sold of maize, kale and milk (ibid., p. 28). Interesting to us, are their findings that land holdings are positively associated with the probability of selling maize, and that increases in land holdings are associated with increases in quantities sold of maize, milk and kale (ibid., p. 28). Also interesting is their finding that cumulative adoption rates of hybrid seeds and fertilizer are (“strongly”) associated with increases in market participation (ibid., p. 28). Precipitation matters too for maize marketing decisions; the highest quintile in maize market participation as well as quantity sold live in areas of high precipitation, compared to the lowest quantile (ibid., pp. 30–31). With policy responses in mind, Olwande et al. advance that reducing transaction costs for smallholder farmers is not enough to encourage market participation as a mechanism to transition out of poverty, since landholdings (as well as owning assets and technology) are important determinants of market participation (ibid., p. 31). Indeed, low productivity might persist in absence of coupled interventions, and income-boosting effects (such as increased commercialization of agriculture) could be prone to be unevenly distributed among the ones who own more and less of land, assets, and technology.

Jedwab and Storeygard study road upgrades over five decades, 1960-2010, in 39 sub-Saharan countries (2021). More specifically, the authors consider the effect from road upgrades on population growth in cities using an instrumental variables approach (Jedwab and Storeygard 2021). Their estimates suggest that on average, over the thirty years after a change in market access, a 10% increase in market access increases city population by 0.8-1.3% (ibid., p. 21). Their road dataset is constructed by drawing from the United States Government’s Digital Charts of the World and Michelin Road maps (with the maps indicating road categories such as "highways", "paved roads", "improved roads" and "dirt roads"). Rural-urban migration is suggested to be one of the main drivers of the found population increase

(Jedwab and Storeygard 2021). The authors also consider whether roads have an impact on economic activity, proxied by night light density, and find that while population effects take time to manifest, positive effects on night lights are quicker and arrive in the first decade after an improvement in market access (ibid., pp. 1, 27–28). Indeed, their study suggests that roads matter for city population and economic activity, in particular the latter. Another study that looks at the effect of roads on economic development, is Casaburi et al., who study the effects of a rural road upgrade program in Sierra Leone on market prices of staple crops (2013). Their findings, obtained via a regression discontinuity design, suggest that rural road upgrades decrease transportation costs and decrease the market price of crops in rural markets, with the strongest price decreases materializing furthest away from the city (Casaburi et al. 2013, pp. 23–26). Their interpretation is that rural road upgrades mainly facilitated access to markets for producers, i.e. farmers (ibid., p. 23).

Finally, Peng and Chen (2021) study the impact of increased market access (using travel time) in Zambia on air pollution, night time light and green space. The authors train an AI to detect the condition of roads ("no road", "dirt road", and "paved road"), and to detect changes in road condition (road improvements) between the years 2009 and 2019 (ibid., pp. 11–12). They also use an AI to detect built-up (new buildings) (ibid., pp. 9–10). The authors furthermore use aerosol optical depth (AOD) measures to detect air pollution and NDVI measures to detect green space (ibid., pp. 9–10). Indeed, Peng and Chen pioneer the usage of remote sensing techniques and AI methods to obtain data for econometric analysis. The authors find that increased market access does not have an impact on income (measured by night light density) except for in primate cities where it decreases slightly, and they interestingly show that increased market access increases built-up in urban areas, leads to increases in air pollution and to increases in deforestation (ibid., pp. 28–31). This leads to their conclusion that increased market access may encourage migration and urbanization without the provision of sufficient employment opportunities within urbanized areas, much in line with the “urbanization without growth” phenomenon that has been studied in Sub-Saharan Africa (ibid., p. 1; see Vernon Henderson and Kriticos 2018, for urbanization in sub-Saharan Africa). Peng and Chen mask out cropland using Global Land Cover 2009 data from the European Space Agency, and then isolate the NDVI differences to areas classified as forests or non-crop land, which leads to their conclusion that increased market access exacerbates deforestation in areas that are not suitable for crops (Peng and Chen 2021, p. 22).

The findings by Peng and Chen support the features described in the “urbanization without growth” literature, namely urbanization in spite of a lack of the traditional notions of “push” and “pull” factors. If indeed such dynamics are at work, it is not evident that allocation of labor out of agriculture as a consequence of improved rural market access, for instance, necessitates urban employment opportunities. In fact, urbanization without growth rather suggests that an allocation of labor out of agriculture could happen in spite of better employment opportunities. Whereas rural road upgrades could provide opportunities to improve agricultural productivity and production as well as to reach better labor markets, they could do neither. As such, it is far from evident that smallholder farmers face a choice between whether to commercialize their agricultural practices versus whether to look for employment in the city, after rural roads have been upgraded. Nonetheless, migration to the city could occur even if employment opportunities are not readily provided.

2.3 Studies of rural economic development set in Mozambique

As suggested by previous research, increased rural accessibility may play an important role in increasing rural economic development. Transportation costs and market access appear to matter substantially for the spatial distribution of labor and the spatial economic activity. However, they are not the only factors that matter. Research by Mather et al. (2013) set out on Kenya, Zambia, and Mozambique, shows that investments that raise land-access and farm-level productivity are important too and might need to be coupled with investments in transportation infrastructure. The authors specifically argue that for Mozambique, it is questionable whether smallholder maize sales will increase unless these other investments are realized (ibid., p. 264). For instance, the authors discuss targeted recommendations about optimal input use and dosage along with public efforts to eliminate "disease constraints" to animal traction (ibid., p. 262). A similar argument, that transportation infrastructure improvements are not enough, is advanced by Boughton et al. (2007), who examine the relationship between asset holdings and market participation in Mozambique. Although Boughton et al. do not claim to estimate a causal relationship, they show that households holding more assets, such as land and livestock, are more likely to participate in the crop market (ibid., p. 66). While the distance to the nearest paved road is positively correlated with cash crop market participation, household assets are even more so (ibid., pp. 54–55). In contrast, the correlation between the distance to the nearest paved road and maize market participation is statistically weaker and economically less meaningful (ibid., pp. 54–55). The authors' interpretation of their probit and OLS regressions is that not all households afford to participate in the market (ibid., p. 36). This is echoed in Barrett (2008), who suggests that substantial barriers to entry exist in staple food grain markets across eastern and southern Africa.

Heltberg and Tarp (2002) study farmer's market participation in Mozambique. They theorize that subsistence farming keeps transaction costs low and mitigates the risk of fluctuating food prices – an insurance that can be achieved at the expense of lower output and sales volume (ibid., pp. 107–108). Heltberg and Tarp identify several obstacles to smallholders' market engagement, including substantial transaction costs, lack of agricultural assets, and presence of a risky, low-yielding environment (Heltberg and Tarp 2002, p. 122). Chiovelli et al. (2018) estimate the effect of demining on economic activity in Mozambique. The authors use night light density as a proxy for economic activity and find the effect to be heterogeneous between urban and rural localities, suggesting that the economic payoffs for urban localities are greater than that for their rural and remote counterparts (ibid., pp. 24–25). One caveat with the study is that night lights as a proxy for economic activity might not be reliable in rural contexts (Yoshino et al. 2018, p. 208). Possibly, if there are positive effects from increased transportation access on economic activity also for the rural population, this might not be captured by night lights data.

Taking these studies together, previous research set in Mozambique indicates that there are reasons to question improvements in (rural) transportation infrastructure to mediate households' transition out of poverty. In short, farmers face several obstacles that hinder them from participating in market exchanges, and transportation costs are as such not the only one.

2.4 The literature gap

Remote sensing data and tools are becoming increasingly available, and there is a growing literature looking to reap the benefits of less costly data collection and (potentially) more accurate data. At the end of Yoshino et al. (ibid.), one is prompted to contemplate the usage of publicly available satellite imagery for ground-level developments in road networks (p. 230-231). The main advantage of publicly available satellite imagery is the availability, whereas the resolution is typically lower than that of commercial satellite imagery.

This paper contributes to existing literature by proposing a new way to leverage publicly available satellite imagery for econometric analyses, potentially opening up for new opportunities for economic research – especially in regions of the world where other data is expensive or unavailable. To our best knowledge, there is no other paper in this field that detects the timing of rural road upgrades, using publicly available satellite imagery, and shares the data generating algorithm with fellow researchers.

Another significant contribution of this paper is the furthering of our understanding of the impact of rural road upgrades on agricultural output in Mozambique. Most relevant previous literature was published over a decade ago, and answered different research questions; while the papers set in Mozambique touched upon different access-related effects on local economic development, they did so without estimating or controlling for the effect of road upgrades. One potential reason for this is that prior to this paper there has not existed any accurate, large-scale data of the completion of rural road upgrades. With the data of road upgrades now available, this paper directly estimates the impact of rural road upgrades on agricultural output. By leveraging remote sensing and machine learning methods, we estimate agricultural effects at scale.

3 Background: Smallholder Farmers in Mozambique

This section attempts to describe the economic realities that smallholder farmers in Mozambique face. We are greatly indebted to the Mozambique CGAP Smallholder Household Survey Report 2016, from which we draw most of the qualitative data reproduced in this section. It should be noted that the CPAG survey encompasses smallholder households throughout Mozambique, i.e. it attempts to capture data about Mozambican smallholder farmers in general and it does not contain geo-coded data. As such, there may be important heterogeneity in answers from farmers in different parts of the country that does not surface in the data. Nevertheless, we use the CGAP survey as a proxy for understanding the economic realities of smallholder farmers of Nampula and Zambézia.

Agriculture is a vital part of the economy of Mozambique, as it contributes to about 25 % of the country's GDP (USAID 2022). About 80 percent of the population is employed in the agricultural sector (CGAP 2016, p. 1; International Potash Institute 2011 claims that about 80 to 90 % of the Mozambican population is involved in agriculture, and largely in subsistence agriculture taking place on less than three hectares; according to the USAID 2022 about 80 % of population works in agriculture). The large majority of the sector is smallholder farmers (ibid.). Furthermore, agriculture in Mozambique is characterized by low production and productivity, in part due to low accessibility to quality agricultural inputs (ibid.). Smallholder agriculture is mainly rain-fed (CGAP 2016, p. 1). For smallholder farmers, agriculture constitutes not only income but livelihood and consumption, which is illustrated by the following figures: 89% of smallholder farmers report consuming their crops, 62% report selling their crops, and 33% report trading their crops (ibid., pp. 15–16). The CGAP report suggests that farmers are committed to and enjoy agriculture and take pride in their work, but that if given the opportunity, many farmers would opt out of (or diversify out of) agriculture (ibid., p. 65).

3.1 Farmers' access to markets

Mozambican smallholder farmers lack access to markets due to limitations in transportation (ibid., p. 3). Farmers are aware that they would be likely to receive higher prices for their goods if they had access to the market (ibid., p. 3). When asked whether they receive the market price for their outputs (crops and livestock), 45% of respondents report selling at the market price, whereas 35% report not getting the market price and 20% answer that they do not know whether they get the market price (CGAP 2016, p. 38). The three top reasons respondents note for not getting the market price are: too few customers (57%), no access to transports to other markets (39%), and that customers take advantage of them (39%) (ibid., p. 38). The two top reasons farmers sell their crops at the market they do, are that they do not have access to transport to other markets (57%) and that they get the best prices at that market (43%) (ibid., p. 38). Indeed, many respondents report selling at a local market (73%), or in the village (72%) (ibid., p. 38). In sum, smallholder farmers lack access to larger markets and as such resort to selling their output in smaller, more local markets where they generally do not receive the market price for their crops and livestock, and where customers "take advantage of them" (ibid., p. 38). Our interpretation is that the local markets where farmers sell exhibit monopsonistic features.

3.2 Farmers' access to inputs

Furthermore, smallholder farmers experience a low level of financial inclusion and as such have a large appetite for financial mechanisms that would enable them to acquire modern-day farming inputs, such as fertilizer, in order to sustain their homesteads (ibid., p. 3). 74% of respondents report buying inputs of some sort already, such as fertilizer (ibid., p. 36-37) and 35% of respondents say that input purchase would be a main reason for borrowing money (ibid., p. 80). Further underlining the importance of inputs and their financing, 47% report that a payment plan for inputs is very important and 46% report that a savings plan for inputs is very important (ibid., p. 80-81). When asked about whether they currently had a payment plan for inputs, 11% report that they currently had it while 48% report that they wanted it (ibid., p. 81). Similarly, 11% report that they currently had a savings plan for inputs while 49% report wanting it (ibid., p. 81). It could also be noted that while formal lending practices are almost non-existing among rural farmers, informal practices are not common either (ibid., p. 48). In sum, we deem it fair to say that the report suggests that farmers are credit constrained with respect to optimizing their input purchases.

3.3 Farmers' future prospects

When asked about their future prospects, the majority of farmers report that they intend to keep on working in agriculture (ibid., p. 19). Nevertheless, it should be noted that the share of respondents planning to keep on working in agriculture falls with decreasing years of agricultural experience (ibid., p. 19). Just over 70 percent of respondents want to expand their agricultural activities (73%), although nearly the same fraction of respondents (71%) claim that they would take on full-time employment if they were offered a job (CGAP 2016, p. 20). Just as pointed out by the CGAP report, what could appear as conflicting responses rather suggest the difficulties of making ends meet on the subsistence farm (ibid., p. 20). These answers shed light on the need for farmers to diligently adjust their life plan to changing circumstances in order to support their homesteads (ibid., p. 19-21). Indeed, many farmers would leave agriculture if a steady wage could be secured elsewhere.

3.4 Size of agricultural holdings

34% of respondents report having an agricultural holding of less than 1 hectare (100x100 m) (ibid., p. 7). 37% report having agricultural holdings of between 1 and 2 hectares, 25% report having between 2 and 5 hectares, while 3% report holding more than 5 hectares (ibid., p. 7). The vast majority of agricultural holdings are thus below five hectares, which means that the grid cells (see section 5.3) that we construct are large enough to hold several households' worth of agricultural land.

Most rural households have small landholdings, which was also reported in Barrett (2008). The study by Barrett shows that the probability of selling maize in Mozambican rural households back in 2001-2002 increased sharply for households holding four or more hectares, which represented about 10% of the sample households (ibid., p. 309).

3.5 Crops

Maize is the most widely produced staple food crop in Mozambique, and it is grown in both the northern and central regions (see e.g. Campenhout 2015, p. 2). It is also one of the most important crops, alongside legumes and cassava (International Potash Institute 2011).

Maize is not an acutely perishable good, but can be stored for some time meaning that it does not need to be produced in immediate adjacency to its divestment (Olwande et al. 2015, p. 23).

4 Theoretical Framework

In this section, we present the theoretical framework. We start by noting that the impact on agricultural output from rural road upgrades is not theoretically obvious. In essence, there are arguments to say that this impact could be either positive or ambiguous. We will go through both cases separately.

Before we dive into the cases, we will illustrate some theoretical implications of transaction cost in the marketplace for smallholder farmers. Consider a good that is both produced and consumed within a household – typically a crop – for which representative households A, B, and C share an identical demand curve. There is a market price for the output good, at P_{market}, but transportation costs amounting to τ means that the good can only be sold at $P_{market} - \tau$, and that the good can only be bought at $P_{market} + \tau$. Transportation costs thus create a band around the market price (for simplicity, we assume that transportation costs are identical for all households). Household B produces the same quantity as it demands, and therefore self-selects out of the market exchange. Household A demands more than it produces, for which reason it is a net buyer. Household C produces more than it demands, which makes it a net producer.

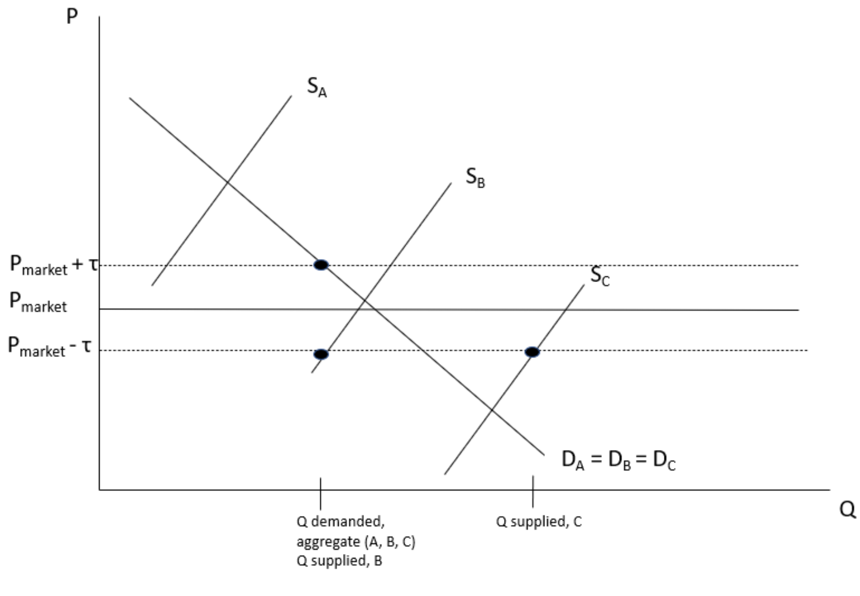

Authors' illustration of Campenhout (2015). The figure depicts supply and demand for a farmed good that is produced and consumed by households A, B, and C. P stands for price and Q for quantity. P_{market} is the market price, τ is a transaction (transportation) cost. S denotes supply curves, D is aggregate demand.

Figure 1: Supply and demand with transaction costs

The model predicts that a household will participate in the market as either net buyer or net seller. Only in the case where demand at the transportation cost adjusted purchase price equals supply at the transportation cost adjusted price, will a household self-select out of the market (household B). This is not necessarily supported in the data; there could be many households that do not participate in the market at all, just as van Campenhout notes (Campenhout 2015, p. 6). However, the model clearly illustrates that transportation costs matter for smallholder households' market participation. As an exercise, we let transportation cost τ diminish, such that the transportation cost adjusted prices approach the market price. Consider household C. Whereas a higher market price for output goods may move household C upwards on their supply curve, household C are also consumers of the good. The income effect may increase their demand for the good such that aggregate demand shifts. Where the demand curve shifts to the right, the household could end up participating in the market to a lesser extent (producing more of the good than previously, but consuming even more of the good than previously). Alternatively, a household could want to consume more leisure as income increases and as such produce less for the market.

The households' individual supply curves are also determinants of the extent of households' market participation. Instead of lowered transportation costs, consider a technological advancement that enables a farm with "few productive assets, rudimentary technology or both" (ibid., p. 6) to increase their output; household A's supply curve shifts out to the right, and in theory, the household may become a net producer. When considering the implications of changes in transportation costs and technological advancements, it is important to keep in mind that households both produce and consume the good in question. If transportation costs decrease, the intersection of $P_{market} + \tau$ happens further down the aggregate demand curve. We should expect aggregate demand to increase, unless individual household demand curves shift to the left and cause an overall shift of the aggregate demand to the left. Where an increased aggregate demand is met with an increased supply, we should see increases in agricultural output. The effect on market participation may however still be ambiguous considering the possibility of supply and demand curve shifts.

In short, it is not theoretically trivial what the effects of rural road upgrades on agricultural output in Mozambique would be. Ultimately, it could be seen as an empirical question.

4.1 Positive effects on agricultural output

In this section we will consider important arguments in favor of positive effects on agricultural output following rural road upgrades.

A road upgrade is interpreted as a reduction in transportation costs for those using the road. For a rural resident working in agriculture, the decreased cost of transportation translates into a decrease in the travel cost adjusted price of acquiring agricultural inputs from the marketplace, such as fertilizer or seeds, and an increase in travel cost adjusted price of outputs, such as crops, given that other factors remain unchanged. We will use a simple framework to demonstrate these effects.

Consider a farmer residing in area i, envisioned as a rational agent seeking to optimize their agricultural output, traveling to the marketplace j to procure inputs:

$$ptotal_i = pinput_j + travelcost_{j,i} \tag{1}$$

(equation inspired by Aggarwal et al. 2018, p. 9)

Where *ptotal* is the total price of the input, *pinput* is the price of the input at the market rate, and *travelcost* is the cost of transportation. The travel cost adjusted total price of an input for farmers living in the area i is the price of input in market j plus the cost of traveling to the market j and returning to the area i with the input acquired.

It follows that a reduction in $travelcost_{j,i}$ reduces the total cost of inputs, $ptotal_i$, holding all other factors fixed. A rational agent residing in area i looking to purchase the input will seek to minimize their cost of doing so. The minimum cost adjusted price for acquiring the input is:

$$ptotal_i = \min\{pinput_j + travelcost_{j,i}\} \tag{2}$$

(see ibid., p. 9)

Let us now turn to the output market. Consider a farmer residing in area i traveling to the marketplace j to sell a crop:

$$ptotal_i = pcrop_j - travelcost_{j,i} \tag{3}$$

(equation inspired by ibid., p. 9)

Where *ptotal* is the total price of the crop, *pcrop* is the price of the crop at the market rate, and *travelcost* is the cost of transportation. The travel cost adjusted selling price of the crop for farmers in area i is the market price of the crop at market j minus the cost of traveling from area i to market j with the crop and from market j back to area i. A reduction in cost of traveling from area i to market j, holding all other factors fixed, increases the total price of the crop. Again, a rational agent residing in area i will seek to maximize the price of their outputs. The maximum price of the crop for farmers in area i can be written as:

$$ptotal_i = \max\{pcrop_j - travelcost_{j,i}\} \tag{4}$$

(see Aggarwal et al. 2018, p. 9)

Indeed, the outcomes described above assume that prices remain fixed following the change in accessibility for remote areas. This need not be the case. For instance, if increased accessibility would induce a large influx of farmers from area i selling their outputs in their less remote, larger market j, there could be reasons to interpret this as a positive supply shift. In a general equilibrium framework, the effect on total price for input and output goods described here may be mitigated by supply and demand shifts caused by the integration of more rural areas into more urban markets. Indeed, if such shifts were to completely mitigate the positive effect on output prices and the negative effect on input prices, we would not expect a positive effect on agricultural NDVI coming from the road upgrades. While the simple framework does not accommodate supply and demand curve shifts, we use it to show the overall impact on input and output prices in the absence of large supply and demand curve shifts. In other words, unless supply and demand curve shifts completely mitigate the above described effect on output and input prices, we should expect to see an increase in NDVI.

For farmers, the input becomes output (e.g. seeds), or the input is expected to be output-enhancing (e.g. fertilizer), which means that rational farmers will take the potential of increased agricultural production and productivity into account when formulating their input and output decisions. As there is opportunity to reach better markets with improved roads, farmers may seek to reap previously unrealizable economic benefits. Some might choose to diversify their production of crops, perhaps by growing cash crops alongside other production (Yoshino et al. 2018, p. 191). The techniques we use to collect our data are not sensitive enough to distinguish between different types of vegetation, for which reason we will not be able to specifically identify whether farmers have opted for strategies to diversify their portfolio of crops. The effects we estimate, may thus be on the lower bound with respect to the value of agricultural output.

Qualitative data suggests that farmers currently purchasing some level of inputs are not acquiring optimal levels of input, but rather a lower than optimal level, as showcased by their demand for a financial mechanism that would enable them to afford inputs to sustain their homestead (see section 3.2). It is also a fact that increasing the value of current agricultural output, would generally be equal to increasing consumption and income among rural Mozambican farmers (see section 3). Taking the simple framework of decreases in transportation costs together with the qualitative data on smallholder farmers' economic lives, we infer that in general, smallholder farmers would want to take advantage of cheaper transportation costs to improve their economic situation. This could be achieved by scaling up agricultural production or increasing agricultural productivity. Even if we were to assume that farmers face limitations in how much they can affect their extensive margin, we would expect to see changes in their intensive margins. Increased agricultural productivity should not mean that farmers choose to cultivate on smaller plots of land, as many do not currently meet their own consumption demands in terms of basic necessities such as food and clothing (CGAP 2016, p. 13). The opportunity to increase agricultural output would become a way of making ends meet for many smallholder farmers, and as such, a rural road upgrade should cause positive changes in agricultural NDVI.

From the World Bank's point of view, rural road improvements should aid local economic development. Even as Mozambique experienced sustained economic growth throughout the years 2005-2015 (World Bank n.d.), rural regions still faced severe poverty, in part due to low agricultural productivity (ibid.). According to the World Bank, limited connectivity between rural and less rural areas was a contributing issue, meaning that rural residents faced limited access to credit markets, agricultural extension services and market information (ibid., p. 1). This underpins the theory of change provided by the World Bank in their launch of the current project, where they expect that rural road upgrades should, in the long term, lead to cultivation of "high value crops and profitable changes in cultivation strategies, including expansion of irrigation [and] cultivated land, consolidation of fragmented plots into larger plots, cultivation outside of the rainy season, shifts to cultivation of new crops" (ibid., p. 2). In short, they advance that rural road upgrades, such as "reconstruction or rehabilitation of bridges and culverts, graveling, surface treatment, [and] routine/periodic maintenance" in the short and intermediate run should improve road surfaces and accessibility, decrease travel time and transportation costs (ibid., p. 2). Such improvements then lead to easier market access, higher output prices combined with cheaper access to inputs, as well as increase access to rural areas for extension agents (ibid., p. 2). The modernized farming methods should in the long term, according to the World Bank, lead to increased incomes and poverty reduction (ibid., p. 2). Although this theory of change favors a positive effect on agricultural output, it considers a longer time horizon, acknowledging that impact might take time.

Using the theory explored in this section, we hypothesize that grid cells that experience a road upgrade in the World Bank project will show a higher level of agricultural output, that is an increase in agricultural NDVI, compared to grid cells that do not experience a road upgrade in the World Bank project.

4.2 Ambiguous effects on agricultural output

In this section we will consider important arguments for ambiguous effects on agricultural output following rural road upgrades. We maintain the assumption that rural road upgrades can translate into decreased costs of transportation of people and goods, but acknowledge that they do not necessarily enable transport of people and goods if other obstacles to transportation prevail. Firstly, however, we consider the possibility that improved access to labor markets outweight detectable effects on agricultural output.

While improved roads mean that there is opportunity to reach better output markets for crop sale, and to procure inputs in a less expensive way, it also means that labor markets are made less expensive to travel to. Improving rural roads could then lead to a reallocation of labor out of agriculture and into other sectors, via labor markets that were previously unfeasible to rural residents (see eg. Shamdasani 2021; Asher and Novosad 2020). Indeed, there is qualitative evidence that suggests that many smallholder farmers in Mozambique would leave agriculture for other employment opportunities to improve their economic situation if they could (see section 3.3). If rural road upgrades mediated this possibility, it is not evident that we would see an increase in agricultural NDVI as a consequence of the road upgrades. The effect would instead be ambiguous, especially as some rural residents may extend or intensify their agricultural activities while others may leave the sector altogether. As discussed in section 2, these effects could very well be spatially distributed. There is theoretical work suggesting that combined effects of reduced transportation costs and increased agricultural productivity could produce the above outcome. Gollin and Rogerson (2014) present a model of a poor economy (comprising "remote agriculture", "near agriculture" and "city manufacturing", p. 40) that displays features of subsistence farming and that has high transportation costs for goods. They calibrate their model to Uganda and show that a 10% increase in agricultural total factor productivity combined with a 10% decrease in transportation costs, leads to a 20% reduction in employment in subsistence agriculture (Gollin and Rogerson 2014, p. 47).

Improved infrastructure is a way of reducing the cost of traveling to markets where inputs such as fertilizer can be procured and outputs such as crops can be sold at the market rate. But rural farmers are generally credit constrained and face additional obstacles (see section 3.2). For example, farmers face limitations in transportation to larger markets (see section 3.1). While the infrastructure itself is generally not in place or is unreliable and of poor quality, it is possible that farmers also face a lack of transportation modes. It indeed seems plausible that financial inclusion, provision of transportation modes, and rural infrastructure improvements would go hand in hand for farmers to experience economic development via a transition out of subsistence farming towards more commercial farming activities. This is somewhat echoed in Yoshino et al. (2018). It is not theoretically evident whether people facing poverty reap the same benefits from better roads as people who have a higher income (Yoshino et al. 2018, pp. 191–192). It could be true that relatively poorer people see greater effects on consumption gains from road improvements compared to the less poor, as the less poor might be able to mitigate the negative effects of not having accessibility and thus not see as great an impact of road improvements (ibid., p. 192). It could also be the other way around: the relatively poorer people still face obstacles even after road improvements (as discussed previously), such as persisting transportation costs, and are thus unable to reap the benefits of an improved road that the relatively less poor can (Yoshino et al. 2018, p. 192). As such, it is not clear that improved rural infrastructure on its own would be enough for farmers to take up more productive agricultural practices, or to extend their agricultural production.

From these insights, our hypothesis differs from that previously posited. If effects are indeed ambiguous we expect that grid cells that experience a road upgrade in the World Bank project will show similar levels of agricultural output, i.e. we will not be able to identify an effect on agricultural NDVI, compared to grid cells that do not experience a road upgrade in the World Bank project.

4.3 When to measure effects

Effects on agricultural practices brought about by road upgrade programs might take time to manifest themselves. The current study looks at upgrades finalized in the autumn of 2021 and attempts to estimate the effects in the early spring of 2022. If indeed effects take time to manifest themselves, we would be unable to detect them in our chosen period of study. There could be reasons to believe that the effects are not instantaneous, perhaps due to adjustment times among market participants. At the same time, there could be reason to argue that effects would be quick to manifest themselves in this particular context: qualitative data on smallholder farmers in rural Mozambique highlight that most farmers find themselves in severe poverty, meaning that they live on or below the poverty line (CGAP 2016, p. 2). Because people find themselves in such conditions, they are forced to continuously adapt, to take advantage of the opportunities that arise. The need to behave in an "opportunistic" (ibid., p. 20) way could manifest itself in rapidly exploiting the upgraded roads – either to seek employment outside of agriculture, or to participate in agricultural market exchanges with lower transaction costs. Indeed, the CGAP report underlines that farmers are also constantly looking for opportunities to improve their situation (ibid., p. 13). We therefore assume that farmers will respond quickly to changes in their environment and take advantage of changes brought about by the newly built roads – if they can (as referred to previously, research suggests that barriers other than transportation infrastructure might hinder farmers' market participation; see e.g. Mather et al. 2013).

Another argument for estimating effects early on, is to say that road upgrades deteriorate with time, unless roads are maintained. Road deterioration is a gradual, more dynamic process than road upgrading. Road upgrading is comparably sharp: roads get upgraded to a certain standard within a short period of time, whereas roads deteriorate during a longer period of time to different levels of quality. Estimating effects of road upgrades "too far" away in time would thus be associated with problems of gradual road deterioration that might be more difficult to identify with the remote sensing techniques presently used – for instance, a higher resolution might be needed for surface quality assessment (Yoshino et al. 2018, discuss using sub-meter resolution imagery for determining the quality of paved roads, see p. 230). A combination of estimating effects both close in time as well as effects further away in time might thus be a promising approach for a holistic impact evaluation. It is unfeasible to us in the context at hand, as the time of writing this paper is the spring of 2022, for which reason we will seek to estimate only the short-term effects.

5 Data

Yoshino et al. (ibid.)) describe how remote sensing methods can be beneficial in contexts commonly known as "data poor" (p. 192). Peng and Chen also discuss this, arguing that applying machine learning methods to satellite data is a promising avenue for economic research in "data poor" contexts (Peng and Chen 2021, pp. 1–2). Helping other researchers discover and use new types of datasets could lower the cost of using such data, as well as contribute to increasing the attention towards countries that do not collect encompassing survey data. The idea of looking into more unconventional data sources in order to study a "data poor" region of the world resonates with the aim of this paper. While researchers such as Yoshino et al. (2018) gathered a combination of conventional and unconventional data types, we face the challenge of relying even more heavily on an unconventional data type as we use remote sensing to collect geospatial data for both dependent and independent variables.

This paper contributes with an innovative approach to detect and classify rural road upgrades. With road coordinates as the only required input, our publicly shared code can be applied to any area on the Earth. We designed our road upgrade detection algorithm with accessibility in mind and took extra measures to accurately detect upgrades on publicly available images, which have a lower resolution than commercial images.

We received a dataset of 219 roads from the World Bank, where each road is stored as a list of coordinates. Using the coordinates in Google Earth Engine (GEE) and Python, we designed three algorithms (1) road upgrade detection, (2) machine learning classification and (3) grid cell creation. Unless otherwise stated, the described method is developed by our own design and intuition.

5.1 Road upgrade detection

To detect road upgrades, we generate composite images (composites) by combining several satellite images taken between the years 2017 and 2021. These underlying satellite images come from the European Space Agency's publicly available Sentinel-2 (Sentinel-2 MultiSpectral Instrument, Level-1C) satellite imagery. Sentinel-2 imagery data is stored in 13 spectral bands, where each band contains data of the spectral reflectance signature (intensity of light reflected from the Earth's surface) from a limited range of wavelengths in the electromagnetic spectrum. For visual exploration, we utilize the Blue (B2), Green (B3), Red (B4) spectral bands, all with a 10-meter/pixel resolution. For road upgrade identification, the algorithm solely relies on the Red band. For every band of a given composite, each pixel has a band value, i.e. an associated reflected-light intensity value. The unit and range of band values differ between satellites and sensors. Sentinel-2 band values vary in the range of 0 to 4,095 light intensity values.

When visualizing multi-band composites, matrices of spectral reflectance can be converted to computer graphics by mapping band values from three bands to a computer's RGB color model. If the Blue, Green, and Red bands are mapped to their respective colors in the RGB model, a true color composite is produced, where the rendered visualization will resemble that of how it appears from space. Composites produced by alternative band-color mappings are called false color composites. Two true color composites of a sample road, from different time periods (before and after road upgrade) are shown below:

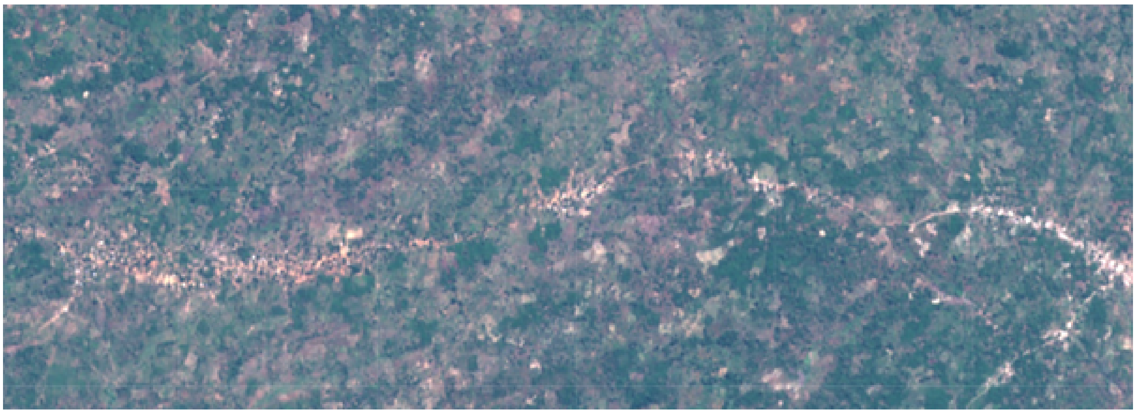

Figure 2: True color composite of a sample road
(2020-05-01 to 2020-07-31)

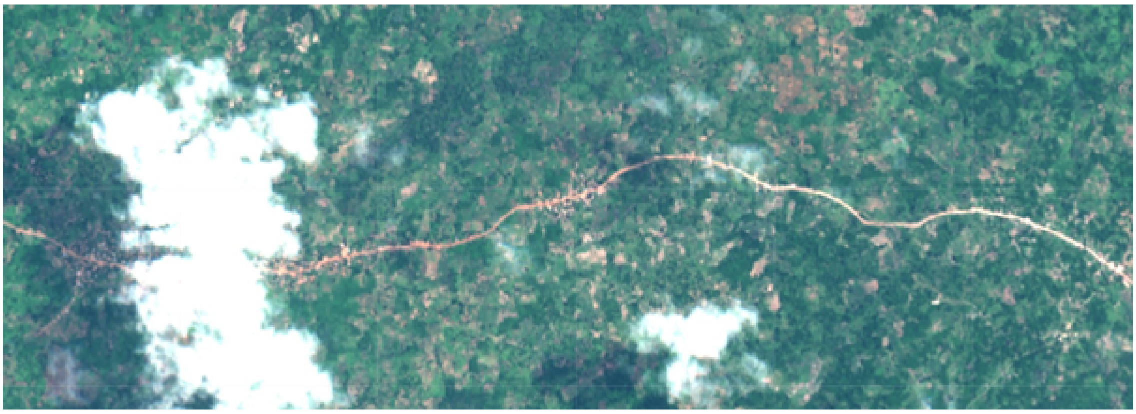

Figure 3: True color composite of a sample road
(2021-01-01 to 2021-03-31)

When visualizing a single band composite, band values from a single band are converted to a panchromatic (grayscale) representation of the image, where bright pixels indicate high spectral reflectance, and dark pixels indicate low reflectance. The same sample road is shown in its single-band visualizations below:

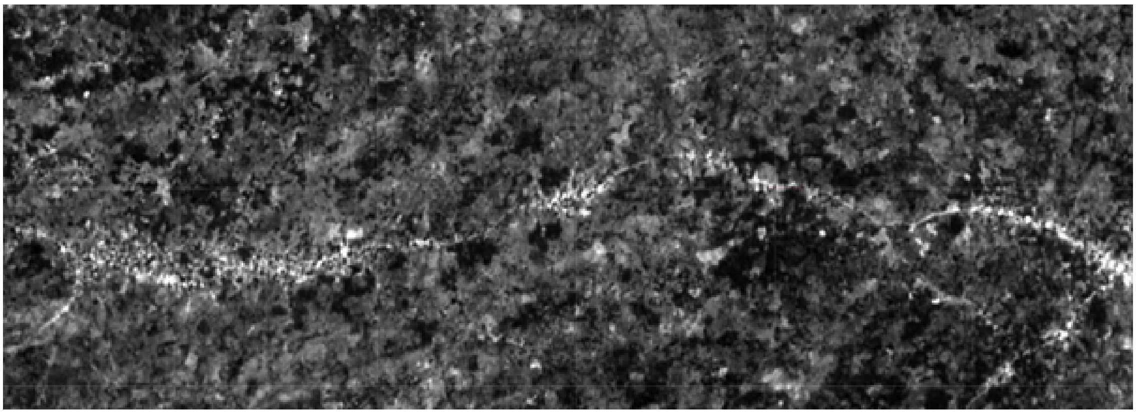

Figure 4: Single-band (Red) composite of a sample road
(2020-05-01 to 2020-07-31)

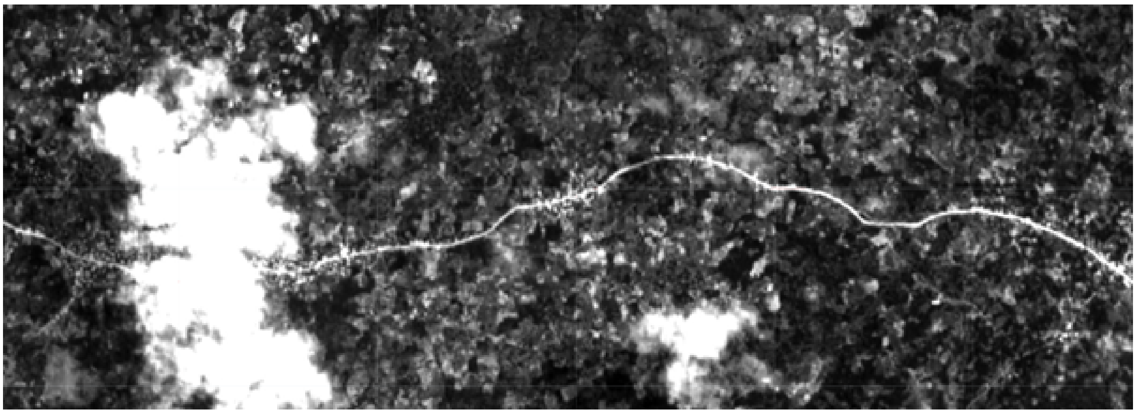

Figure 5: Single-band (Red) composite of a sample road
(2021-01-01 to 2021-03-31)

As in the images above, a road upgrade can be detected visually; the road appears much brighter in the latter image than in the former. With composite images being matrices of spectral-reflectance intensities, a difference of images can be computed by subtracting the band-values of respective pixels in the former image from the latter. Computing the differenced image from the two single-band composites above will itself yield a new single-band composite, shown below:

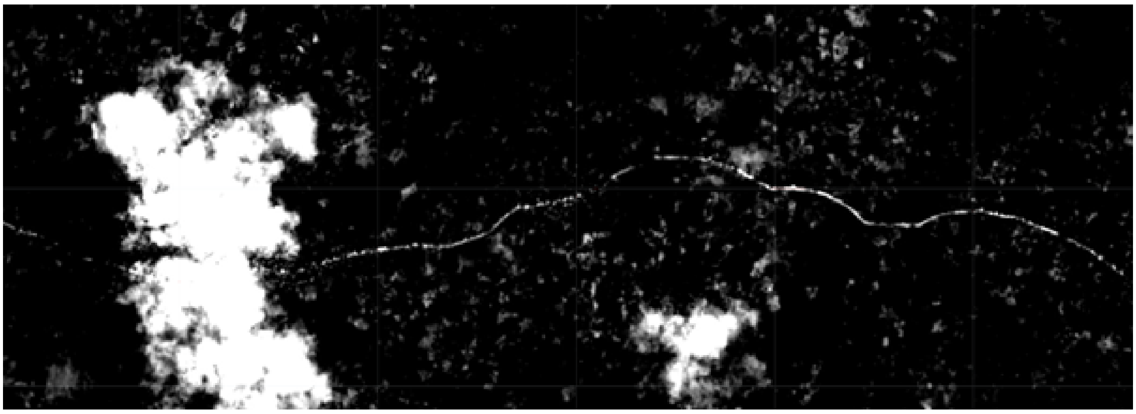

Figure 6: Single-band (Red) differenced image of a sample road
(2021-01-01 to 2021-03-31 vs 2020-05-01 to 2020-07-31)

The visualization of the differenced composite can be interpreted as bright pixels indicating brighter features that are present in the latter composite, but not in the former. The brightness of the pixel represents the magnitude of the difference, i.e. white pixels signify a more prominent change than gray pixels. Pixels with no change, or change in the other direction (strong reflectance in former composite, but weak in latter) appear as black pixels. The selected sample composites are particularly illustrative in showing this, as the clouds from the latter composite are clearly shown in the differenced composite. Similarly, the clear outline of the road indicates that the road is significantly brighter in the latter image, suggesting that the road has been upgraded by leveling. In the detection of road upgrades, we employ this described method of image differencing to detect roads that have been upgraded by leveling. By inverting the band values of the differenced composite, one can produce a composite highlighting features that are brighter in the former composite, than in the latter. As such, the method can be used to detect roads that have been upgraded by paving.

The road upgrade detection algorithm employs the aforementioned leveling and paving detection methods by differencing single-band composites on the Red (B4) band. In order to distinguish between upgraded and non-upgraded roads, the algorithm places points on the road, every 20 meters. For each of these points, control points are placed 25, 50, and 75 meters away, on both sides orthogonal to the road direction. In the illustration below, the points on the road are shown in green, and the 75-meter control points are blue and red. The 25 and 50-meter control points are excluded for clarity:

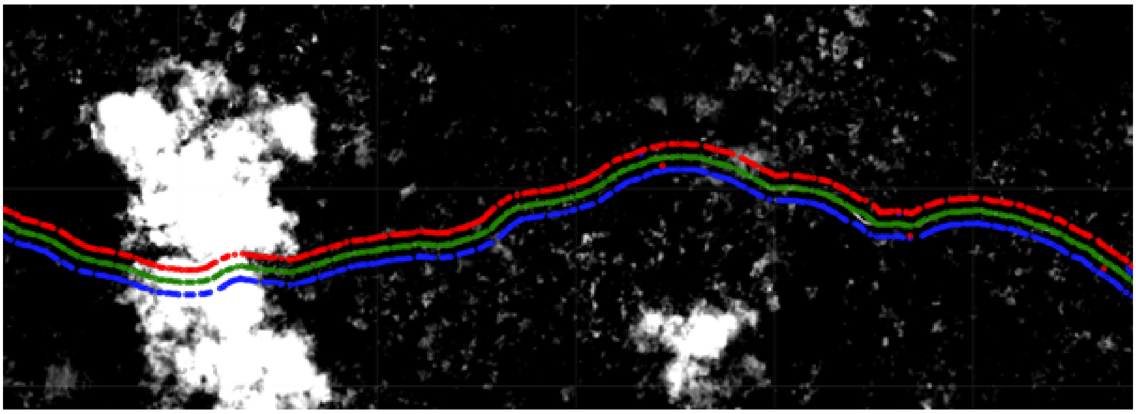

Figure 7: Placing of points on, and parallel to, road in differenced image of sample road

At each point, and accompanying control points, we measure the change in band values. The classification of a road as upgraded will depend on the relationship between the change in band values in points on the road, compared to the control points. The classification employs machine learning methods and is described in section 5.2.

When gathering data on the change in band values on the road and among control points, there are two main challenges: (1) the provided coordinates do not necessarily lie exactly on the road, and (2) the placement of a road may differ between composites. The former challenge is a result of data-entry errors, while the latter, commonly known as a correspondence problem, arises as a result of satellite images being taken from different angles, due to satellite orbits. In a process known as image rectification, satellite images are transformed to fit on a standardized coordinate system. This is done by mapping corresponding features to their standard coordinates. In remote regions, where it is harder to identify unique features, correspondence problems are more likely to arise. While the magnitude of these errors is at most a couple of meters, measurements taken on single pixels will be inaccurate.

We address both of the aforementioned challenges by buffering the points and control points by 10 meters. This means that instead of recording the change in band values at a single point, GEE aggregates the change in a surrounding circle with a 10-meter radius. We download the minimum, maximum, and mean change in band value. Any mentions hereafter of band value changes at a point or control point, will refer to the aggregated changes in the 10-meter buffer. The mitigation of errors arising from the first challenge, by employing this method, is illustrated below:

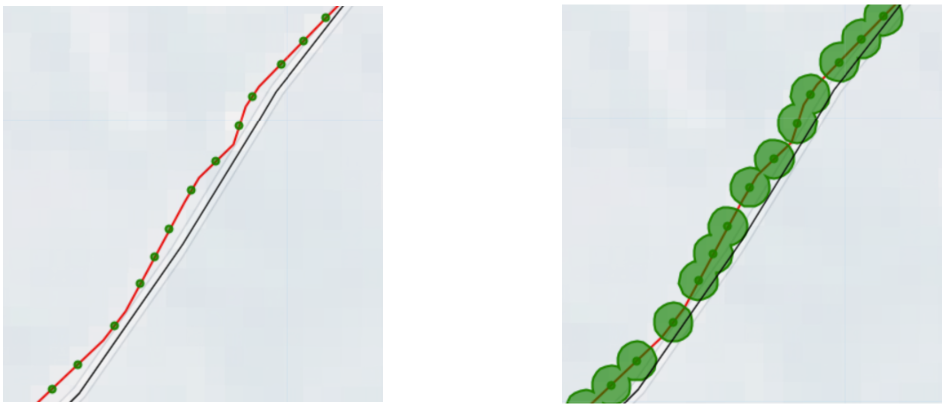

(a) provided coordinates not on road

(b) buffered points overlap road

Figure 8: point buffers correct for data-entry errors

In panel (a) of figure 8 we see that the road formed by coordinates from the provided data, shown in red, does not lie on top of the true location of the road, shown in black. Changes in band values on the green points, placed on the red line, will not represent changes in brightness of the actual road. By buffering the green points, as done in panel b, we see that the formed circles overlap the true (black) road. By recording the maximum change in band value in each circle, we will capture the change in brightness on the true road, if it indeed is upgraded.

In addressing the second challenge we modify the computations done in the buffered points. By following the image differencing method as described so far, a road that is not upgraded, but shifted as a result of imperfect rectification, will appear as an upgraded road in the differenced image i.e we will see a clear white line. An example is shown below:

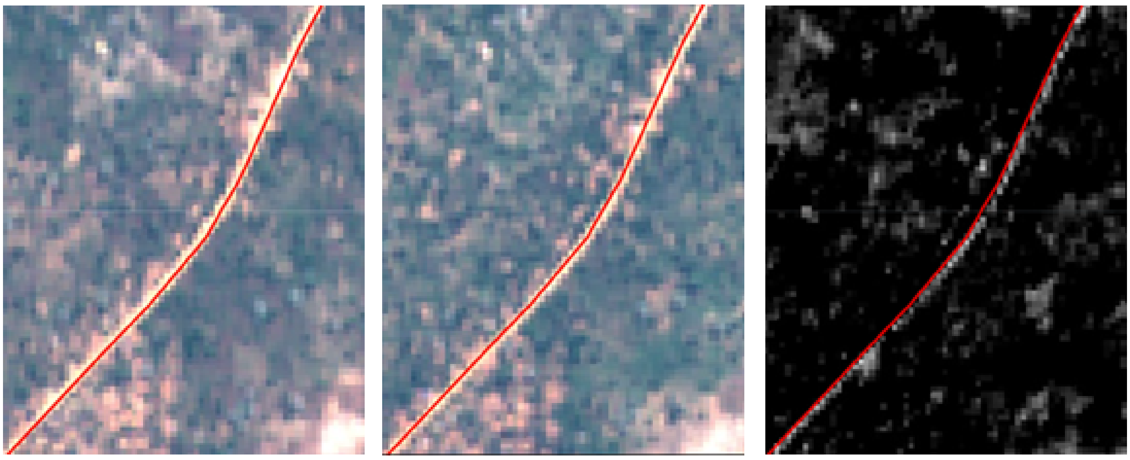

(a) earlier composite (b) later composite (c) differenced composite

Figure 9: shifts between composites result in false positives

In the panels of figure 9, we see a sample road where the road shifts between the earlier (a) and later (b) composites. With the red line staying in the same position throughout all images, we see that in the earlier composite a larger share of the road lies to the left of the line, while in the later the larger share lies to the right. In the differenced image (c), a gray line appears to the right of the red line, representing the positive difference between the bright road in the later composite, compared to the trees in the earlier. With buffering around points placed on the red line, as described earlier, this road would be classified as a false positive upgrade. To control for these false positives, we redesign the computations done in the buffered regions. Rather than measuring the minimum, maximum, and average band values at each point in the differenced image, we compute the minimum, maximum and average band values of each composite and take their difference. By doing this, we allow for small shifts in the position of a road while maintaining the ability to discriminate between upgraded and non-upgraded roads; a brighter road will be detected by a change in the maximum band value, while a wider road will be detected by a change in the average band value.

Finally, the overall brightness of composite images can differ. Given significant differences in brightness between composites, an upgraded road might not be correctly classified; while the road is brighter in absolute terms, the change relative to that of the control points (which also become brighter) will be diminished. To control for this, we normalize the measured differences in band values at each point, by adjusting for the difference in median band value of the two composites.

Employing the data-collection method described in this section, we download data for 146,587 road points and 879,552 control points.

These points cover all 219 roads in our dataset, in 455 unique segments, which we call subroads. These subroads arise as a result of the road coordinates from the World Bank being split into multiple segments. With a significant number of segments being noise, we only kept segments longer than 500 meters.

For each point, we download the change in minimum, maximum, and average band values within the point's buffered region, between the years 2016/2017, 2017/2018, 2018/2019, and 2020/2021. In each of these year intervals, we measure the change in band value in 5 overlapping composite periods, i.e. we generate composites from images taken in the three-month periods: April-June, May-July, June-August, July-September, and August-October. For each period, we compute the change in band value between the years. This amounts to 20 comparisons for each road in the four-year period. This overlapping, or "rolling", composite method has the effect of a smoothing kernel, which allows us to reduce image noise while maintaining the ability to discriminate between time periods to identify when an upgrade took place. We compare each three-month period with itself across years to avoid noise from seasonal differences. Furthermore, we do not compare composites between November and March as they feature many clouds due to the rainy season.

5.2 Machine learning classification

While the image-differencing method described in the previous section can roughly identify road upgrades, machine learning methods can be leveraged to improve the accuracy and speed of more detailed road upgrade detection. The motivations for employing machine learning methods are strong: (1) Visually identification of upgrades among the 455 subroads in the 20 comparison time periods, a total of 3,100 comparisons, is not feasible. (2) A well-trained machine learning model may detect upgrades missed in the visual identification process. (3) we intend to develop a tool that can be useful beyond this study. By aggregating the road and control point data to the road level, we generate a dataset suitable for road-upgrade classification, with which we train a random forest model.

For the training and assessment of the model, we visually identified a subset of road upgrades between 2020 and 2021, specifically looking at composites from the months of May-July. We selected these composites as they featured a high number of upgrades identified roughly in the road upgrade detection process. Among the 455 subroads in this time period, we visually identified 156 upgraded subroads. This visual identification was conducted by looking at each subroad in GEE. Images of the assessed subroads along with their visual classification, can be viewed in the digital Appendix (link found in Appendix B). All identified upgrades were leveling upgrades, as such we have no training data for paving upgrades. Thus, within the scope of this paper we do not detect upgrades by paving. However, our published code can be modified to do so. When describing the classification model in this section, the term road and subroad will be used interchangeably.

In our random forest model, we classify a road as either upgraded or not upgraded based on a number of variables, or features. The features *road_max*, *left_max_avg*, and *right_max_avg* capture the mean of the maximum change in Red (B4) band values at each of the road and control points on a road. For the latter two features, the average is computed of the means at the 25-, 50-, and 75-meter control points. The feature *road_mean* captures the mean of the mean change in band value at the road-points. While *road_max* will capture a road becoming brighter at any point in the 10-meter buffer around the point, *road_mean* will capture upgraded roads that are not necessarily becoming brighter, but significantly wider.

To further discriminate between upgraded and non-upgraded roads, we aggregate data of "upgraded points" for each road. We define an upgraded point as a road point that exhibits a specific band value change relative to its neighboring control points, setting itself apart from non-upgraded points. Whether or not a point is considered upgraded therefore depends on a threshold of relative band value change. We refer to this threshold as the *relative threshold* and define it as the maximum percentage of band value change experienced on the road point allowed to be experienced on the control points, on average, for a road point to be considered upgraded. This means that at a relative threshold of 50%, road points will be considered upgraded if their corresponding control points exhibit band value change of a magnitude half or less of that experienced on the road point, i.e. the road point is experiencing twice as much band value change. At a relative threshold of 100% all road points where the corresponding control points do not experience more band value change, on average, than that of the road point, are considered upgraded. At a relative threshold of 0% no points are considered upgraded and the feature is just noise. There is no theoretical motivation for how much brighter the change of a road point should be compared to the control point averages (other than that it should not exceed 100%) for it to be considered upgraded. Therefore we test the entire range of possible thresholds between 0% and 100%. The best threshold will be the one that defines upgraded and non-upgraded points, such that when aggregated they contribute the most to distinguishing between upgraded and non-upgraded roads.

In our random forest model, the feature *share* captures the share of points on the road that are upgraded. The feature *dist_mean* represents the mean distance between upgraded points on a road. To compute the distance between upgraded points, we sort all points in their order along the road. For each upgraded point, we compute the number of points between itself and the previous upgraded point. Finally, the feature *diff_factor* stores the band value change applied in the composite normalization, described at the end of the previous section.

With these features, we train the random forest classifier with 5-fold cross validation to reduce overfitting. To avoid data leakage, we train the model on a 90% training set and evaluate the model on the remaining 10% of the data. We tune the model's hyperparameters by iterating through all combinations in a finite range of values and selecting the best combination. We repeat these steps, training the model in this manner, for every 5 percentage point increment in relative threshold. Out of these 21 models, we select the model with the highest accuracy.

To evaluate our model, we compare it to a rule-based heuristic model. We base the heuristic solely on the feature *share* and classify roads with a share of upgraded points above a given level as upgraded. For every iteration across the range of relative thresholds, we compute the heuristic that best discriminates between upgraded and non-upgraded roads. It is this heuristic with the highest accuracy that we compare with our random forest model.

The variation in model and heuristic accuracy across the range of relative thresholds is shown in the figure below:

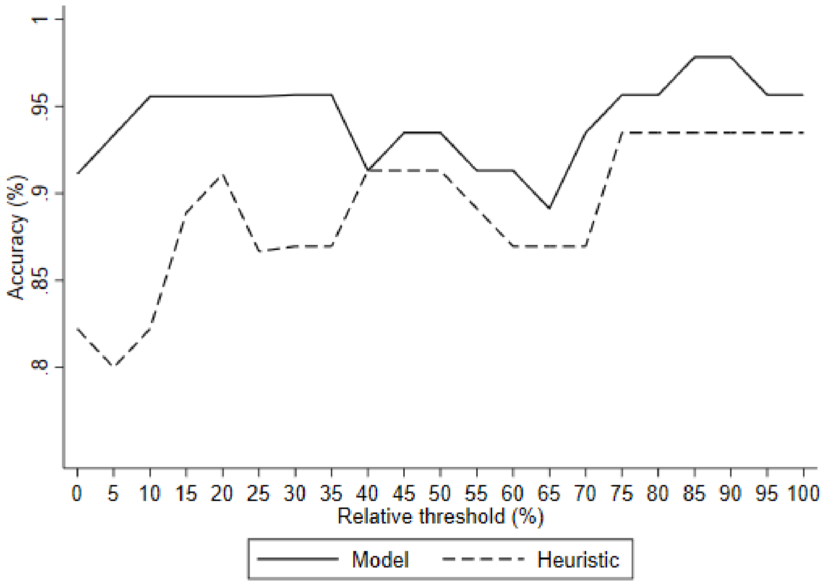

Figure 10: Variation in accuracy across relative thresholds

By varying the relative threshold of average band value change on control points with respect the road points, the accuracy of the model and heuristic differs. This means that both our model and heuristic are sensitive to changes in the relative threshold. Without the features *share* and *dist_mean*, our model, independent of point-level discrimination, performs with an accuracy of 93.48%. From the graph we see that when including the features *share* and *dist_mean*, the model can perform better or worse depending on the relative threshold.

When interpreting the graph, we recall that the relative threshold determines how many points are classified as upgraded. With a higher threshold, more points are considered upgraded. With more upgraded points the share, captured in the feature *share*, increases while the mean distance between upgraded points, *dist_mean*, decreases. With a low threshold, there are fewer upgraded points on each subroad. We see in the graph that the random forest model performs best in the relative threshold intervals of 10-35% and 75-100%. In these intervals the model classifies subroads with an accuracy of 95.56%, with the exception of thresholds 85 and 90%, where the accuracy is 97.83%. The model performs worst, with an accuracy of 89.13%, at the 65% relative threshold. This accuracy is lower than that of the model without the features *share* and *dist_mean*, which means that at such a relative threshold the upgraded point aggregations do not help discriminate between upgraded and non-upgraded roads, but rather contribute with noise.

The accuracy of the heuristic is similarly sensitive to changes in relative threshold. At low relative thresholds, where the share of upgraded points is low, there is less variation in the share of upgraded points between roads. In such cases, it is difficult to optimize a level of *share* for the heuristic that accurately discriminates between upgraded and non-upgraded roads. At higher relative thresholds, the rule-based heuristic discriminates between upgraded and non-upgraded roads with an accuracy of 93.48%. This accuracy is the same of that as the random forest model without the features *share* and *dist_mean*. A high-performing heuristic model suggests that the relative threshold is positively contributing to the discrimination between upgraded and non-upgraded roads.

Following this interpretation of figure 4, we select the relative threshold of 90% when defining upgraded points. Our subsequent random forest model performs with an accuracy of 97.83% compared to a heuristic of 93.48%. This translates to a 67% reduction in error rate.

The strength of a random forest model lies in its ability to identify complex relationships between the provided features. We can get an idea of these relationships by looking at the importance of each feature determined by the model. For each feature in the model the following shares of importance are reported:

road_max: 11.19%, *left_max_avg*: 3.34%, *right_max_avg*: 3.50%, *road_mean*: 56.89%, *dist_mean*: 2.41%, *share*: 20.33%, *diff_factor*: 2.34%,

From the feature importance report we see that the features *road_mean*, *share* and *road_max* have the highest importance. This is reasonable as the features capture the changes in width, balance, and brightness of road upgrades respectively. Nonetheless, the other features still play an important role in the capturing of complex relationships between the features which we would miss without machine learning.

When adapting this model to new contexts it will be important to train the model with new training data. While our trained model will be accessible in our publicly shared code repository, upgrades in other contexts may differ in appearance, and in extension differ in the relationships between features.

With our finally trained model, we predict upgrades for all subroads in all periods. While the initial visual classification of the subroad from one period took four hours to complete, the machine learning classification for subroads in all remaining 19 periods took less than a minute. Aside from the significant boost in accuracy of our random forest model, the time saved compared to manual classification is astounding.

With the rolling composite method described in the previous section, we attempt to control for any false positives predicted by the model by dropping upgrades that only appear in one of the five overlapping composite periods for a given year. Furthermore, we compute what share of an entire road has been upgraded by summing up the lengths of all upgraded subroads belonging to a road. We only consider roads as upgraded if at least 90% of the road length has been upgraded. With the classification of upgraded and non-upgraded roads now complete, we proceed to create grid cells.

5.3 Grid cell creation

The unit of observation in our regressions is a grid cell. Grid cells are generated by splitting a given subroad into strips of 1 kilometer length. From each strip, as well as the remainder following the split, a quadrangle is formed by placing points 5 kilometers to each side of every road strip node, orthogonal to the average direction of the subroad. With this strip height of 10 kilometers, a grid cell can have a maximum area of 10 square kilometers, when the strip is a 1 km straight line. With increasing curvature of the strip, the area of the corresponding grid cell decreases.

The area of the grid cell is important as it is the basis for our aggregations. If the area of the grid cell is too small and smallholder farmers who experience changes their agricultural productivity reside beyond the grid cell, then the effect that we are trying to estimate will not be captured. Similarly, if the area of the grid cell is too large, there will likely be noise for unaffected areas distorting our coefficients. Following the anecdotal evidence in section 3 with regards to smallholder farmers' access to markets and size of agricultural holdings, we believe a 10 kilometer grid cell height is appropriate. Controlling for the possibility that it should it not be the case, we collect data from a range grid cells, varying in grid cell heights, from 1 kilometer to 15 kilometers, at 1 kilometer increments. Our findings from this heterogeneity analysis are presented in section 8.4.1.

The placement of grid cells on a sample road is illustrated below. The orange grid cells have a grid cell height of 10 kilometers, while the blue grid cells have a height of 1 kilometer. The road is shown in red.

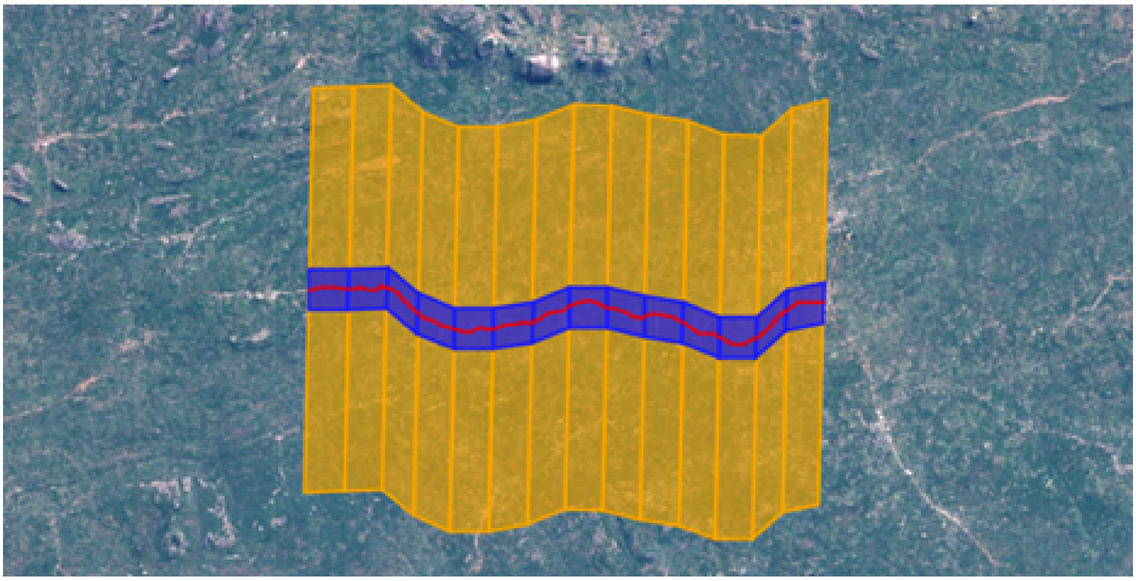

Figure 11: Grid cells placed along road

For each grid cell, several measurements are made. Agricultural output is estimated by measures of Normalized Difference Vegetation Index (NDVI). The index is calculated accordingly:

$$NDVI = \frac{(R_{NIR} - R_{red})}{(R_{NIR} + R_{red})} \tag{5}$$

where R_{NIR} and R_{red} are the reflectances of near-infrared and red light respectively. Vegetation exhibits high reflectance of near-infrared light and high absorption of red light (absorbtion of red light is what makes plants appear green to the human eye). At scale, increased vegetation should entail increased NDVI.

We collect NDVI averages computed and provided by NASA from their MODIS satellites. To isolate changes in vegetation as a result of agriculture, we calculate the change in NDVI value by comparing the beginning of the growing season with the planting season (see section 6.4). We call this agricultural NDVI. For each grid cell we also aggregate averaged monthly rainfall from the Copernicus Climate Data Store.

Furthermore, for each grid cell, we compute a measure of population density, provided by the Facebook Data For Good High Resolution Population Density Maps. The measure is computed for 2019. Finally we also measure the distance to the nearest city, from data downloaded from OpenStreetMap.

6 Method

In this section, we detail our study design and empirical strategy, as formulated econometrically. We also identify threats to the validity of our identifying assumptions, and potential remedies to issues we may come across.

We employ a differences-in-differences identification strategy to estimate the average treatment effect of receiving a road upgrade program on agricultural output on areas of land (grid cells) surrounding a road.

6.1 Econometric specification

The differences-in-differences specification is as follows:

$$\begin{aligned} NVDI_{it} = \beta_0 + \beta_1 \times upgrade_i + \beta_2 \times post_t + \lambda(upgrade \times post)_{it} \\ + X_{it} + \delta_i + \omega_t + \theta_i + \epsilon_{it} \end{aligned} \tag{6}$$

Where $NVDI_{it}$ is agricultural NDVI for grid cell i at time t, β_0 is the intercept, $upgrade_i$ is a dummy that takes in the value 1 if the grid cell is in the treatment group, $post_t$ is a dummy variable that takes on value 1 if the observation is is in the post-treatment time period. The difference-in-differences estimator is λ, X_{it} is a vector of control variables, δ_i are grid cell fixed effects, ω_t are year fixed effects, θ_i are province fixed effects, and ϵ_{it} is the error term.

Note that most of our controls are not used simultaneously as our grid cell fixed effects due to the fact that most controls are time-invariant. We also cannot use grid cell fixed effects at the same time as region fixed effects due to collinearity issues.

6.2 Panel data

We collect a panel of bimonthly NDVI data, and year-end as well as start of the year rainfall data covering the years 2010-2022. From 2016 and onwards, Sentinel data is available. We have identified road upgrades from 2018 and onwards, as the launch of the rural road upgrade program in 2018 should mean that roads are generally upgraded after this date. For our difference-in-differences estimations, we therefore use data from the year 2018 and onwards.

NDVI and rainfall data change for each year of our panel, while data on proximity to the closest city and road length remain constant throughout. The data on grid cell population represents only the year 2019.

Our panel is completely balanced (it contains all of our observations in every year), but it does contain missing values for some of the control variables.

6.3 Unit of observation

The unit of observation is spatial. We primarily use grid cells of approximately 1 km x 10 km in size (sizes differ due to an individual road's curvature) as our unit of observation. This measure is constructed at the road segment level, which means that all units of observation are associated with a unique road segment, that is in turn associated with a road that can be upgraded or not upgraded during the time of study.

The difference between *observations* and *road segments*, is that the number of observations comprises all years of study while a road segment is one unit of observation in a single year.

6.4 Outcome variable

The outcome variable is a measure of yearly NDVI difference, constructed by subtracting the NDVI composite of the previous fall/winter's NDVI from the NDVI of the early spring. The measure is constructed to control for vegetation that is non-agricultural, i.e. persists throughout the year. As such, we want to take the difference in vegetation index between the growing season (where NDVI should peak) and the planting season, where NDVI should be low. Our approach mimics Asher and Novosad (2020, p. 803). We use NDVI data for November-December for the fall/winter NDVI and for NDVI data for January-February for the early spring NDVI.

According to the FAO, Mozambique generally has a planting season for maize in between October and December, and a growing season during January to March, with harvest starting in the middle of March (Food and Agriculture Organization of the United Nations 2017). According to Agrilinks, central Mozambique has a planting season for maize in between November and December, with a vegetative season for maize in between January and March and harvest starting in the middle of March (Agrilinks 2020). The region of Zambézia would qualify as central Mozambique, whereas Nampula would qualify as central/northern Mozambique (manual comparison of maps provided by Agirlinks and Google Earth).

The Mozambique Food Security Update, dated 2006, says that planting in the northern regions of the country takes place in between November and January, with harvest in between April and July, meaning that growing can be inferred to take place between February and April (ReliefWeb 2006). Central Mozambique, on the other hand is described to plant in between October and January with immediate subsequent harvesting in between February and May (ibid.). The Food Security Update speaks generally of crops, among them maize, and specifically about the planting season of years 2006-2007 (ibid.). It does also refer to Zambézia as belonging to the central part of the country, and Nampula as belonging to the northern part, which is in line with what we concluded based on the Agirlink map (ibid.).

6.5 Independent variables

We have the collected the following independent variables: distance to the nearest city (km), road length (km), population in 2019 (number of people), and average daily precipitation (mm) for months November-February. In addition, we have stored information on grid cell size (km^2).

6.6 Further independent variables

We have kindly been given information on the prioritization score of our observations. This means that we have obtained information on the score on prioritization for upgrade given to a certain road in an assessment by the Mozambique Road Authority. Before individual roads were assigned a score, districts had been chosen according to "economic potential, network centrality and vulnerability" (von Carnap-Bornheim, personal communication, April 22, 2022). The prioritization exercise then addressed the following points: costs and economic potential, including agricultural potential, fishing, and tourism (ibid). The prioritization scores were subsequently used to select certain roads for upgrade within the frames of the World Bank project, with compliance between prioritization score and selecting a road for upgrade being almost perfect (ibid).

The information on prioritization score has been given to us in a “jittered” fashion. This means that we do not have information on the road's geographical location once we merge to our dataset the information on prioritization score. Instead, we use a random road ID to still be able to understand which observations belong to the same road. This is important for an appropriate handling of standard errors (see section 6.9). The prioritization scores range between –35.2 to 27.3 in our cleaned data set.

In addition to information on prioritization score, we have been given information on our observations' status as "selected" or "not selected" for an upgrade within the frames of the World Bank project. This variable is a dummy that takes on value one for selected roads, and zero otherwise. The information is conveyed to us in the same manner as information on the prioritization score, in a jittered fashion such that we no longer have information on the individual road's coordinates.

There is not an explicit relationship between prioritization score and agricultural potential. If there were, we would have a selection problem where roads that systematically differed in agricultural potential were prioritized higher and more often chosen to receive upgrades. While assessing agricultural potential could be rather difficult, it would be a strong assumption on our part to suggest that rural infrastructure placement would not attempt to take agricultural potential into account, especially when rural roads are upgraded in part to lift smallholder farmers out of poverty by enabling them to commercialize their agricultural practices. Recalling the aim of the prioritization exercise, it is evident that agricultural potential was addressed even if it remains unclear to what extent and in what way. We might thus worry that roads (or communities around the roads) that were selected to receive an upgrade share characteristics that are unobservable to us but correlated with agricultural potential, and that these roads (communities) systematically differ in this respect to roads (communities) that were not selected to receive an upgrade. If this is the case, it could be that the variable *selected* (a dummy that equals one for roads that were selected to receive an upgrade within the frames of the World Bank project, and zero otherwise) or the variable *prioritization score* holds some information about these factors. In other words, we think that worries we have with respect to systematic agricultural potential differences across our treatment and control roads might be captured to some extent by the variable *selected* or the variable *prioritization score*. To remedy the potential selection bias, we utilize the information in these variables to construct two secondary sets of treatment and control groups, see section 6.7.3.

Finally, we were given information about the roads' belonging to the province Nampula and Zambézia, respectively, and to which district the road belonged.

6.7 Treatment and control groups

In this section, we discuss the selection of treatment and control groups.

6.7.1 Treatment group

Out of our 219 roads, we identified 89 individual roads as partly upgraded with leveling between 2018 and 2021. Out of these 89 roads, 58 were (partly) upgraded between the spring of 2020 and the fall of 2021. 1 road was identified as a duplicate and was removed from the dataset.

We classified all roads with more than 90% of the road length upgraded as fully upgraded. In order to obtain the number of fully upgraded roads, we calculated road percentage upgrade using information on individual subroads' upgrade status. In terms of full upgrades, 73 roads were fully upgraded by 2021 and 54 out of these roads were upgraded between 2020-2021.

As the majority of roads were upgraded during 2020-2021, we deleted all observations associated with roads that were previously upgraded, as they could not be used as control roads nor treatment roads in our set-up. This led to a deletion of 19 roads. Roads that had a percentage upgrade over 10 but below 90 were removed from the data set, as they did not qualify as neither control nor treatment roads. This led to a deletion of 10 roads. One implication is that roads where 10% of the road was identified as upgraded, will be considered not upgraded. This is symmetrical to our definition of an upgraded road.

A treatment dummy was introduced. We let it equal one for all observations in the cleaned dataset identified as "upgraded", and zero for all observations in the cleaned dataset identified as "not upgraded". Our primary treatment group consists of 54 roads.

For the regressions with our two secondary control groups, we restrict the sample with respect to the treatment group, meaning that we have two secondary treatment groups. They consist of roads classified as both "treated" and as "selected", and roads classified as "treated" with positive prioritization score (see section 6.7.3).

6.7.2 Primary control group

The first control group comprises all the remaining road segments in the cleaned dataset not classified as treated (treatment dummy equal to zero). As such, our primary control group consists of 135 individual roads.

6.7.3 Secondary control group

For our secondary control groups, we utilized information of roads' prioritization score and "selected for upgrade" status. Along with information on a road's prioritization score, we were given information on whether a road was selected for upgrade.

The prioritization exercise is carried out prior to roads being selected for upgrade; as such, it can be seen as a guide for the allocation of rural road upgrades. A positive prioritization score is associated with a higher probability of being selected for upgrade (data not shown).

We firstly restrict our primary sample such that it only contains roads that were selected for upgrade. We may then compare “selected for upgrade” roads that are identified as upgraded by us (treatment roads) to “selected for upgrade” roads that are not identified as upgraded by us (control roads). The latter group is referred to as the “not yet treated” group, and performing this type of analysis may be warranted when there is reason to believe that "not-yet-treated" observations behave differently than "never-treated" observations (Callaway and Sant’Anna 2021, p. 205). Furthermore, the usage of the not-yet-treated as a control group, is consistent with an approach used to account for the endogeneity in aid project placement, in research studying the impact of aid projects on institutional quality (Isaksson and Durevall 2021). Our rationale for restricting the sample is that the restriction should leave us with more comparable treatment and control groups, in terms of agricultural potential, at the cost of decreasing sample size.

Secondly, we restrict the primary sample to only compare roads that have a positive prioritization score. This is also done in an attempt to weed out roads (both treatment and control) that are perhaps not suitable for comparison due to differing agricultural potential. While comparing "treated" roads to "not-yet-treated" roads is consistent with approaches used by scholars, the restriction proposed with respect to prioritization score ignores whether a road is "never-treated" or "not-yet-treated". In fact, our primary control group also cuts across this dimension, as it consists of roads in both the "never-treated" and the "not-yet-treated" group. It is not evident whether the endogeneity issue is best addressed by restricting the sample to only include roads selected for upgrade, or to restrict it according to prioritization score.

6.8 Fixed effects

For some of our specifications, we use fixed effects (FE) to control for non-observable factors that do not change across time. We primarily run specifications with grid cell fixed effects and year fixed effects. Grid cell fixed effects control for grid cell specific characteristics that do not change across time, and year fixed effects control for changes that are common across observations in each year.

The conflict in the province of Cabo Delgado in the northern part of Mozambique has seen many displaced civilians, with tens of thousands coming to Nampula, according to the UN (United Nations 2022a). Because of the conflict, we also run regressions using province fixed effects, as there is a possibility that the precarious situation affects agricultural outcomes in the Nampula province differently than it affects agricultural outcomes in Zambézia. Due to collinearity issues, we were only able to implement province fixed effects when removing grid cell fixed effects.

6.9 Standard errors

The observations consist of grid cells that are placed along each road segment. Grid cells along the same road can thus not be treated as independent observations. The status of treatment is allocated and identified at the road level, and as such the treatment statuses of spatially adjacent observations correlate perfectly. The experimental design, where treatment is assigned on a different level (road level) than the level we estimate effects on (grid cell level) motivates the usage of clustered standard errors (Abadie et al. 2017). In our case, we cluster at the road level as treatment is assigned at the road level and we expect to see that residuals correlate for observations belonging to the same road (clustering in the assignment).

Bertrand et al. (2004) shows that conventional standard errors from difference-in-differences estimates leads to an over-rejection of the null hypothesis. One remedy suggested is the clustering of standard errors, which fits well with our approach and the structure of our data that contains many groups (roads) to provide a sufficient number of clusters (Bertrand et al. 2004). Our clustering of standard errors is thus also important as an attempt to mitigate the risk of over-rejection of the null hypothesis.

6.10 Causality

There would be concerns if control and treatment observations prior to the intervention of the road upgrade program differ in agricultural potential. The challenge is to account for this endogeneity problem using a differences-in-differences set-up. While agricultural potential can be difficult to assess, a proxy might be to look at NDVI in observations prior to the launch of the road upgrade program. This way of dealing with agricultural potential ties neatly into the parallel trends assessment. If we find indications of violations in parallel trends prior to treatment, this could be an indication of issues with differing agricultural potential. Parallel trends will be discussed further on in section 8.3.

Another way of handling differences in agricultural potential prior to treatment, is to look at the prioritization score assigned to roads by the Mozambique Road Authority, and their subsequent status as "selected" or "not selected" for upgrade. One worry is that roads with similar prioritization scores are to some degree similar in characteristics used to assess their agricultural potential. Using roads' prioritization score and/or "selected for treatment" status, we weed out control (and treatment) roads that may not be suitable for comparison. This is done in the construction of our secondary treatment and control groups, see section 6.7.3.

7 Results

This section presents the results of our main regressions, our sensitivity analyses as well as our heterogeneity analyses. It starts by introducing some descriptive statistics. When NDVI is used in tables and figures, it should be interpreted as "agricultural NDVI".

7.1 Descriptive statistics

This section presents descriptive statistics for the pre- and post-treatment time periods, for our primary treatment and control groups respectively.

Certainly, when testing the difference of multiple variable means – in this case, as many as eight variables at a time – we would expect some statistical significance to be displayed just by chance. For instance, one in ten differences should exhibit statistical significance at the 10% level.

We decided to include a measure of difference over the standard deviation within the control group, to facilitate interpretation of the size of differences.

7.1.1 Descriptive statistics for primary treatment and control group

The below table shows descriptive statistics for our primary treatment and control observations, respectively, in the pre-treatment years of 2018-2020.

Table 1: Descriptive statistics for 2018-2020

	Treatment		Control		Difference	
	Mean	SD	Mean	SD	Diff	Diff/SD$_{Control}$
NDVI	0.21	0.09	0.19	0.09	0.02***	0.20
Nearest city	11.02	7.37	13.64	9.63	-2.61***	-0.27
Road length in kilometers	19.33	11.34	20.63	15.53	-1.29***	-0.08
Grid cell size	7.96	2.37	7.71	2.34	0.25***	0.11
Population in 2019	4,711.42	6,843.31	4,605.28	6,544.72	106.15	0.02
Grid cell population density	662.48	1,353.51	660.42	1,277.72	2.06	0.00
Prioritization score	6.00	6.18	-4.50	9.68	10.50***	1.09
Precipitation in mm	5.86	1.13	5.72	1.09	0.15***	0.13
Observations	2,358		5,706		8,064	

* p<0.05, ** p<0.01, *** p<0.001

Descriptive statistics for the pre-treatment time period by primary treatment and control group.

As is shown in the above table, agricultural NDVI and precipitation is somewhat higher in the treatment group compared to the control group. On average, observations in the treatment group are also located closer to cities and have a higher prioritization score. Road length is similar between the groups, albeit on average slightly higher in the control group. Population is similar too, with the treatment group having somewhat higher population, although the difference is not statistically significant at any conventional level. It also appears that grid cells are somewhat bigger in the treatment group compared to the control group. This suggests that treatment roads display systematically less curvature than control roads. Standard deviations for nearest city, road length, grid cell population and prioritization score are quite high in both groups, suggesting there is a lot of variation in the data.

The below table presents descriptive statistics for our treatment and control observations in the post-treatment time period, 2022.

Table 2: Descriptive statistics for 2022

	Treatment		Control		Difference	
	Mean	SD	Mean	SD	Diff	Diff/SD$_{Control}$
NDVI	0.27	0.09	0.23	0.09	0.04***	0.43
Precipitation in mm	7.23	1.44	6.68	1.74	0.55***	0.32
Observations	786		1,902		2,688	

* p<0.05, ** p<0.01, *** p<0.001

Descriptive statistics for the post-treatment time period by primary treatment and control group.

Average agricultural NDVI and precipitation are, just as in the pre-treatment time period, higher in the treatment group compared to the control group.

7.1.2 Prioritization score and “selected for upgrade” status on road level

In this section, we illustrate the allocation of prioritization scores and "selected for upgrade" status among the primary treatment and control groups. We do this for the primary sample.

Firstly, we display the mean and standard deviation of prioritization score, positive prioritization score and status as selected for upgrade, for the treatment and control group respectively. The data is presented on the road level.

Table 3: Prioritization score and "selected" status, road level

	Treatment		Control	
	Mean	SD	Mean	SD
Prioritization score	6.77	7.90	-3.50	11.22
Positive prioritization score	0.89	0.32	0.34	0.48
Selected	0.94	0.23	0.31	0.46
Number of roads	54		135	

Summary statistics of prioritization score, positive prioritization score and "selected for upgrade" status for primary treatment and control roads, respectively.

The mean of prioritization score is approximately 6.77 for roads in the treatment group and -3.50 for roads in the control group. For the dummy variables *positive prioritization score* and *selected*, the mean is closer to one for roads in the treatment group whereas it is below 0.5 for roads in the control group.

Table 4: "Treatment" group and "selected" status (road level)

	Selected for upgrade		Not selected for upgrade	
	Treatment	Control	Control	Treatment
Number of roads	51	42	93	3

The figure summarizes the number of roads identified as upgraded by us (primary treatment roads), by "selected for upgrade" status. The second column make up the not-yet-treated roads.

Figure 4 unveils a key factor that contributes to raising the mean of the variables in the control group in figure 3, namely that just short of half the roads selected for upgrade are not yet identified as upgraded by us, meaning that they make up the not-yet-treated group of observations. A positive prioritization score is associated with a higher probability of being classified as selected for upgrade (data not shown).

7.1.3 Descriptive statistics for secondary treatment and control group

This section presents descriptive statistics for the secondary treatment and control groups.

The appendix (appendix A) shows descriptive statistics for observations belonging to selected roads only, divided by treatment and control groups, in the pre-treatment years of 2018-2020. There is a larger difference between the control group mean of agricultural NDVI and the treatment group mean of agricultural NDVI for this sub-sample, compared to the primary sample. Potentially, road upgrading was decided to commence for the roads that had the highest level of agricultural NDVI within the group of selected roads. Other than this, observations that are selected for upgrade appear to be more similar in terms of proximity to city and prioritization score, than the observations in the primary sample are. Treatment observations have somewhat higher precipitation and lower population than control observations do. Road length is similar, although treatment observations on average belong to a somewhat longer road compared to control observations.

In the appendix (appendix A), we also show descriptive statistics for observations with positive prioritization score only, divided by treatment and control groups, in the pre-treatment years of 2018-2020. Observations with positive prioritization score are more similar in terms of average agricultural NDVI than observations in the selected only category. Treatment observations, on average, experience more precipitation and are located closer to cities. They have a lower prioritization score than control roads, slightly lower population density, and are associated with on average shorter roads.

Furthermore, the appendix (appendix A) presents descriptive statistics for selected roads only, divided by treatment and control groups, in the post-treatment year 2022. The difference in agricultural NDVI means between treatment and control observations is large in the post treatment time period.

Finally, the appendix (appendix A) presents descriptive statistics for observations with positive prioritization score only, divided by treatment and control groups, in the post-treatment year 2022. Similarly to the selected roads only, there is a large difference in agricultural NDVI means between treatment and control observations in the post-treatment time period.

7.2 Pre-treatment trends in NDVI measures and rainfall

An important assumption in differences-in-differences estimations is the parallel trends assumption. In this section, we illustrate the pre-treatment trends for our primary and secondary sample.

7.2.1 Pre-treatment trends for primary sample

Below, we present pre-treatment trends in agricultural NDVI for the primary treatment and control groups. The data covers the years 2011-2020.

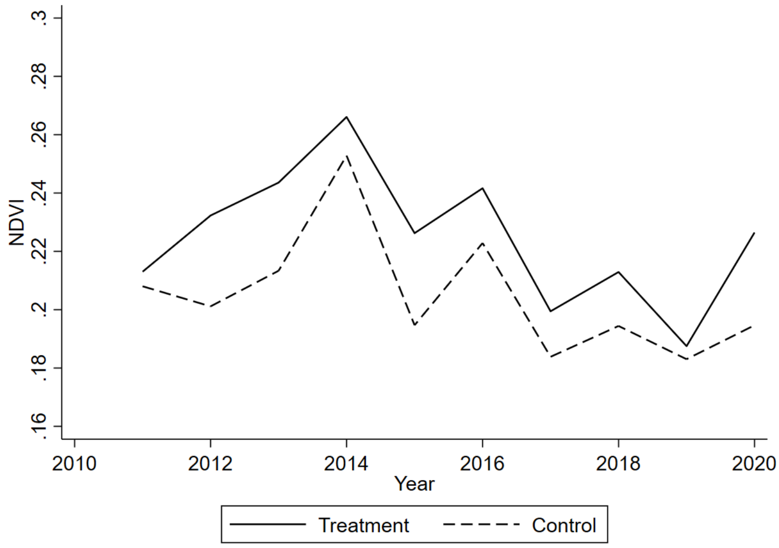

Means in agricultural NDVI over the years 2011-2020 for the primary treatment and control group.

Figure 12: Pre-treatment trends in agricultural NDVI

The trends in agricultural NDVI are fairly similar for the treatment and control group, from the year 2013 and onwards. The level of agricultural NDVI is bigger for the treatment group throughout the years 2011-2020, and appears to spike somewhat more abruptly than the control group between the years 2019-2020.

In order to eliminate some of the fluctuation in agricultural NDVI, we decided to show the pre-treatment trends with variation from rainfall removed; conditional pre-treatment trends. Since precipitation matters greatly for variation in NDVI, we ran a regression of agricultural NDVI on rainfall with standard errors clustered at the road level. The residuals were collected and plotted against the same x-axis. The rainfall-adjusted trends are represented graphically in the below figure.

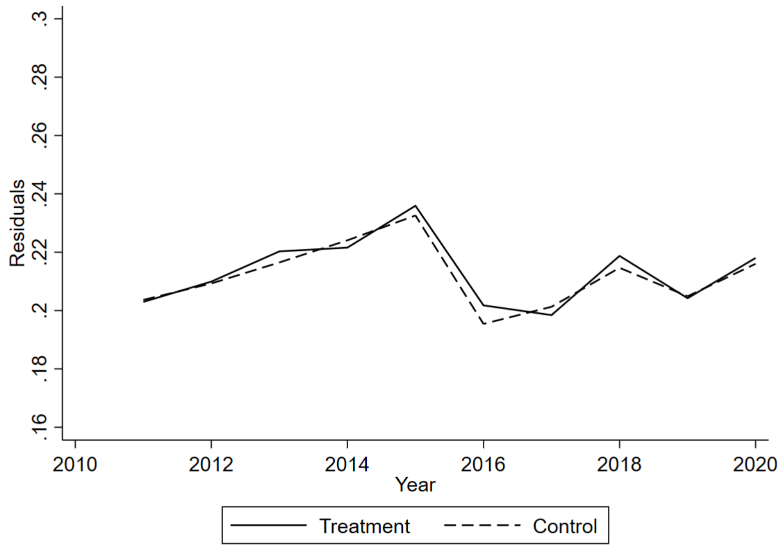

Means of residuals collected from regression of agricultural NDVI on rainfall for the primary treatment and control group.

Figure 13: Pre-treatment trends adjusted for precipitation

The variation in means of agricultural NDVI is reduced when rainfall is taken into account, and centered on approximately 0.19-0.24. The two lines follow each other closely, the treatment group in general having somewhat higher agricultural NDVI. Compared to figure 12, trends now look somewhat parallel also for the years 2011-2013, although there is a spike for the treatment group in year 2013 and a subsequent drop in year 2014. There is a somewhat lower drop in 2016 for the control group compared to the treatment group, and some subsequent oscillation in the two lines over the years 2017-2019.

7.2.2 Pre-treatment trends for secondary sample

As the pre-treatment trends adjusted for rainfall smooth the variation in agricultural NDVI but preserves rather parallel trends, we reproduce this figure for our secondary samples as well. As such, figure 14 shows pre-treatment precipitation adjusted trends for the secondary control and treatment groups.

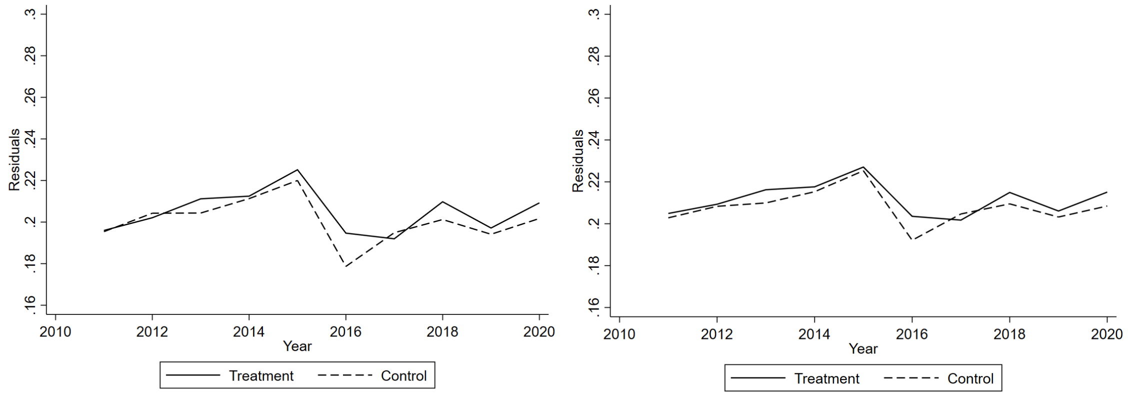

(a) Selected roads

(b) Roads with positive prioritization score

Figure 14: Secondary treatment and control group pre-treatment precipitation-adjusted trends

The trends are similar to the ones shown in figure 13. They appear somewhat more parallel in panel (b) compared to panel (a). They appear somewhat less parallel than the trends in figure 13.

7.3 Difference-in-differences estimation

7.3.1 Primary estimation (main results)

Our differences-in-differences estimates for the primary treatment and control groups are shown in the below table. Specification uses (1) grid cell fixed effects and year fixed effects. Regression (2) adds a covariate for rainfall to the grid cell and year fixed effects. Regression (3) is similar to regression (2), but instead of grid cell fixed effects, it uses province fixed effects. Regression (4) uses the a set of covariates (road length in km, distance in km to nearest city, population in 2019), year fixed effects and province fixed effects, but no grid cell fixed effects. We also ran regressions (without covariates) using pooled OLS and grid cell fixed effects only, with the results effectively identical to regression (1).

Table 5: Primary difference-in-differences

	(1) NDVI	(2) NDVI	(3) NDVI	(4) NDVI
DiD	0.0223*	0.0146	0.0178	0.0175
	(0.0123)	(0.0118)	(0.0119)	(0.0118)
Precipitation in mm		0.0187***	0.0107***	0.0114***
		(0.0030)	(0.0026)	(0.0026)
Grid cell FE	Yes	Yes	No	No
Year FE	Yes	Yes	Yes	Yes
Set of controls	No	No	No	Yes
Province fixed effects	No	No	Yes	Yes
Constant	0.1998***	0.0870***	0.1429***	0.1186***
	(0.0038)	(0.0197)	(0.0176)	(0.0193)
N	10,752	10,736	10,672	10,672

Standard errors in parentheses
* p<0.1, ** p<0.05, *** p<0.01

The coefficient of interest is the difference-in-differences estimator, DiD. In terms of direction of the estimate, DiD is consistently positive throughout the specifications, although it decreases in magnitude when rainfall is added to the regression. The relative size of standard errors increases as we add rainfall to the regression, meaning that the estimated coefficients on DiD no longer are significant at the 10% level.[5,6]

[5]As our sample size is quite substantial, a statistical significance at 10% should be interpreted with caution, when it comes to assessing the size of statistical significance.

[6]We also ran the regressions with district fixed effects, specifically regressions with year FE and district FE. When adding rainfall to the regression, the coefficient decreases in magnitude and relative size of standard errors increase. One motivation for including district fixed effects, is that we worry that our precipitation measure is too coarse to account for rainfall differences within districts. When district fixed effects are added, grid cell fixed effects are removed due to collinearity issues. The estimated coefficients and their standard errors do not change substantially, but rather display large similarities with the estimates shown above.

7.3.2 Secondary estimation (main results)

Below, we restrict the primary sample to (1) only include the observations that belong to a road selected for upgrade, and to (2) only include observations that have a prioritization score above zero. We proceed with our preferred specification from table 5, specification (2), that used grid cell fixed effects, year fixed effects and our precipitation variable, since precipitation appeared to matter for the difference-in-differences estimator.

Table 6: Secondary difference-in-differences

	(1) Selected NDVI	(2) Positive score NDVI
DiD	0.0348**	0.0386***
	(0.0162)	(0.0138)
Precipitation in mm	0.0178***	0.0185***
	(0.0047)	(0.0045)
Grid cell FE	Yes	Yes
Year FE	Yes	Yes
Constant	0.0859***	0.0946***
	(0.0309)	(0.0292)
N	5,100	5,128

Standard errors in parentheses
* p<0.1, ** p<0.05, *** p<0.01

The estimated coefficients of interest are positive, just as in table 5, but larger in magnitude. They are statistically significant at the 5% and 1% level, respectively. The estimate from regression (2) is higher than the estimate from regression (1), and its relative size of standard errors is smaller. We also note that the sample size has decreased substantially compared to what it was in table 5.

7.4 Sensitivity analysis

This section presents a sensitivity analysis, that is comprised by two regressions and is aimed at investigating anticipatory effects and performing a placebo test, respectively. Primary as well as secondary treatment and control groups are used.

Regression (1), (3) and (5) show the test of "no anticipation". Since the World Bank grant for the road upgrade project was approved in May, 2018, regression (1), (3) and (5) use year 2019 as the first treatment year, and introduces a dummy for year 2019 interacted with the treatment dummy – this is the coefficient named *Anticipation effect*. We perform our preferred specification for the years 2018-2022. Absent anticipatory effects, we should expect not to see any statistically nor economically significant coefficient for the interaction dummy. Regression (2), (4) and (6) perform a placebo test. We introduce a placebo treatment between year 2015 and 2016, and code observations in years 2013-2015 as pre-treatment and observations in years 2016-2018 as post-treatment. The placebo test should not yield economically nor statistically significant results, if it were to support the main regressions and our "true" treatment event.

Table 7: Sensitivity analysis

	(1)	(2)	(3)	(4)	(5)	(6)
	Primary sample		Selected roads		Positive prioritization score	
	NDVI	NDVI	NDVI	NDVI	NDVI	NDVI
DiD, anticipation test	0.0094		0.0215		0.0243	
	(0.0137)		(0.0178)		(0.0174)	
Anticipation effect	-0.0161		-0.0470***		-0.0280*	
	(0.0126)		(0.0155)		(0.0154)	
DiD, placebo test		-0.0084		-0.0240**		-0.0314***
		(0.0070)		(0.0100)		(0.0099)
Precipitation in mm	0.0193***	0.0128***	0.0186***	0.0125***	0.0213***	0.0130***
	(0.0028)	(0.0017)	(0.0043)	(0.0031)	(0.0040)	(0.0027)
Grid cell FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Constant	0.0836***	0.1435***	0.0810***	0.1302***	0.0781***	0.1341***
	(0.0182)	(0.0119)	(0.0294)	(0.0208)	(0.0268)	(0.0182)
N	13,420	16,104	6,375	7,650	6,410	7,692

Standard errors in parentheses

* p<0.1, ** p<0.05, *** p<0.01

For the primary sample, shown in column (1) and (2), the coefficient on *Anticipation effect* is between the difference-in-difference estimates from regressions (2) to (4) in table 5 in magnitude, although negative and not statistically different from zero at any conventional level. The coefficient of interest from the placebo test is comparably small in magnitude, negative and not statistically different from zero at any conventional level.

Regression (1) shows an estimate in the negative direction. If there were in fact anticipatory behavior, we would rather expect to see a positive coefficient here. Nevertheless, it cannot be rejected on any conventional level that the coefficient is equal to zero, which supports the notion of no anticipatory behavior. Similarly, it cannot be rejected on any conventional level in regression (2) that the estimate is in fact equal to zero. The coefficient itself is economically small and negative. If it were greater in size or if its relative standard errors were smaller, we would be worried that the parallel trends assumption was violated.

Moving on to the secondary treatment and control groups, the estimates change. The anticipation test in (3) uses only roads selected for upgrade. It shows a negative coefficient large in magnitude and statistically different from zero on the 1% level. The placebo test for this subsample, regression (4), gives a negative coefficient of sizeable magnitude, significant at the 5% level. For the sample where we only include roads with a positive prioritization score, the coefficient on *Anticipation effect* in regression (5) is of non-negligible size, negative and statistically significant at the 10% level. The placebo test, shown in regression (6), yields, again, a negative coefficient of non-negligible size, statistically significant at the 1% level. The results for our secondary samples indicate some issues with our identification strategy, and will be further elaborated in section 8.3.1.

7.5 Heterogeneity analyses

In this section, we present two heterogeneity analyses: varying grid cell size and decomposing the difference-in-differences estimator by distance to nearest city. Primary as well as secondary treatment and control groups are used.

7.5.1 Difference-in-differences estimation with varying grid cell size

In our first heterogeneity analysis, we change the size of our grid cells, letting the height of the grid cell vary between 1 and 15 km, in increments of 1 km. The width of grid cells are constantly preserved at (approximately) 1 km. We perform our preferred specification, regression (2) from table 5, with varying grid cell height and plot the coefficient of interest together with its associated 95% confidence interval, calculated by adding and subtracting $1.96 \times se$, standard error of the coefficient, from the estimates. We do this as we are concerned that effects on agricultural NDVI closer to the road could be attenuated by large grid cell sizes, if there are less or no agricultural effects further away from the road.

7.5.2 Primary sample

Firstly, we show the results of the grid cell size heterogeneity analysis for our primary sample.

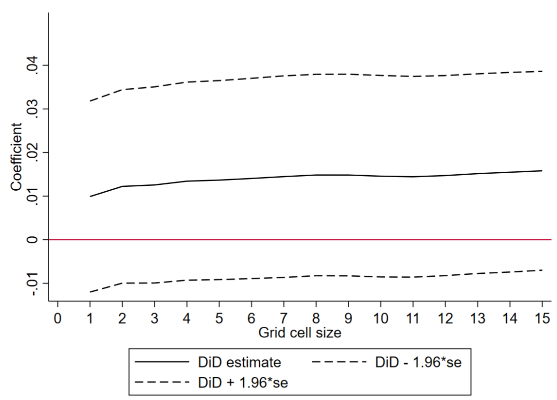

(a) Confidence interval at 95%

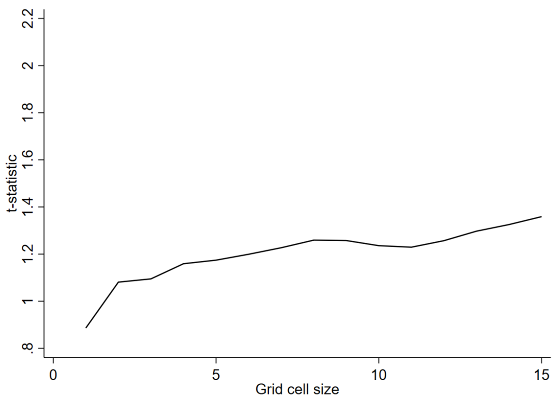

(b) t-statistic over grid cell sizes

Figure 15: DiD estimates for differing grid cell sizes

Our difference-in-differences estimates increase as grid cell size increases, as is seen in panel (a). While the slope is quite steep between the smallest grid cell sizes, it appears to flatten out at 2 by 1 km. For grid cell sizes larger than 2 by 1 km, the change in size of coefficient is modest. As grid cell size reaches 15 by 1 km, the lower bound of the 95% confidence interval is at the closest of an intersection with the vertical line placed at zero. The coefficient is not statistically different from zero on the 5% level for any grid cell size. Panel (b) shows how the t-statistic of the estimated coefficient changes with varying grid cell size.

The concern that effects in agricultural NDVI would be attenuated by our original grid cell size is not supported by figure 15.

7.5.3 Secondary sample

In the below figures, we perform the analyses described in section 7.5.1, this time using our secondary sample. We start by showing the results for selected roads and proceed by showing the results for roads with positive prioritization score. The estimated coffeficients are generally larger for our secondary sample, compared to our primary sample.

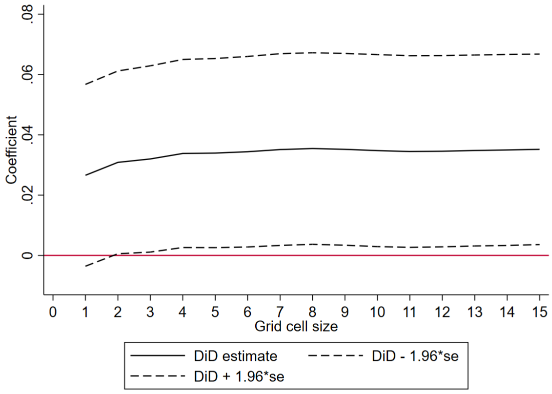

(a) Confidence interval at 95%

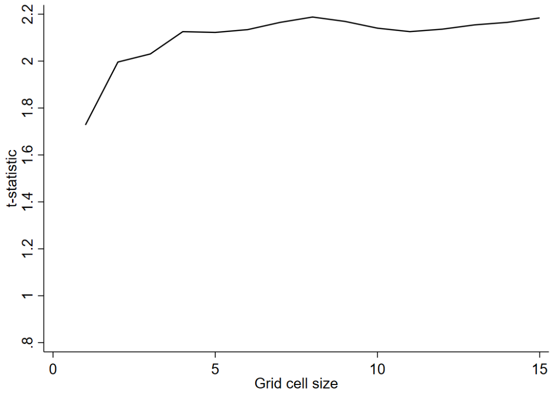

(b) t-statistic over grid cell sizes, selected roads only

Figure 16: DiD estimates for differing grid cell sizes, selected roads

Similarly to the primary sample, our estimated coefficients are most sensitive for grid cell sizes in the smallest range, although the slope in panel (a) is quite steep not only for heights bewteen 1 and 2 km, but between 1 and 4 km. The estimates are quite stable as grid cell size increases, and they are statistically different from zero at the 5% level for grid cell sizes above 2 by 1 km. The t-statistic increases steeply for small grid cell sizes but flattens out as grid cell height is approximately 5 km.

The concern that effects in agricultural NDVI would be attenuated by our standard 10 by 1 km grid cell size is not supported for selected roads only.

The below figure finally shows the results with varying grid cell size for roads with a positive prioritization score.

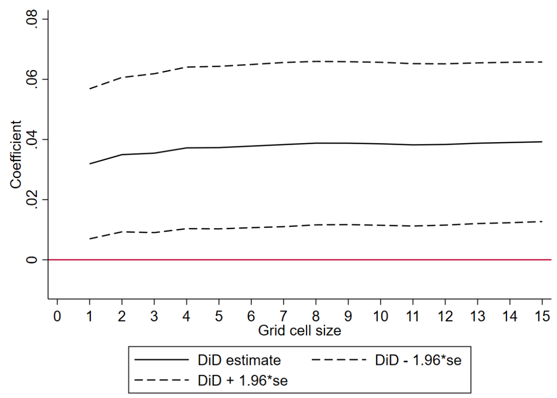

(a) Confidence interval at 95%

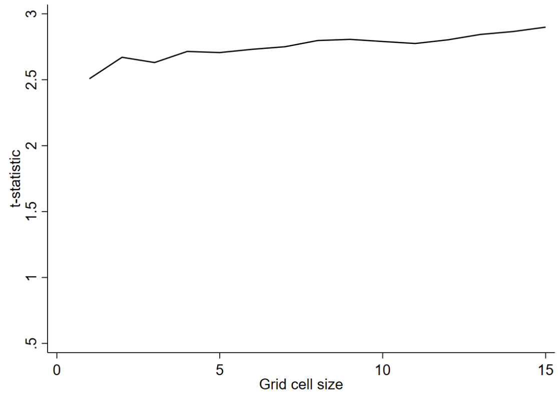

(b) t-statistic over grid cell sizes, positive prioritization score only

Figure 17: DiD estimates for differing grid cell sizes, positive prioritization score only

Figure 17 exhibits resemblance with both previous figures; the slope on coefficient is again quite steep for small grid cell heights, especially between 1 and 2 km as well as between 2 and 4 km. It appears to flatten out when grid cell height is approximately 4 km. What sets this figure apart from the previous, is that even the smallest grid cell sizes display estimates that are statistically different from zero. The t-statistic is generally upward sloping, although the slope is modest.

Again, we find no evidence suggesting that the concern of large grid cell sizes attenuating agricultural NDVI effects would be valid.

7.5.4 Difference-in-differences estimation by distance to nearest city

In our second heterogeneity analysis, we categorize our observations according to their distance to the nearest city into three groups: within 10 km from the city, between 10 and 20 km, and further away than 20 km from the city. We then proceed with our preferred specification, regression (2) from table 5, and interact the difference-in-differences variable with dummy variables for distance categories, preserving the closest distance as baseline. We do this to see whether observations at a particular distance appear to drive the results of our main regressions. As advanced in Fafchamps and Shilpi (2003), economic activity could very well be spatially distributed, such that farmers in our sample only within a certain distance from cities would be likely to commercializing agricultural practices following rural road upgrades.

Below, we present the difference-in-differences estimator, decomposed with respect to distance to the nearest city.

Table 8: Difference-in-differences by decomposition according to the nearest city

	(1) Primary sample NDVI	(2) Selected roads NDVI	(3) Positive score NDVI
DiD	0.0004	0.0198	0.0270*
	(0.0130)	(0.0173)	(0.0153)
DiD $\times$ Nearest city $\geq$ 10 km and $<$ 20 km	0.0271*	0.0281*	0.0210
	(0.0157)	(0.0162)	(0.0168)
DiD $\times$ Nearest city $>$ 20 km	0.0282	0.0293	0.0278
	(0.0293)	(0.0299)	(0.0314)
Precipitation in mm	0.0188***	0.0180***	0.0188***
	(0.0031)	(0.0049)	(0.0047)
Grid cell FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Constant	0.0863***	0.0846***	0.0929***
	(0.0201)	(0.0318)	(0.0303)
N	10,736	5,100	5,128

Standard errors in parentheses

* p<0.1, ** p<0.05, *** p<0.01

For the primary sample, the estimated coefficients are economically larger for distances further away from the city, although only the coefficient on the difference-in-differences estimator interacted with a distance dummy in the second category is statistically significant at a conventional level (10%), and given the sample size this level of significance is not strong. Furthermore, we are not able to reject the null hypothesis that the joint significance of the parameters of interest is equal to zero, as the p-value of this F-test is 0.2324.

When it comes to the secondary sample, there are both similarities and discrepancies. Starting with specification (2), that uses the selected roads only, the findings are similar to those of specification (1), with the exception that the coefficient on the difference-in-difference estimator without interactions is of much larger magnitude. The p-value of the F-test for joint significance of the three parameters of interest is 0.0671, which indicates that we may reject the null hypothesis of no joint significance at the 10% level. The results from this specification suggest that observations within 10 and 20 km may drive the results of our main regressions.

Specification (3), on the other hand, tells a different story. Now the coefficient is smallest when interacted with the second distance dummy. Only the difference-in-difference estimator without interaction terms is statistically significant at the 10%. The p-value of the F-test for joint significance of the three variables is 0.0374, indicating that we may reject the null hypothesis of no joint significance at the 5% level. The results from specification (3) suggest that observations closest to the city may drive the main results.

8 Discussion

8.1 Primary results

Our primary regressions of the main estimations show slightly varied results. The regression that do not control for rainfall, regression (1), show that post-treatment, grid cells associated with upgraded roads experience an increase in agricultural NDVI compared to non-upgraded roads and compared to pre-treatment time periods, statistically significant at the 10% level. These estimates are likely to be biased, as the coefficients of interest change in magnitude and statistical significance when rainfall is taken into account. Regressions (2) to (4) display somewhat smaller estimates than (1), and they are not statistically significant at a conventional level: we cannot reject that the estimates of interest in regressions (2) to (4) are equal to zero. Regression (2) uses grid cell fixed effects, year fixed effects and the rainfall control, and it is our preferred specification as it controls for grid cell characteristics that remain constant through time, as well as year effects, and since precipitation appeared to matter for the difference-in-differences estimator.

Even though we cannot reject that the coefficient of interest from our preferred specification is in fact equal to zero, we note that all our DiD estimates are consistently positive throughout the primary specifications. The estimated increase in agricultural NDVI is of order 1/10 of the average agricultural NDVI, as can be seen when comparing the size of the coefficients of interest in table 5 and reading off the y-axis in figures 12 and 13. For regression (2), the size of the coefficient on DiD is approximately similar in size to the coefficient on the precipitation variable (although smaller). Reflecting on the magnitudes of these coefficients, it means that the effect on agricultural NDVI from receiving a road upgrade, almost corresponds to an increase in daily average rainfall by 1 mm.

The directions of the coefficients of interest in table 5 are consistent with the theory presented in section 4.1, positive effects on agricultural output. When rainfall is added to the regression, the coefficients of interest decrease in size and the relative size of standard errors increase (see table 5), which indicates that there is some correlation between the difference-in-differences estimator and the rainfall variable. This is puzzling, as it indicates that precipitation would be correlated with the treatment itself. From the regression, it appears that when rainfall is left in the error term, estimates are biased, the direction of bias being positive. Precipitation is positively correlated with agricultural NDVI, and as bias is positive, it should be positively correlated with the treatment. The fact that the coefficient on rainfall displays strong statistical significance is not surprising, as rainfall affects agricultural vegetation. One explanation for the correlation between rainfall and the treatment, would be that it was decided to commence upgrade work on roads that received more rain.

8.2 Secondary results

The results from our secondary estimation of the main regression, presented in table 6, are consistent with the results from the primary estimations in terms of direction of coefficients, although the coefficients are larger in magnitude and the relative size of standard errors is smaller. Indeed, we reject the null hypothesis that the difference-in-differences estimates are equal to zero on the 5% and 1% level, respectively. Theoretically, there are reasons to believe that the secondary samples are more suitable to proceed with in the analyses, as the treatment and control groups should be more comparable. While it is an accepted approach to compare the "treated" to the "not-yet-treated" and in doing so exclude the "never-treated" from analysis, the exclusion of observations based on prioritization score is less conventional. While the group of "selected" observations and the observations with positive prioritization score have great overlap, the rationale for considering the latter subsample was the prioritization exercise carried out by the Mozambique Road Authority. Visually, the parallel trends assumption appear to be somewhat better justified for this subsample than for the selected roads, as is shown in figure 14, although trends look less parallel than the primary pre-treatment trends in figure 13 (primary sample). If we think that there might be some endogeneity concerns with respect to agricultural potential, they should be mitigated to some degree when the comparison considers roads that are selected for upgrade only, or potentially roads that scored positively on prioritization. Of course, roads with a prioritization scores just above zero might be more similar to roads with a prioritization score just below zero, than to roads with a very high prioritization score. This would be an argument against treating observations with positive score as more similar to each other in terms of agricultural potential. However, it is also not completely evident which roads were classified as "selected", as the compliance with positive prioritization score is not perfect. Looking at the subsample consisting of roads with positive score only, could be seen as a complementary analysis in this respect.

The descriptive statistics reveal some interesting facts. Firstly, there are quite meaningful differences in average agricultural NDVI pre-treatment for the group that comprises selected roads only. On the other hand, prioritization scores are much more similar for this subsample than between the treatment and control groups of the primary sample. Secondly, prioritization scores are actually higher in the control group when we consider roads with positive prioritization score only. Differences in agricultural NDVI pre-treatment are also larger for this subsample than for the primary sample. The primary sample hence presents more comparable groups in terms of previous means of agricultural NDVI, whereas the secondary samples present more comparable groups in terms of prioritization score.

8.3 Difference-in-differences assumptions

When it comes to assumptions underpinning the difference-in-differences design, Callaway and Sant'Anna discuss irreversibility of treatment, random sampling, limited treatment anticipation and conditional parallel trends (2021). We will go through them in turn, devoting separate sections for the anticipation assumption and the parallel trends assumption. Firstly, irreversibility of treatment would merit special attention if we were to investigate long or medium term impact of rural road upgrades. It could indeed be an issue that roads that were upgraded in one period deteriorated in a later period, such that roads that were once classified as "upgraded" and "treated" should no longer be classified as "upgraded" and "treated" in a later period. We deem the issue of deteriorating roads to be plausible but limited in the setting at hand, where we look at short-term effects of rural road upgrades. We have furthermore ensured that observations classified as belonging to the "treatment" group are not coded as "control" observations in any time period. Likewise, we have ensured that no observations used in the analysis are classified as "treated" in any pre-treatment time period. Secondly, random sampling is not explicitly addressed in our research design, that primarily uses all roads subject to the prioritization exercise by the Mozambique Road Authority within the frames of the World Bank project.

8.3.1 Limited treatment anticipation

Limited treatment anticipation is addressed in our sensitivity analysis, where we code year 2019 as an anticipation year. The coefficient of interest is the interaction term of the dummy for year the 2019 and the treatment dummy. For out primary sample, this coefficient is about the same absolute size as our primary estimates, although negative, and not statistically different from zero on a conventional level. This suggests that there is no anticipatory behaviour in the primary sample.

Turning to the secondary sample, the results are different. When we use roads selected for treatment only, the coefficient is large, statistically significant at the 1% level and negative. When we use roads with positive prioritization score only, the coefficient decreases in magnitude and statistical significance, although it is still large compared to our primary results and of negative sign.

The analysis of anticipation effects for our secondary samples do not suggest anticipatory behavior, as the sign of coefficients are difficult to understand from this perspective. A negative coefficient is what we could potentially observe if many people left agriculture as a consequence of the road upgrades. Given that roads are yet not upgraded by 2019, it is hard to argue that we would see a reallocation of labor out of agriculture as access to labor markets is not yet improved. A positive anticipatory effect, on the other hand, could indicate that people who are able to do so prepare for improved market access by increasing agricultural output. Given that many people find themselves facing constraints in form of finance and access to markets, however, we would be more likely to expect no anticipatory behavior. These findings are thus concerning with respect to the difference-in-differences assumptions.

8.3.2 Pre-treatment trends for primary sample

A critical assumption in difference-in-differences estimations, is the parallel trends assumption (Callaway and Sant'Anna 2021). The assumption is that absent the event known as treatment, the "average outcomes for treated and comparison groups would have followed parallel paths over time" (ibid., p. 200). The counterfactual is of course not observed, but assumed and supported by, for instance, parallel trends pre-treatment.

The pre-treatment trends for our primary sample are firstly shown in figure 12. Due to the large variations in agricultural NDVI across years, we also employ a comparison of pre-treatment trends adjusted for rainfall (figure 13). The variation in agricultural NDVI is reduced when rainfall is controlled for. Pre-treatment trends appear largely parallel when we control for rainfall, although there are some intersections between the lines that otherwise move in tandem. The pre-treatment trends are similar in terms of being parallel without the control for rainfall (figure 12), but not entirely given the steep slope between 2019 and 2020 for the treatment group, and given the U-shaped curve between years 2011-2013 for the control group and the inverted U-shaped curve for the treatment group over the same period. The parallel trends assumption could very well be challenged.

While there might be reason to question the pre-treatment parallel trends assumption, our placebo test that yields a small coefficient (-0.0084) not statistically different from zero on any conventional level, draws some support for the assumption of pre-treatment parallel trends (see table 7).

8.3.3 Pre-treatment trends for secondary sample

Figure 14 shows parallel trends for our secondary treatment and control group, using pre-treatment trends adjusted for rainfall. They are quite similar to figure 13, albeit visually less parallel. There are also some intersections between the lines, which potentially challenges the parallel trends assumption for the secondary treatment and control groups as well.

The placebo tests indeed suggest that parallel trends for our secondary treatment and control groups might be violated. For the selected roads, the coefficient of interest if -0.0240, significant at the 5% level, and correspondingly, -0.0314 for roads with positive prioritization score, significant at the 1% level. For the placebo test to lend support for our parallel trends assumption, these coefficients should not be of economic nor statistical significance.

8.4 Varying grid cell size and distance to city

In this section, we discuss the findings from our heterogeneity analyses.

8.4.1 Varying grid cell size

It is true for both the primary and secondary samples, that the coefficients of interest and their standard errors appear to be somewhat sensitive to grid cell size, and especially sensitive to grid cell sizes that are in the smallest range. It could be that grid cells below a certain size do not capture enough agricultural land for effects in agricultural output to be identified. It could also be that farms most responsive to road upgrades are located further away from the road, and are only included in the analysis when observations grow large enough.

As soon as grid cell size increases to the mid-range, the estimates are however rather stable. Our concern that effects on agricultural NDVI would be attenuated by large grid cells does not appear to be supported. Instead, the estimates appear to be rather robust to differing grid cell sizes above a certain height.

8.4.2 Distance to city

First of all, it should be noted that the sizes of our groups participating in the analysis differ. This warrants a cautious stance when interpreting the standard errors on the coefficients of interest. For the primary sample, there are most observations in the distance category closest to the city, somewhat fewer observations in the mid-distance category and the fewest number of observations are located further away than 20 km from the city. For selected roads, there are most observations in the closest distance category, fewer the the mid-distance category and the lowest number of observations are in the category above 20 km away from the city. For roads with positive prioritization score, the pattern is similar. One way to make groups more similar in size without using terciles[7], would be to just look at observations closer than 10 km and observations further away from the city. We decided to divide our sample into three rather than two categories, because we wanted to see whether there are discrepancies in response to treatment on more levels of remoteness.

For the primary sample, the results suggests that our main regression results could be driven by observations located between 10 and 20 km from the city, although the coefficient displays only a low level of statistical significance. Furthermore, we cannot the reject that the coefficients on the closest and furthest distances are zero. Also, we are unable to reject the null hypothesis of no joint significance of the parameters of interest in this test. Indeed, we lack statistical support to say that distance to city matters for treatment effects. Interestingly, however, we note that the size and direction of the estimates are what we would expect if people living closer to the city allocate out of agriculture as a consequence of rural road upgrades, while people living further away from the city shift agricultural production towards increased commercialization, which is consistent with previous research. Nevertheless, we are unable to draw such a conclusion from our results.

The results from the selected roads tell a similar story, although the coefficient for the group of observations closest to the city is higher than for the primary sample (still, it is not statistically different from zero). The null hypothesis of no joint significance of our parameters of interest is rejected at the 10% level, which means that they appear more statistically significant than for the primary sample, even though 10% is generally considered too low a level for it do be dubbed statistical significance.

The results from the roads with positive prioritization score only somewhat changes the position. The joint significance of our parameters of interest appear more statistically justified (null hypothesis rejected at the 5% level), but now observations closest to the city display a coefficient that is statistically different from zero, albeit only at a low level of significance (10%). In short, the F-test suggest that distance appears more important for roads with positive prioritization score. But if observations closest to the city indeed drive the results, it is not in line with what we would expect if people closer to the city allocated out of agriculture and people further away from the city commercialized their agricultural practices.

[7] We also did the regressions with tercile categories, which did not change previous conclusions of joint significance, but rendered none of the coefficients different from zero at any conventional level. The means of distance to the nearest city are quite similar to the thresholds we set.

8.5 Do rural road upgrades increase agricultural output?

The primary and secondary results appear to contradict each other, but ultimately point to the same conclusion: the rural road upgrades in Nampula and Zambézia had no short-term effect on agricultural output. While our primary results lack statistical significance, they are robust to sensitivity analyses with anticipation and placebo tests. Our secondary results, on the other hand, show larger estimates and display a higher level of statistical significance. They are, however, not robust to our sensitivity specifications with anticipation and placebo tests. Taking this together, we cannot reject that the rural road upgrades we studied had no effect on local agricultural output in the short term.

The hypothesis from the theory suggesting positive effects on agricultural output, is not supported by our results. The hypothesis formulated upon the theory suggesting ambiguous effects on agricultural output, is not rejected by our results. However, in a short-term assessment, solely using remote sensing data that captures visible physical changes, it is not feasible for us to say anything detailed about the mechanisms producing our result.

8.6 Limitations of our study

There are a number of limitations to our study. We will address some key issues in this section.

Firstly, and as was already discussed in section 4.3, we set out to study only the short-term effects of the rural road upgrades in Mozambique in this paper. Even as short-term effects are important to investigate and as there are valid arguments in favor of an estimation using only one year of post-treatment data, we would ideally follow our estimation up with data in the years to come, extending our post-treatment period. Constrained by the fact that such data is simply unavailable at the time being, we leave the question of middle-term and long-term effects to future research(ers).

Secondly, there is the question of how one should view roads. This paper treats roads as individual pieces of infrastructure, rather than a network of infrastructural elements. While we take into account the distance from a grid cell to the nearest city and cluster our standard errors at the road level, we do not acknowledge the fact that roads are connected to each other. One implication is that we might neglect heterogeneity in treatment with respect to network centrality. Another implication is that we do not exercise control of potentially inadvertently treated roads. The first issue we attempt to address in our decomposition of the difference-in-differences estimator with respect to the nearest city. Certainly, this measure does not correspond perfectly to a a measure for network centrality, but it holds information about an observation's location in relation to an urban area. It does not hold information about an observation's proximity to a range of urban centers, on the other hand, which a network centrality variable should hold. An observation could be close to one city but otherwise remote, whereas another observation might be close to many cities. Similarly, our measure of the distance to the nearest city does not take into account the size or "importance" of the city in question. There might be reasons to believe that agricultural and economic activities are different for well-connected areas and less connected areas, and it is a topic that future research could justifiably address. The second issue, of inadvertently treated areas, is not addressed in the scope of this thesis. Observations that lie in between roads, where the closest road is not upgraded but the second-closest road is, could experience some level of treatment effects. It would be interesting to construct such a measure, for a future researcher, to vary the treatment group to also include the inadvertently treated observations.

Thirdly, this thesis sets out to study what is sometimes dubbed a "data poor" region of the world. Collecting agricultural and infrastructural survey data is time consuming and costly for researchers, in comparison to utilizing remote sensing data. A drawback in relying heavily on satellite imagery, which we do for our road upgrade detection and for our agricultural NDVI measures, is that we merely observe that which is "visual". We do not observe prices and quantities of maize sold in the marketplace, for instance. We need to trust that our road upgrade detection works well enough for us to identify treated and untreated roads, and that our agricultural NDVI measure is constructed to appropriately account for agricultural output. In the first case, we see a that the vast majority of roads classified as treatment roads were also selected for upgrade. This should support the functionality of our road upgrade detection algorithm. In the second case, our agricultural NDVI is constructed with inspiration from Asher and Novosad (2020). However, an important limitation to our measure of agricultural NDVI is that we do not discriminate between land types. Ideally, one employs a land cover mask, as done in Peng and Chen (2021), to separate changes in NDVI for different types of land. While there are several published land mask datasets, we were unable to find any products that fit in the time period of this study. While NDVI without land cover masks can be useful in contexts such as the Midwestern United States (Lobell, Thau, et al. 2015), where fields can be a clearly identified, it has proved much more difficult to interpret in contexts such as Uganda, requiring complementary coordinate data of plots to control for problems with "geolocation accuracies and mixed pixels in Sentinel-2" (Lobell, Azzari, et al. 2020). To refine our measure of NDVI, we decided to construct agricultural NDVI as the difference between growing season NDVI and planting season NDVI. There might be some issues with different areas having slightly different planting and growing season; both between Nampula and Zambézia, and within the provinces. We think that fine-tuning the outcome measure to different regions could be a promising avenue for further research on this topic, to accurately capture agricultural output. Survey data is also likely to be a promising complement to remote sensing data.

Finally, if this World Bank project was to be evaluated over the medium or long term, there are modifications to the research design that could be considered. If it would prove difficult to identify a theoretically sound control group satisfying the parallel trends assumption, researchers could consider using a different identification strategy. For instance, panel data can be used in regressions that take fixed effects and control variables into account, instead of using different groups in an experimental set-up. While we found theoretical motivations for using our secondary samples, the empirical analysis conducted in our sensitivity checks rather suggested that the primary sample provided better comparison groups.

9 Conclusion

The role of infrastructure for spatial economic distribution and development has been reinforced and acknowledged in applied economic research, but it still remains contested and contextual as to what it entails. This thesis has identified rural road upgrades in Mozambique over the years 2018-2021, and studied their impact on agricultural output in 2022. We add to the knowledge in the fields of geospatial and development economics by contributing with our road upgrade detection algorithm and a workflow for large-scale data collection. The novelty brought about is the further exploitation of unconventional data sources and techniques to enable economists to study areas of the world that previously might have been overlooked.

In our study, we have not been able to identify short-term effects on agricultural output from rural road upgrades in northern Mozambique. Our primary sample displays a difference-in-differences estimate of low or no statistical significance, although the coefficient is in line with the hypothesis of positive effects on agricultural output. These results are robust to sensitivity analyses. Our heterogeneity analysis does not suggest that distance matters for treatment effects, but then again, we were unable to identify an increase (or decrease) in agricultural output. The secondary samples, that theoretically should provide for good comparisons with respect to agricultural potential, show estimated effects of receiving an upgrade that are both economically larger and statistically more meaningful. These results are however not robust to sensitivity analyses, which suggests issues with the identifying assumptions of no treatment anticipation and conditional parallel trends. As we do not trust the main regressions for our secondary samples, the decomposition with respect to city for the difference-in-differences estimator appears less interesting at this stage.

These results are relevant for policymakers considering road upgrades as a mechanism for subsistence farms to scale up production and transition out of poverty.

Effects on agricultural output are not immediate, although an evaluation considering a longer time period is needed for a complete understanding, and to best formulate policy in this regard. Furthermore, it is likely that policymakers will need the aid of additional data and analyses to discern the mechanisms at play.

The results presented in this thesis were derived from our desks in Stockholm. As geospatial methods are becoming more advanced and as the quality of publicly available imagery improves, the potential for similar studies will only increase. At the same time, evidence collected in the field will always be valuable and complement the data collected remotely. By developing our own road upgrade detection algorithm, we have created a means to collect data that is of great value. With road coordinates being the only required input, our algorithm and data collection process can be applied to any region in the world.

To shed light on the role of rural transport infrastructure in local economic development, more research is surely warranted. Our study suggests that rural road upgrades do not increase agricultural output in the short term. Additional data for years to come in the period post road upgrades, is needed to complete the picture of the agricultural impact of this particular World Bank project.

References

Abadie, Alberto et al. “When should you adjust standard errors for clustering?” In: *NBER Working Paper Series* (2017).

Aggarwal, Shilpa et al. “Market access, trade costs, and technology adoption: Evidence from Northern Tanzania.” In: *NBER Working Paper Series* (2018).

Agrilinks. *How farmers shifted from subsistence to profit without leaving the village*. 2022. URL: https://www.agrilinks.org/post/how-farmers-shifted-subsistence-profit-without-leaving-village (visited on 05/05/2022).

— *Sub-national crop calendars*. 2020. URL: https://www.agrilinks.org/covid-19-mitigation-and-crop-cycle-models/sub-national-crop-calendars?f[0]=sub%7B%5C_%7Dnational%7B%5C_%7Dname%7B%5C%%7D3AMozambique.

Alder, Simon. “Chinese roads in India: The effect of transport infrastructure on economic development.” In: *Working Paper* (2019).

Asher, Sam and Paul Novosad. “Rural roads and local economic development†.” In: *American Economic Review* 110.3 (2020), pp. 797–823.

Baez, Javier E., German Caruso, and Chiyu Niu. “Extreme weather and poverty risk: Evidence from multiple shocks in Mozambique.” In: *Economics of Disasters and Climate Change* 4 (2019), pp. 103–127.

Banerjee, Abhijit, Esther Duflo, and Nancy Qian. “On the road : Access to transportation infrastructure and economic growth in China.” In: *Journal of Development Economics* 145 (2020), p. 102442.

Barrett, Christopher B. “Smallholder market participation: Concepts and evidence from eastern and southern Africa.” In: *Food Policy* 33.4 (2008), pp. 299–317.

Baum-Snow, Nathaniel et al. “Does investment in national highways help or hurt hinterland city growth?” In: *Journal of Urban Economics* 115 (2020), pp. 103–124.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan. “How much should we trust differences-in-differences estimates?” In: *The Quarterly Journal of Economics* 119.1 (2004), pp. 249–275.

Boughton, Duncan et al. “Market participation by rural households in a low-income country: An asset-based approach applied to Mozambique.” In: *Faith and Economics* 50 (2007), pp. 64–101.

Callaway, Brantly and Pedro H.C. Sant’Anna. “Difference-in-Differences with multiple time periods.” In: *Journal of Econometrics* 225.2 (2021), pp. 200–230.

Campenhout, Bjorn van. “The dynamics of smallholder marketing behavior: Explorations using Ugandan and Mozambican panel data.” In: *International Food Policy Research Institute* November (2015).

Casaburi, Lorenzo, Rachel Glennerster, and Tavneet Suri. “Rural roads and intermediated trade: Regression discontinuity evidence from Sierra Leone.” In: *Unpublished* (2013).

CGAP. *National survey and segmentation of smallholder households in Mozambique.* Tech. rep. March. 2016, p. 114. URL: https://www.cgap.org/sites/default/files/Working-Paper-Survey-and-Segmentation-Smallholders-Coted%7B%5C%%7D27Ivoire-Jul-2017.pdf.

Chandra, Amitabh and Eric Thompson. “Does public infrastructure affect economic activity? Evidence from the rural interstate highway system.” In: *Regional Science and Urban Economics* 30.4 (2000), pp. 457–490.

Chiovelli, Giorgio, Stelios Michalopoulos, and Elias Papaioannou. “Landmines and spatial development.” In: *NBER Working Paper Series* (2018).

Donaldson, Dave. “Railroads of the Raj: Estimating the impact of transportation infrastructure.” In: *American Economic Review* 108.4-5 (2018), pp. 899–934.

Donaldson, Dave and Richard Hornbeck. “Railroads and American economic growth: A "market access" approach.” In: *The Quarterly Journal of Economics* 131.2 (2016), pp. 799–858.

Eaton, Jonathan and Samuel Kortum. “Technology, geography, and trade.” In: *Econometrica* 70.5 (2002), pp. 1741–1779.

Faber, Benjamin. “Trade integration, market size and industrialization: Evidence from China’s national trunk highway system.” In: *The Review of economic studies* 81.3 (2014), pp. 1046–1070.

Fafchamps, Marcel and Forhad Shilpi. “The spatial division of labor in Nepal.” In: *The Journal of Development Studies* 39.6 (2003), pp. 23–66.

Food and Agriculture Organization of the United Nations. *Country brief on Mozambique.* 2017. URL: http://www.fao.org/giews/countrybrief/country.jsp?code=MOZ (visited on 05/05/2022).

Glaeser, Edward L. and Janet E. Kohlhase. “Cities, regions and the decline of transport costs.” In: *Papers in Regional Science* 83.1 (2004), pp. 197–228.

Gollin, Douglas and Richard Rogerson. “Productivity, transport costs and subsistence agriculture.” In: *Journal of Development Economics* 107 (2014), pp. 38–48.

Heltberg, Rasmus and Finn Tarp. "Agricultural supply response and poverty in Mozambique." In: *Food Policy* 27.2 (2002), pp. 103–124.

Iimi, Atsushi et al. "Social and economic impacts of rural road improvements in the state of Tocantins, Brazil." In: *Policy Research Working Paper Series* April (2015).

International Potash Institute. *Maize intensification in Mozambique: Demonstrating to farmers the benefits of better land and crop management.* Tech. rep. 2011. URL: https://www.ipipotash.org/uploads/udocs/eifc%7B%5C_%7Dno28-rf2.pdf.

Isaksson, Ann-Sofie and Dick Durevall. "Aid and institutions: Local effects of World Bank aid on perceived institutional quality in Africa." In: *IFN Working Paper* (2021).

Jacoby, Hanan G. "Access to markets and the benefits of rural roads." In: *Economic Journal* 110.465 (2000), pp. 713–737.

Jedwab, Remi and Adam Storeygard. "The average and heterogeneous effects of transportation investments: Evidence from sub-Saharan Africa 1960-2010." In: *NBER Working Paper Series* (2021).

Krugman, Paul. "Increasing returns and economic geography." In: *Journal of Political Economy* 99.3 (1991), pp. 483–499.

Lindgren, Erik, Per Pettersson-Lidbom, and Björn Tyrefors. "The causal effect of transport infrastructure: Evidence from a new historical database." In: *IFN Working Paper* 1407 (2021), pp. 1–25.

Lobell, David B., George Azzari, et al. "Eyes in the sky, boots on the ground: Assessing satellite- and ground-based approaches to crop yield measurement and analysis." In: *American Journal of Agricultural Economics* 102.1 (2020), pp. 202–219.

Lobell, David B., David Thau, et al. "A scalable satellite-based crop yield mapper." In: *Remote Sensing of Environment* 164 (2015), pp. 324–333.

Mather, David, Duncan Boughton, and T. S. Jayne. "Explaining smallholder maize marketing in southern and eastern Africa: The roles of market access, technology and household resource endowments." In: *Food Policy* 43 (2013), pp. 248–266.

Michaels, Guy. "The effect of trade on the demand for skill: Evidence from the interstate highway system." In: *The Review of Economics and Statistics* 90.4 (2008), pp. 683–701.

Moneke, Niclas. "Can big push infrastructure unlock development? Evidence from Ethiopia." In: *Unpublished* June (2020).

Olwande, John et al. “Agricultural marketing by smallholders in Kenya: A comparison of maize, kale and dairy.” In: *Food Policy* 52 (2015), pp. 22–32.

Peng, Cong and Wenfan Chen. “Roads to development? Examining the Zambian context using AI-Sat.” In: *Unpublished* (2021).

Redding, Stephen J. “The empirics of new economic geography.” In: *Journal of Regional Science* 50.1 (2010), pp. 297–311.

Redding, Stephen J. and Matthew A. Turner. “Transportation costs and the spatial organization of economic activity.” In: *Handbook of Regional and Urban Economics* 5.1 (2015), pp. 1339–1398.

ReliefWeb. *Mozambique food security update, December 2006*. 2006. URL: https://reliefweb.int/report/mozambique/mozambique-food-security-update-december-2006.

Shamdasani, Yogita. “Rural road infrastructure & agricultural production: Evidence from India.” In: *Journal of Development Economics* 152.February 2020 (2021), p. 102686.

Statista. *Mozambique’s population by province 2020*. 2021. URL: https://www.statista.com/statistics/1183370/china-population-density-by-region-province/ (visited on 05/05/2022).

Suri, B Y Tavneet. “Selection and comparative advantage in technology adoption.” In: *Econometrica* 79.1 (2011), pp. 159–209.

United Nations. *Food and agriculture organization of the United Nations*. Tech. rep. 2014. URL: https://www.un.org/en/ecosoc/integration/pdf/foodandagricultureorganization.pdf.

— *Mozambique: Thousands continue to flee violence in Cabo Delgado*. URL: https://news.un.org/en/story/2022/03/1114412 (visited on 05/05/2022).

— *Subsistence agriculture for green economy and subsistence life*. URL: https://sdgs.un.org/partnerships/subsistence-agriculture-green-economy-and-subsistence-life (visited on 05/05/2022).

— *Unmasking disparities by ethnicity, caste and gender*. Tech. rep. 2021. URL: https://hdr.undp.org/sites/default/files/2021%7B%5C_%7Dmpi%7B%5C_%7Dreport%7B%5C_%7Den.pdf.

USAID. *Fact sheets - Mozambique*. URL: https://www.usaid.gov/mozambique/newsroom/fact-sheets (visited on 05/05/2022).

Utrikespolitiska institutet [The Swedish Institute of International Affairs]. *Landguiden: Moçambique [The country guide: Mozambique]*. Swedish. Tech. rep. 2022, pp. 1–46. URL: https://www.ui.se/landguiden/lander-och-omraden/afrika/mocambique/.

Vernon Henderson, J. and Sebastian Kriticos. "The development of the African system of cities." In: *Annual Review of Economics* 10.1 (2018), pp. 287–314.

World Bank. *Estimating the impact of road network improvements*. URL: www.worldbank.org/en/research/dime/brief/transport.

Yoshino, Naoyuki, Matthias Helble, and Umid Abidhadjaev, eds. *Financing infrastructure in Asia and the Pacific*. Asian Development Bank Institute, 2018, p. 504.

Appendices

A Descriptive Statistics for Secondary Samples

Table 9: Descriptive statistics for 2018-2020, selected roads only

	Treatment		Control		Difference	
	Mean	SD	Mean	SD	Diff	Diff/$SD_{Control}$
NDVI	0.21	0.09	0.16	0.08	0.05***	0.71
Nearest city	11.15	7.39	11.75	11.53	-0.60	0.05
Road length in kilometers	19.64	11.31	18.87	12.94	0.77	0.06
Grid cell size	7.97	2.36	7.47	2.55	0.51***	0.20
Population in 2019	4,546.97	6,750.46	5,465.76	8,620.27	-918.79***	-0.11
Grid cell population density	639.62	1,352.72	833.09	2,002.59	-193.47***	-0.10
Prioritization score	6.33	5.83	5.53	8.96	0.80**	0.09
Precipitation in mm	5.86	1.14	5.29	1.02	0.58***	0.57
Observations	2,298		1,527		3,825	

* p<0.05, ** p<0.01, *** p<0.001

Descriptive statistics for the pre-treatment time period by secondary treatment and control group.

Table 10: Descriptive statistics for 2018-2020, positive prioritization score only

	Treatment		Control		Difference	
	Mean	SD	Mean	SD	Diff	Diff/$SD_{Control}$
NDVI	0.21	0.09	0.18	0.09	0.02***	0.28
Nearest city	10.81	7.43	14.50	13.00	-3.69***	-0.28
Road length in kilometers	19.94	11.50	21.93	21.92	-1.99***	-0.09
Grid cell size	7.99	2.36	7.75	2.30	0.24**	0.01
Population in 2019	4,651.73	6,861.44	5,242.99	6,907.63	-591.26**	-0.09
Grid cell population density	656.24	1,384.35	785.43	1,851.83	-129.18*	-0.07
Prioritization score	6.74	5.74	8.67	7.16	-1.94***	-0.27
Precipitation in mm	5.85	1.13	5.26	1.08	0.58***	0.54
Observations	2,175		1,671		3,846	

* p<0.05, ** p<0.01, *** p<0.001

Descriptive statistics for the pre-treatment time period by secondary treatment and control group.

Table 11: Descriptive statistics for 2022, selected roads only

	Treatment		Control		Difference	
	Mean	SD	Mean	SD	Diff	Diff/SD$_{Control}$
NDVI	0.27	0.09	0.18	0.08	0.09***	1.13
Precipitation in mm	7.21	1.45	6.47	1.75	0.74***	0.42
Observations	766		509		1,275	

* p<0.05, ** p<0.01, *** p<0.001

Descriptive statistics for the post-treatment time period by secondary treatment and control group.

Table 12: Descriptive statistics for 2022, positive prioritization score only

	Treatment		Control		Difference	
	Mean	SD	Mean	SD	Diff	Diff/SD$_{Control}$
NDVI	0.27	0.09	0.19	0.10	0.08***	0.80
Precipitation in mm	7.15	1.47	5.87	1.72	1.28***	0.74
Observations	725		557		1,282	

* p<0.05, ** p<0.01, *** p<0.001

Descriptive statistics for the post-treatment time period by secondary treatment and control group.

B Digital Appendix: Visual Classification

A report of the visual classifications of roads in our training dataset is available at:

https://github.com/jeffrey-clark/road_upgrades/blob/main/media/appendix_road_composites.pdf