



STEG WORKING PAPER SERIES

REMOTELY-SENSED MARKET ACTIVITY AS A SHORT-RUN ECONOMIC INDICATOR IN RURAL AREAS OF DEVELOPING COUNTRIES

Tillmann von Carnap

FEBRUARY 2022
REVISED JUNE 2022
STEG WP002
SRG 64

Remotely-sensed market activity as a short-run economic indicator in rural areas of developing countries

Tillmann von Carnap^a

June 17, 2022

- [Latest version here](#) -

Abstract

Effective targeting of spatial policies and their rigorous evaluation require relevant and accurate data. With the majority of the world's poor depending on agriculture and informal businesses for their livelihoods, information on these sectors is particularly valuable. I use high-frequency satellite imagery to map rural marketplaces across large geographies and track activity within them in real-time. I show that the method accurately detects existing markets and that measured activity not only correlates with alternative indicators, but also expands their temporal and geographical detail. Focusing on Kenya and Ethiopia, I present applications of the novel method to the effects of lockdowns and violent conflict on market activity.

^aInstitute for International Economic Studies, Stockholm University; tillmann.voncarnap@iies.su.se.

I thank Anna Tompsett, Jakob Svensson, Konrad Burchardi and Jonathan de Quidt for their guidance and support with this project. I also thank Lauren Falcao Bergquist, Joshua Blumenstock, Marshall Burke, Andrea Guariso, Johannes Haushofer, Sol Hsiang, Dahyeon Jeong, John Loeser, Edward Miguel, Lukas Neuhauser, Dennis Ochieng, Moritz Poll and Andrew Wabwire for feedback, technical support or sharing data. Audiences at the Stockholm Development Tea, Berkeley Development Lunch, Berkeley Global Policy Breakfast, Tinbergen Spatial Economics Seminar, Kiel Institute Research Seminar, German Development Economics Conference and ACM Compass Conference provided helpful feedback. I gratefully acknowledge financial support by the STEG initiative.

1 Introduction

There is growing awareness among policymakers and researchers that effective spatial targeting of policies as well as their rigorous evaluation require localized, timely and relevant data (World Bank, 2021). With the majority of the world’s poor depending on agriculture and informal businesses for their livelihoods (ILO, 2018), information on these sectors across space and time is particularly valuable. However, traditional sources often have limited applicability to many policy and research questions as they are collected infrequently (e.g. nationally representative household surveys), do not cover the informal sector (e.g. tax records) or are not able to pick-up short-term fluctuations (e.g. remotely-sensed nightlights or other wealth indicators). This data scarcity is particularly salient for policies that are implemented over large areas, such as road improvements or social safety nets. Here, both local and aggregate effects may be of interest, but unless dedicated data collections are implemented, the coverage of available data may be insufficiently sparse.

In this paper, I present a scalable method generating a high-frequency economic activity measure for rural areas of developing countries. I focus on changes in attendance of rural periodic markets using satellite imagery. Periodic markets are a common and persistent feature of the rural economy across low- and middle-income countries. Typically, buyers and sellers meet once or twice per week in a public space, trading a wide range of goods and services such as vegetables, clothing or kitchenware (Mukwaya, 2016; Kithuka et al., 2020; Balineau et al., 2021). By aggregating otherwise thin supply of rural products and demand for urban goods, these markets function as key locations of trade across regions. Beyond direct trading, marketplaces are where small-scale businesses in sectors such as food processing or tailoring operate, taking advantage of periodic population gatherings there.

Due to markets’ close link with rural economies, trends in their activity - i.e. changes in the presence of buyers and sellers - are informative about local economic conditions: busier marketplaces reflect that attendees either have more income to spend, more goods to sell, or both. Observing this activity at scale across large geographies thus holds potential to better understand local and aggregate impacts of spatial policies. Data collection using remote sensing furthermore allows researchers and policymakers to understand on-the-ground conditions when information is otherwise hard to come by, such as during violent conflict or after extreme weather events.

Tracking rural market activity requires knowledge of marketplaces’ locations. However, comprehensive and up-to-date market maps are rarely publicly available. The method addresses this challenge by first screening imagery of locations where markets may be expected to exist for indicative reflectance patterns. I exploit two characteristic features of periodic markets for their detection. Firstly, stalls, vehicles and crowds are distinctly bright in images taken on market days when compared to the bare ground of the marketplace and surrounding areas on other days. Secondly, markets’ regular occurrence provides a temporal signal to distinguish them from other, idiosyncratic changes in the imagery over time.

The high frequency and deep stack of the imagery I use allows me to exploit these regularities: in

essence, the method screens stacks of imagery for contiguous areas within candidate locations for periodic changes in brightness. I perform various exercises to show that the method indeed (i) reliably identifies marketplaces where they exist and (ii) does not pick up other periodic events such as religious gatherings. While I focus below on applications in Kenya and Ethiopia, the method can in principle be applied in any relatively cloud-free context where open-air periodic markets exist.

Inherent to the novelty of the derived market maps is the absence of ideal validation data. Optimally, one would have data indicating market locations, their days of operation and a measure of their extent for a representative sample of all periodic markets in a given area. To my knowledge however, no such data exist. In order to still assess the accuracy of market detection, I first examine whether the method indeed identifies weekly markets and not other periodic events such as religious gatherings – a low false positive rate of market detection. I show that detected locations of periodic activity are centered on roads or village squares, as opposed to around known places of worship. Secondly, I show that the method identifies a high share of actual markets – a high true positive rate – by comparing the detected markets and market days against a validated ground sample in Western Kenya (Bergquist and Dinerstein, 2020). Currently available ground-truth data does not allow me to test, however, whether the true negative rate – not detecting a market when there is none – is also high, as maps of locations without markets are not available. It is reassuring in this respect that, conditional on detecting a market in the validation dataset, I always confirm the stated market day and do not detect features on other days. On such non-market days, market locations should look similar to places without markets and can thus approximate a sample of the latter.

Equipped with a sample of marketplaces in a region of interest – in my case East Africa –, I then turn to measuring market activity. For this, I measure the density of participants within the detected market area by extracting the median brightness across all pixels contained therein from each image of a given location. I interpret this measure as tracking changes in local GDP – places with larger markets may not necessarily have higher GDP levels than places with smaller or without markets; for a given market, however, days on which measured market activity is high compared to other days likely reflect increased goods exchange and revenue by market participants.

The PlanetScope imagery I use for detection and tracking is available since 2016 and made available in real-time. In Ethiopia and Kenya, I gather between two and five observations for market activity per marketplace and month, depending on seasonal cloud coverage and market frequency. To illustrate that the remotely-sensed market activity can indeed be interpreted as a measure of changes in economic activity, I present three applications using data from 954 marketplaces in Kenya, Ethiopia and Uganda. First, I correlate quarterly changes in market activity between when imagery becomes first available in 2016 and the onset of the COVID-19 pandemic with estimates for sectoral GDP in Kenya. I find that the indicator tracks trends for agricultural GDP both qualitatively and quantitatively, whereas correlations with figures for the trade and transport sectors are less clear. Second, I show how in areas with rainfed, small-scale agriculture, market activity is seasonal overall, and levels of harvest-season market activity

are affected by rainfall during the previous growing season. Thirdly, I illustrate the effects of external events that can be expected to impede economic conditions on market activity, focusing on government-mandated lockdowns during the COVID-19 pandemic and the ongoing war in Tigray and other regions of Ethiopia.

Measuring rural market activity expands the remote-sensing toolkit for data collection in developing countries in various dimensions (see Donaldson and Storeygard (2016) for a review). Compared to commonly used nightlights as an indicator of economic wealth (Henderson et al., 2012; Addison and Stewart, 2015), market activity is available at a sub-yearly frequency, is more likely to pick up short-term changes in economic conditions including downward adjustments, and tracks even non-electrified places (addressing a criticism against nightlights raised by Gibson et al. (2021)). Beyond nightlights, other recent work has constructed local poverty estimates from very-high resolution imagery using machine learning (Jean et al., 2016; Huang et al., 2021; Rolf et al., 2021). Here, market activity provides an indicator related to short-term income flows, as opposed to more long-term wealth proxied by village structures and rooftop materials which are commonly used for wealth prediction in this literature. Finally, rural market activity provides a useful complement to the upstream measures of production provided by the agricultural yield estimation literature (Lobell et al., 2019).

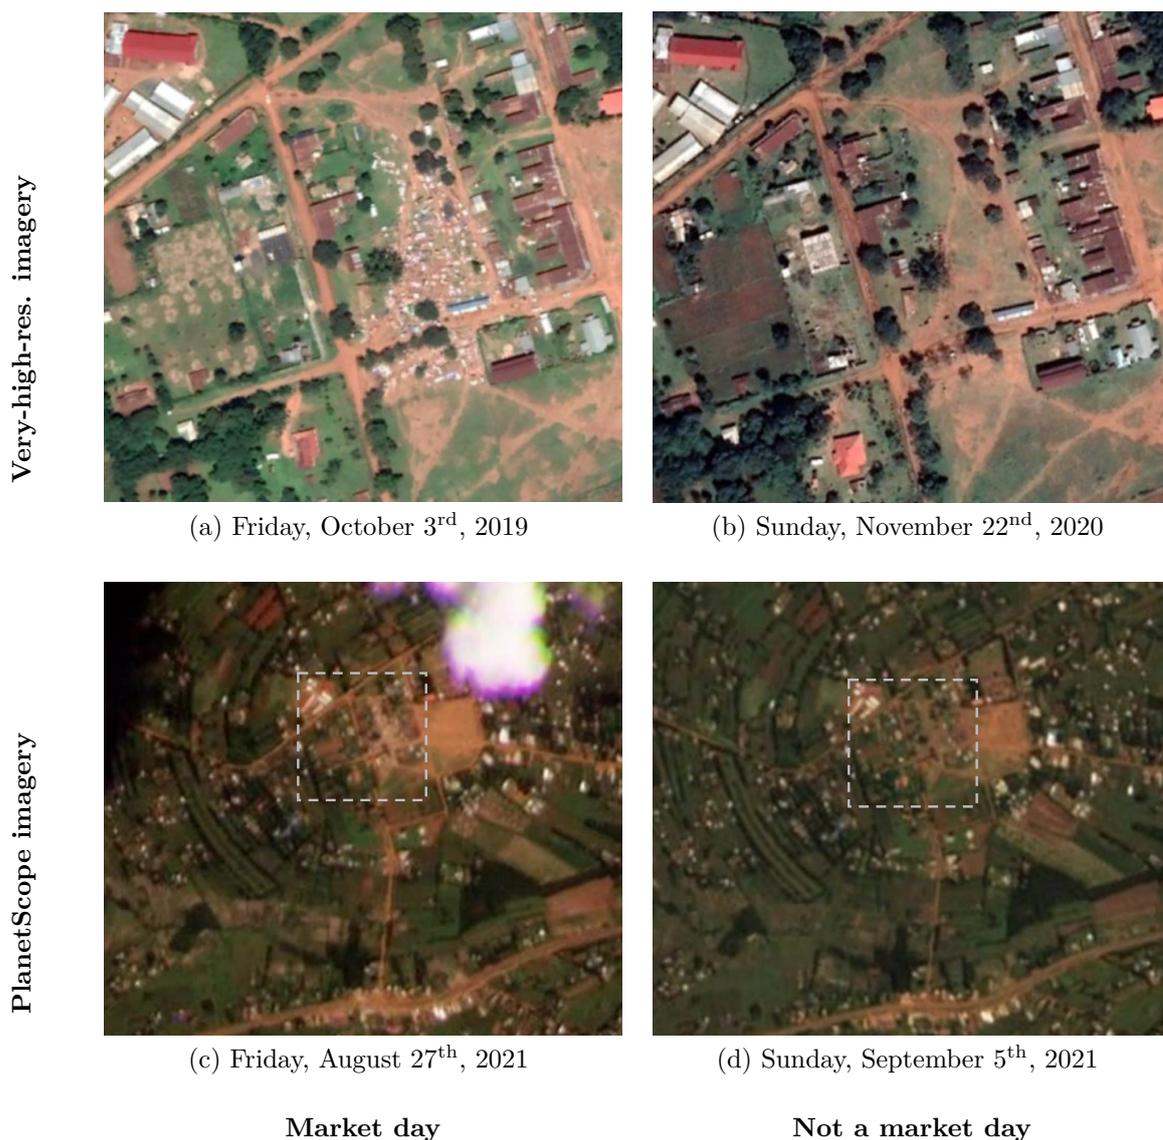
Besides providing a novel economic indicator in otherwise data-scarce environments, the method also generates data to study the role of marketplaces and their interactions more broadly. Previous research on the role of these marketplaces in economic development was hindered by a lack of data. An earlier, cross-disciplinary wave of interest in markets from economists, geographers and anthropologists (see e.g. Hill (1963); Skinner (1965); Wood (1975); Good (1975); Bromley et al. (1975)) died out, with Hill (1963) explicitly mentioning that 'there ought to be much more mapping of marketplaces'. Since then, work on rural marketplaces has been limited to a small number of case studies (e.g., Mukwaya, 2016; Kithuka et al, 2020). Beyond mapping and regarding market-level economic data, available price monitoring programs such as those maintained by the World Food Program or IFPRI typically focus on sets of larger wholesale markets, and lack quantity data. By providing large-scale coverage and local measures of market activity, the method can contribute to our understanding of spatial patterns of economic development.

The rest of the paper proceeds as follows. Section 2 presents the method and validates the market detection. Section 3 presents exercises illustrating the validity and usefulness of the novel activity measure. Section 4 provides an illustration of the method's costs, discusses limitations and concludes.

2 Finding markets & tracking their activity

In the following, I present an overview of the method. I begin by describing how I detect otherwise unmapped periodic marketplaces and explain how I assess the accuracy of the method. I then detail how market activity is measured within the detected markets' extents.

Figure 1: Marketplaces in satellite imagery



Panels (a) and (b) show Tongaren in Bungoma County, Kenya, in very-high resolution imagery from the GoogleEarth archive. Panels (c) and (d) show a different extent of the same location in PlanetScope imagery. The grey squares in panels (c) and (d) indicate the extent shown in panels (a) and (b)

2.1 Finding markets

Figure 1 illustrates the visual pattern underlying the market detection method. It shows in the top row two very-high-resolution images from the Google Earth archive for a Kenyan village, acquired on a Friday and a Sunday. In panel (a), the village square is covered in white, blue and red structures – such as stalls, vehicles and tarps on which goods are displayed – that are typical of periodic markets in the context. While in principle it would be possible to scan an archive of similar imagery for places that look like marketplaces using machine learning, in practice images at the required resolution are only infrequently acquired and made publicly available.

Infrequent captures imply that the few available images may not show a market if they are not taken on market day. This is the case in panel (b): here, the village square appears only as bare ground,

Figure 2: Creating a median composite to resemble a non-market day image



The figure illustrates how a sufficiently large stack of individual images can be collapsed into a representation of a typical day by means of a median composite. Since markets typically take place on less than half of the days of the week, this median composite will resemble a non-market day.

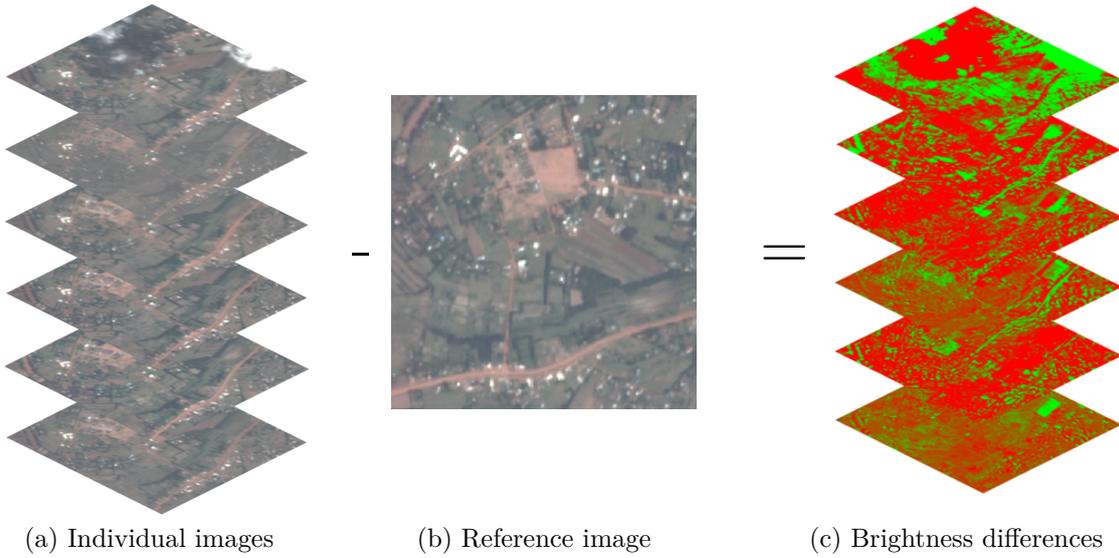
indistinguishable from other open common areas around it. My method therefore uses PlanetScope imagery which compared to images in GoogleEarth has a slightly lower resolution (3 meters per pixel) but a higher, up to daily revisit frequency at around 11am local time¹. This allows me to exploit the relative brightness of markets on market days – evident from comparing panels (a) and (b) – as well as their periodic nature – e.g. taking place every Friday, but not on Sundays – in commercially available imagery with global coverage.

Panels (c) and (d) of Figure 1 show examples of the PlanetScope imagery I employ. While the market is not clearly discernible with the bare eye due to the imagery’s lower resolution, comparing the area within the grey dashed squares in the two images still reveals a brighter patch in the image taken on a Friday compared to the one taken on a Saturday. This again illustrates the basic idea to screen images for changes in brightness that - unlike the patch of cloud visible in the Friday image - occur at a regular frequency.

The premise to look for relative changes in brightness between market and non-market days necessitates the definition of a reference image, ideally showing a day on which the market is not held. Given the lack of comprehensive catalogues of market locations, there is also no widespread information on market schedules, including days where markets are not held. To make progress, I exploit the fact that periodic markets are usually held on less than half of the days in a week. Previewing the mapping exercise across Ethiopia and Kenya, I find that most markets are held on one day (55% of markets in Ethiopia, 70% in Kenya) or two days (37% and 27%) per week, while it is less common for them to occur on three days per week (8% and 3%). The dominance of markets held on less than half the days per week aligns with available evidence from earlier studies of periodic markets in the region (Wood, 1973; Bromley et al., 1975). Since periodic markets thus typically occur at less than half of all days in a week, a median

¹For a given location of interest, I access all imagery taken since the inception of the program in mid-2016. I filter out images with high cloud cover and apply a set of processing steps to make scenes comparable over time.

Figure 3: Create brightness difference representation from raw images



The figure illustrates how the median composite from Figure ?? is subtracted from each individual image to get representations of brightness differences per band. Green (red) pixels indicate areas of positive (negative) brightness deviations relative to the median composite.

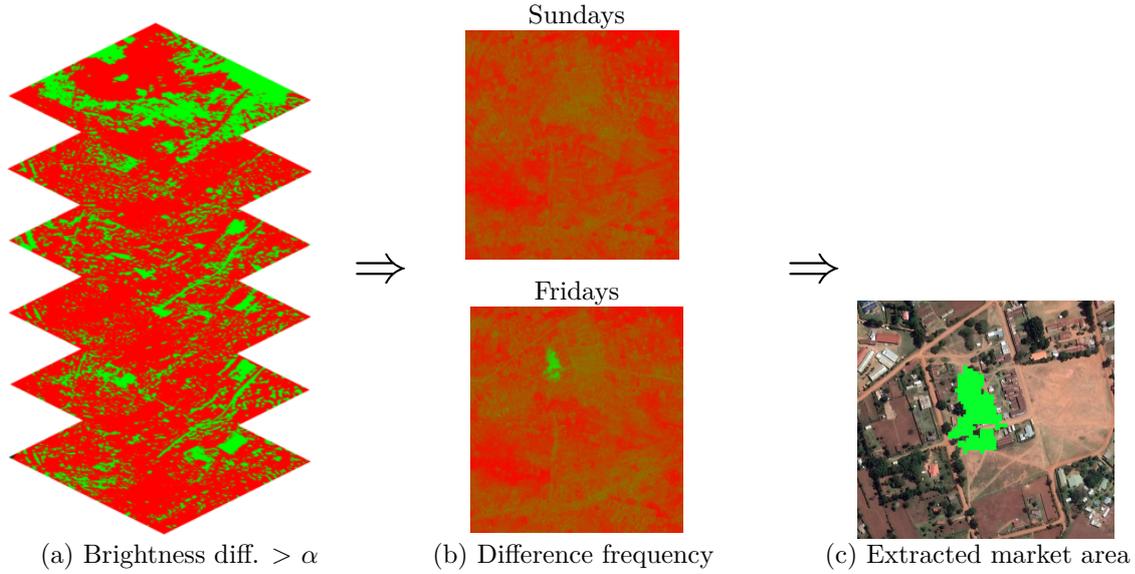
composite of all images within a time interval (e.g. a quarter) will – given a large enough set of images to construct the composite – look like a ‘typical’ day which for the vast majority of places is a non-market day (Figure 2). The composition within a set time interval takes into account that relative brightness of market structures relative to the bare ground may differ across seasons.

Using the median composite as the reference image, I can then define relative changes in brightness for each image in the sample. I here match the bands in the visible spectrum (red, green, blue) between individual and reference images and subtract the latter from the former. This is illustrated for one band in Figure 3. The individual difference images contain both low-level noise in light green – coming from differing image lighting conditions or sensor properties – and high-level differences in bright green. The latter include the market signal, but also idiosyncratic variation between images stemming from, e.g., remaining patches of clouds after pre-processing or harvested fields.

In order to reduce the influence of noise, I simultaneously filter out the low-level noise – coming from image conditions – and reduce the impact of high-level noise – coming primarily from very bright clouds and reflecting rooftops – by converting the continuous difference measure into a dummy representation. The dummy here indicates whether the pixel-level difference exceeds a threshold value α . In order to account for the observation that e.g. a metal rooftop changes reflectance much more in absolute terms than e.g. fields, I define relative strength in terms of pixel-level standard deviations of the differences. Section 2.2 describes how the brightness difference parameter is chosen.

To distinguish between brightness changes from regular market activity or other sources, I rely on the assumption that only markets appear systematically in the same location within an image at a constant periodicity. I can thus filter out remaining high-intensity noise by averaging the relative brightness

Figure 4: Identification of areas with frequent changes in brightness



The figure illustrates how brightness deviations from Figure 3 (column a) are collapsed by day of the week at which they are taken (column b). Green pixels here indicated areas that are relatively bright on a high share of a given weekday in the sample. Pixels exceeding a threshold value are extracted as market outlines, as shown in column c.

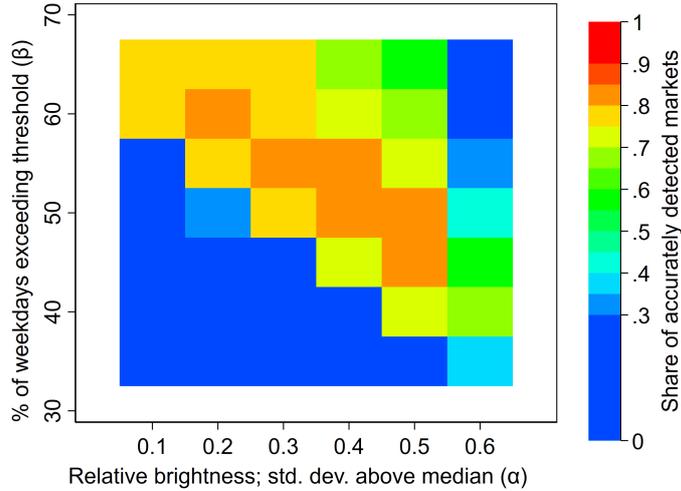
dummies across all incidents of a given weekday in the sample, as illustrated in Figure 4. Here, areas marked in green are those where pixels are relatively bright on a high share of images taken on that weekday. In order to convert these heatmaps into polygons outlining the market areas, I define a second parameter β , defined as the share of weekdays in the sample on which a given brightness threshold is exceeded. Section 2.2 describes how I calibrate this parameter and evaluates whether assuming markets to be the only visible periodic changes is valid.

2.2 Calibration & validation of market detection

Ideally, the method outlined in the previous section would, when deployed over a large number of locations, detect a high share of existing markets – a high true positive rate – and not detect markets in a large share of locations that do not have markets – a high true negative rate. In order to assess this, I optimally would have access to a validation dataset containing precise coordinates, timings and sizes of a representative sample of periodic markets in a given region, as well as a sample of locations without periodic markets. To the best of my knowledge, this is not available, at least publicly for any African country where periodic markets are common. Publicly available market monitoring dataset such as those maintained by actors like the WFP or IFPRI focus on aggregation markets in district capitals or other larger towns. Furthermore, the published data typically do not have measures of market extent or timing.

Without this ideal validation data, I rely for calibration of the method’s parameters and its validation on the set of 60 markets studied in Bergquist and Dinerstein (2020) in western Kenya. The authors specifically sampled periodic markets as opposed to other places of trading and recorded their location

Figure 5: Detection accuracy by parameter combination



The figure shows on the horizontal axis a range of parameter values of what may constitute a brightness change relative to the median composite, measured in standard deviations of brightness per pixel. The vertical axis shows a range of what may constitute a frequent change in brightness over all instances of a given weekday in the sample. The color scale represents detection accuracy in the validation sample from Bergquist and Dinerstein (2020) for a given parameter combination.

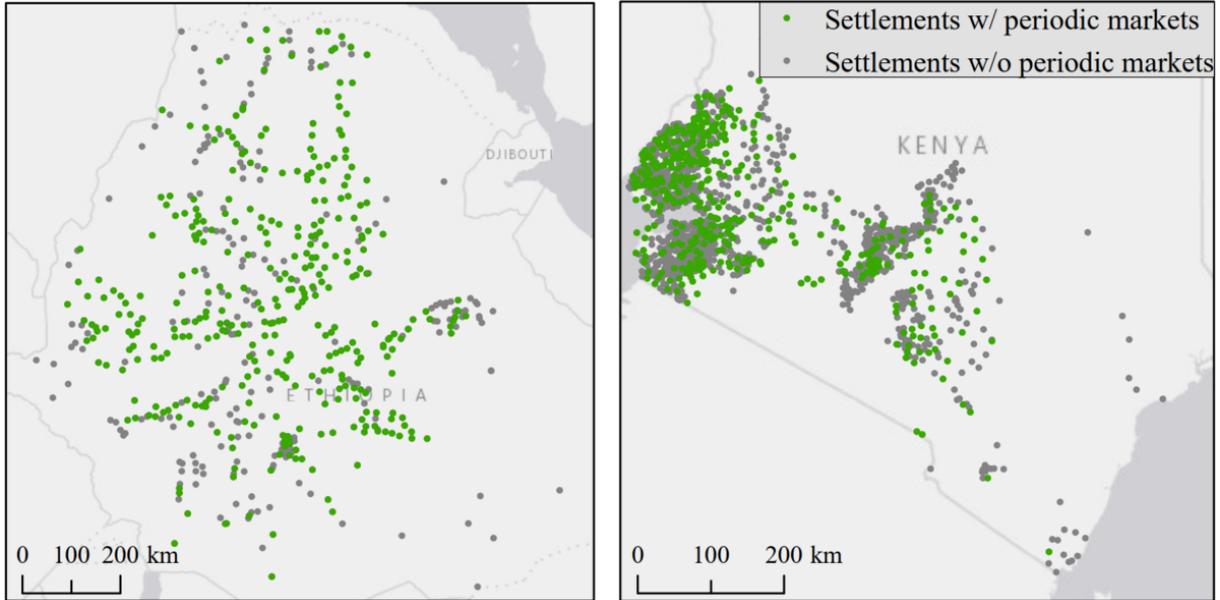
and days of operation². The goal here is to find the parameter combination of α and β that maximizes detection accuracy. I assess parameter combinations as 'good' if the method detects market outlines on the correct days in a large share of locations ('true positives'), and does not detect shapes on other days ('true negatives'). More specifically, I define detection success for each parameter combination as the share of locations in which at least a subset of the stated market days are confirmed and no other days than these.

$$\{\text{detected market days}\}_i \subseteq \{\text{validation market days}\}_i$$

I perform a grid search over a range of possible parameter combinations of market brightness and brightness frequency, summarized in Figure 5. On the axes, I combine various values of the parameters. The colors of the cells illustrate for each parameter combination the share of sample markets I confirm according to the metric above. There is an intermediate parameter range in which the method confirms around 85% of the markets. For high values of α and β , the method does not detect a larger number of markets as image - exclusion errors - whereas for lower values, inclusion errors increase, typically caused by large rooftops or other reflecting surfaces. I henceforth work with the parameter combination giving the highest accuracy in the validation sample ($\alpha = 3$, $\beta = 55$). This parameter combination may not be optimal in contexts with different atmospheric conditions or where periodic markets consist of differently looking structures. To ensure that the method returns a representative sample of market locations, it may be necessary to calibrate the parameters against a context-specific validation sample, collected either on the ground or from secondary sources such as very-high resolution imagery or administrative records.

²The authors selected the markets in their sample in 2016. Moritz Poll kindly collected and shared information on which of these markets were still in operation in 2021. I drop from the validation sample four markets that were found to have permanently ceased operations.

Figure 6: Candidate locations by detection status across Kenya and Ethiopia

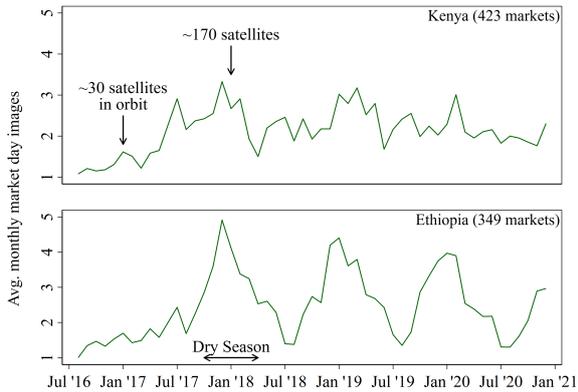


Each dot represents a candidate location where the method has been applied to screen for the existence of a market. I selected candidate locations based on their population size and location along the road network, as identified in very-high-resolution imagery.

I show the detected shapes for each location in the validation dataset in Appendix A, where the underlying very high resolution image suggests that it is indeed a market and not another regular pattern that is identified. This is particularly supported by the observation that if more than one market day is detected for a given location, the extents for the various days typically overlap each other closely. Beyond this evidence, a specific confounder may be places of worship that are also frequented on a weekly basis. To address this concern, I calculate the distance from all 879 registered churches and mosques in the Open Street Map database that fall within the outlines of candidate locations to the detected marketplaces. None of the listed religious buildings falls into the detected market extent. This suggests that religious activities do not regularly spill onto surrounding areas where they may be visible in satellite imagery and confound the market detection.

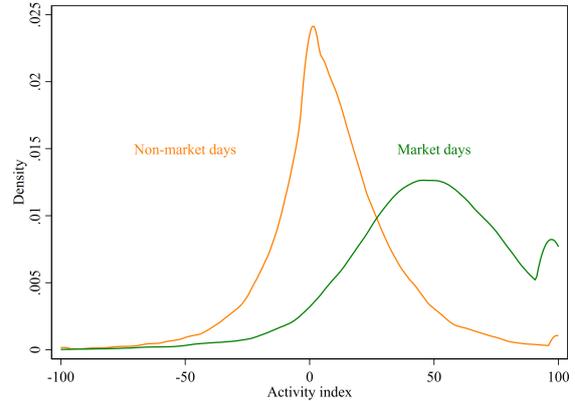
After choosing the method's parameters based on Figure 5, I apply the method to a large set of candidate market locations in Kenya and Ethiopia using GoogleEarth Engine. I select candidate locations based on visual inspection of the GoogleMaps basemap, focusing on clusters of houses along roads. In principle, the definition of candidate locations can be automated using secondary datasets, such as high-resolution population grids and road networks. It is worth noting that the sample is not necessarily representative of all periodic markets in the study countries in a statistical sense, but that it is geographically broad. As of June 2022, the sample covers 439 markets in Ethiopia, 452 markets in Kenya and 63 markets in Uganda, illustrated in Figure 6.

Figure 7: Monthly observations per detected marketplace



The lines indicate the number of monthly observations per market in the sample. Markets with multiple market days can have more than one observation per week.

Figure 8: Measured activity on market and non-market days



The lines show densities of measured market activity between market and non-market days in 2018 and 2019 for Kenyan markets. The index is capped at the 95th percentile across market days, hence the excess mass at the right tail.

2.3 Extracting market activity

After having collected information on market locations and their schedules, I now describe the construction of a measure of market activity over time. Here, I again rely on the observation from very-high-resolution imagery that a marketplace is relatively brighter when more cars, tarps and people are present, compared to when the marketplace is empty. I therefore extract the value underlying the market outline detection above, the maximum deviation in brightness measured in standard deviations for each day (including non-market days) across the three bands. Considering maximum instead of average deviations across bands takes into account that some marketplaces may appear bright on non-market days in visual imagery if the ground consists of bright sand. In these cases colorful market structures may not stand out on average across bands, but should still be markedly different in at least some of them. In order to obtain one observation per day and market location, I compute the median of these deviations across all pixels within each image of the detected market area. I hence also obtain measures of 'activity', i.e. brightness deviations, on non-market days which can be used to normalize activity on market days within a given time period.

An advantage of remotely sensed market activity is its availability throughout the year, weather permitting. Figure 7 shows for the set of markets in Figure 6 the average number of market activity observations obtained per month. The imagery is relatively infrequently available from mid-2016 onwards and increase with the deployment of a large number of satellites in early 2017. It is furthermore evident that there is a seasonality in image availability, stemming from cloud conditions and particularly so in Ethiopia. In turn though, more markets with two or more market days per week in that country increase the likelihood for successful captures during relatively cloud-free months.

The measure is most useful when interpreting its changes over relatively short periods where fluctuations due to external factors are likely to dominate potential secular trends towards fewer and less busy

markets. In particular, a location with a large market is not necessarily more developed than one without if economic activity in the second location is in other sectors. Furthermore, as economies develop, trade may be taking place through more formal networks than those exemplified by periodic markets. Both factors caution against interpreting the level of market activity.

Given the above and in order to make changes comparable across different markets, I index each market’s activity measures to the 95th percentile of activity across observations between mid-2017 and late 2020 and cap the measure at that value. Figure 8 shows density plots of the market activity index for the Ethiopian and Kenyan markets in 2019, separately for measures obtained on market days and on non-market days. It is evident how measured market activity is generally higher on market days, but also that the measure picks up substantial variation within either group that is unlikely to be explained by actual market events. Rather, it reflects residual variation in image conditions after pre-processing, such as cloud shadows or haze. This noise can be addressed by aggregating over multiple markets in an area of interest such as a district, or averaging over longer time periods.

3 Validation of market activity measure

Analogous to the market mapping validation above, the optimal validation data for the market activity tracking is to my knowledge not available at the required temporal and geographical scale. Ideally, one would have panels of attendance counts for a large sample of periodic markets in a region of interest, together with measures of traded quantities and prices in order to translate attendance into an easily interpretable economic quantity. Short of this, I present below three exercises in which the measured indicator displays intuitive variation and where its applicability in terms of temporal and geographic detail are evident.

3.1 Remotely-sensed market activity and sectoral GDP

One possible interpretation for the market activity measure is as a proxy for rural GDP. If rural producers valorize their goods and, symmetrically, rural consumers buy their goods in periodic markets to a large extent, then attendance can be interpreted as an indicator of that process. As a first exercise, I hence compare the market activity measure with quarterly statistics on sectoral GDP provided by the Central Bank of Kenya³. While only available on a national level and not disaggregated by rural and urban origins, this exercise can still shed light on which parts of the economy are reflected in the market activity data.

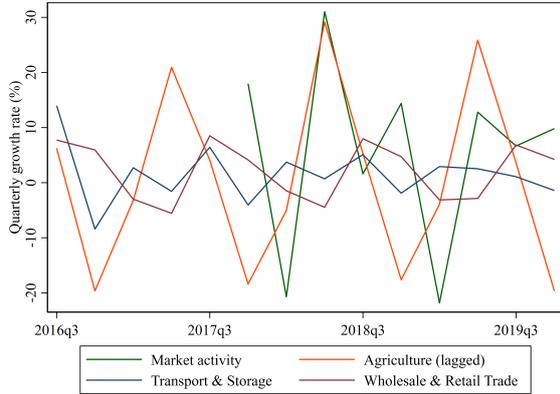
I calculate quarterly growth rates of market activity using the following regression:

$$\text{MktAct}_{t,m} = \sum_{q=2016Q1}^{2020Q4} (\beta_1^q \text{mktD}_{t,m} * \beta_2^q I(t \in T_q)) + \mu_m + \lambda_q + \epsilon_{m,t}$$

$\text{MktAct}_{t,m}$ is the indexed activity in market m on day t , mktD is a dummy indicating whether t is a

³<https://www.centralbank.go.ke/statistics/national-accounts-statistics/>

Figure 9: Quarterly growth rates of market activity and sectoral GDP



The lines indicate quarter-on-quarter growth from national statistics and for remotely-sensed market activity.

Table 1: Correlations between quarterly market activity and sectoral GDP growth rates

Sector	2016Q3-2019Q4
Agriculture	-.7
lagged	.26
Transport & Storage	-.42
lagged	.87
Wholesale & Retail	.27
lagged	-.16

The table indicates raw correlations between pre-pandemic market activity and sectoral growth rates.

market day for location m and $I(t \in T_q)$ assigns days to their respective quarter of a year. Finally, all regressions include a market-fixed effect μ_m and a quarter-fixed effect λ_q . The variables of interest are the β_1^q , yielding quarterly market activity growth rates. I use the market activity measure computed from the detected non-market days to normalize the market-day activity measure in order to account for any seasonal differences in marketplace appearance that may be due to e.g. varying image conditions or vegetation.

Figure 9 presents the estimates compared with quarterly growth rates by sector from the national statistics. I focus on quarters before the onset of the COVID-19 pandemic. It is evident that market activity and agriculture display significant seasonality. I find that the former correlates most with the lagged form of the latter, possibly because of the timing of when official statistics register agricultural production as opposed to when it gets traded. It is striking how the amplitude of agricultural and market activity growth align. I furthermore note how the market activity measure peaks in the second and fourth quarter, whereas agricultural production is not estimated to be relatively high in the latter.

Table ?? shows that market activity also correlates with estimated growth rates in the related transport and retail sectors. However, the precise values here are strongly affected by the short sample period. It is also worth noting that periodic markets are even in rural areas only a part of the trading and transportation sectors and hence correlations cannot be expected to be perfect.

3.2 Remotely-sensed market activity and growing-season rainfall

Absent large-scale irrigation infrastructure, smallholder agriculture continues to be rain-dependent in large parts of Africa, including Kenya and Ethiopia. Rainfall in a given year is thus a key determinant of agricultural production and, consequently, incomes. This link has been used to get at arguably exogenous variation in household income in the development economics literature (e.g. Björkman-Nyqvist (2013); Sarsons (2015)). Periodic markets are where much of agricultural output gets traded and where rural

populations buy necessary goods from farming and other income. One can hence expect market activity to react to rainfall shocks.

As a first step for this analysis, I define locally accurate growing seasons. For this, I rely on the time series of remotely-sensed Normalized Difference Vegetation Index (NDVI) from 2012 until 2020 provided by NASA (Didan and Barreto, 2018). I extract the index for all land within 5km from a market, aggregate markets by pre-2010 Kenyan provinces and Ethiopian regions, and define as the growing season the four months per year where NDVI is highest. I then define as the harvest season the three months following each set of consecutive growing season months. In case of a bimodal rainfall distributions - as is common in parts of both countries - this procedure returns two distinct growing and related harvest seasons.

In order to define rainfall shocks, I match markets to the 1990-2020 time series of remotely-sensed rainfall using the TAMSAT product (Maidment et al., 2017). Rainfall shocks are defined as standard deviations from the long-term mean occurring during the previously defined growing season in a 1km buffer around each market location.

I analyze the relationship using the following regression at the market-season level,

$$\text{MktAct}_{t,m} = \beta_1 \text{mktD}_{t,m} + \beta_2 Rf_{t,m} + \beta_3 Rf_{t,m} * \text{mktD}_{t,m} + \mu_m + \sum_{y=2016}^{2020} (\gamma_y \text{mktD}_{t,m} * I(t \in T_y)) + \epsilon_{m,t}$$

where $\text{MktAct}_{t,m}$ is the median market activity for all measured days within harvesting season t at market m and $Rf_{t,m}$ is rainfall during the preceding growing season. Effects of rainfall are estimated separately for activity measures taken on market and on non-market days, in order to account for any effects of changing market surface due to rainfall. Regressions include market-fixed effects and year-effects separately for market days and non-market days. Note that this is a conservative approach to controlling for market surface conditions since failure to detect some, possibly minor, market days for a given location would also increase estimates.

Table 2 confirms the existence of a positive and significant relationship for Kenya and for Ethiopia where widespread unimodal rainfall implies fewer observations and thus a less precise estimate. A one standard deviation increase in growing-season rainfall is associated with markets being $\frac{1.843}{30.94} \sim 6\%$ more active. I find that the relationship is robust to controlling for overall changes in reflectance related to rainfall and lagged precipitation, mitigating concerns about general rainfall pattern sequences driving the result.

3.3 Additional illustrations - COVID-19 and Tigray War

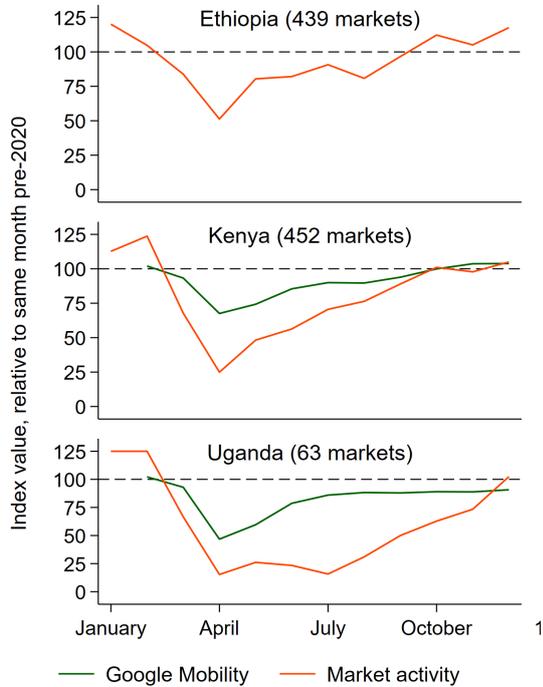
The remote collection of monitoring data has obvious advantages when movement on the ground is constrained. Two examples of such episodes are the COVID-19 lockdowns and the ongoing war in Tigray and other regions of Ethiopia. In such contexts, remote sensing may be used to provide timely and localized information, for example for the allocation of food aid.

Table 2: Growing season rainfall and following harvest season market activity

Dep. Var. Country	Market activity in harvest season			
	Kenya		Ethiopia	
	(1)	(2)	(3)	(4)
Market Day	30.94*** (3.087)	30.63*** (3.091)	44.01*** (2.021)	44.11*** (2.464)
Growing-season rain	0.580* (0.311)	0.629* (0.336)	-0.835* (0.457)	-0.450 (0.480)
Market Day x Growing-season rain	1.843*** (0.622)	1.393** (0.665)	1.625* (0.848)	1.617* (0.890)
Growing-season rain (L1)		0.0236 (0.300)		1.810*** (0.458)
Market Day x Growing-season rain (L1)		-1.069* (0.556)		-0.0725 (0.879)
Non-market day (Constant)	4.004*** (1.506)	3.961*** (1.507)	8.191 (9.763)	8.115 (9.624)
Observations	5,380	5,115	3,505	3,505
R-squared	0.573	0.565	0.657	0.659
Market-FE	yes	yes	yes	yes

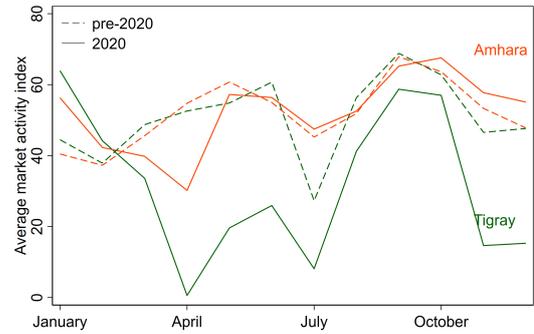
The table reports regression estimates of measured market activity during harvest season in Western Kenya in 2017-2020 on rainfall during the current and preceding rainy season. 'Market day' is a dummy indicating whether a given observation falls onto a market day or not. Robust standard errors in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Figure 10: Market activity and urban mobility during 2020



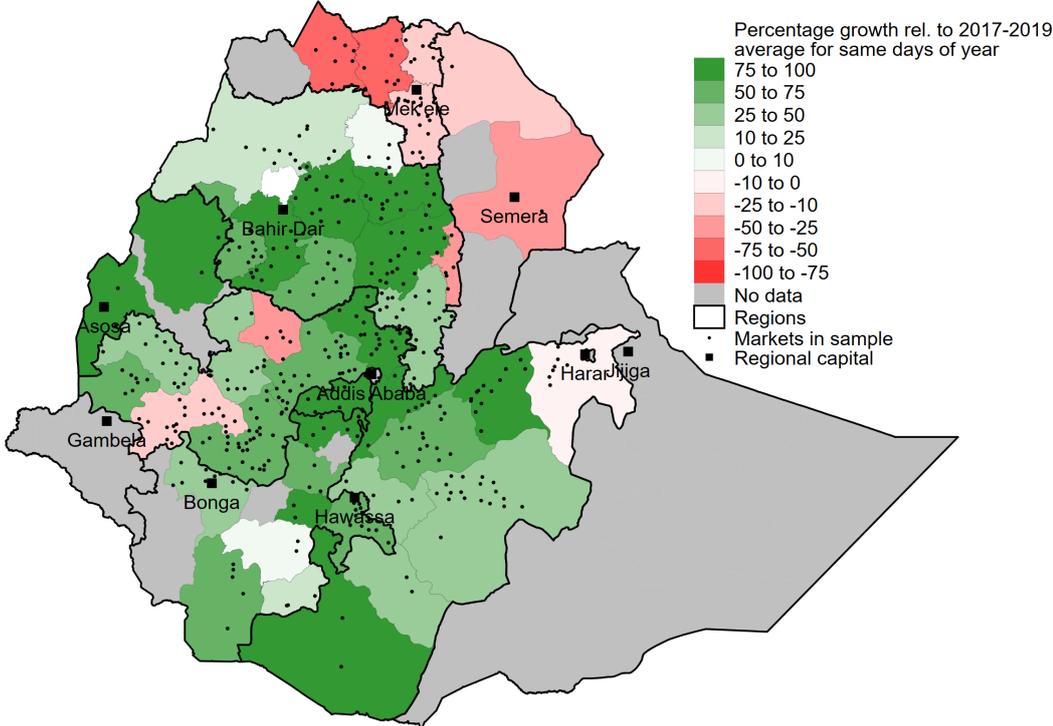
Lines indicate time series for remotely-sensed market activity and Google Mobility data (for workplace, retail and transport visits) during 2020. Values are indexed to pre-pandemic levels, in the case of market activity the 2018-2019 average.

Figure 11: Market activity in Tigray and Amhara, during 2020 and previous years



Lines indicate monthly measures of market activity for markets in the Tigray and Amhara regions of Ethiopia. The dashed lines are averages over 2018 and 2019.

Figure 12: Changes in market activity in Ethiopian zones between Jan 5th and Feb 2nd, 2022 relative to pre-2020 levels



Zones in gray either have less than two detected markets or insufficient images for the given period. Administrative borders are taken from the FAO Global Administrative Unit Layers dataset. Some minor neighboring zones are merged within the same region to increase interpretability.

Figure 10 shows time series of activity across markets in Ethiopia, Kenya and Uganda together with data on imposed movement restrictions (Hale et al., 2021) and measured mobility around workplaces and shopping locations (Google LLC, 2021). It is evident how activity dropped sharply in all three countries in March and April 2020 when movement restrictions and other control measures were implemented. This mirrors data from phone surveys conducted during the early months of the pandemic in Kenya where 67% of respondents in Western Kenya experienced difficulties in accessing markets and overall food insecurity increased (Egger et al., 2021). Importantly, market activity decreased more than otherwise measured mobility and, especially in Uganda, remained subdued for longer.

In Ethiopia, movement restrictions were implemented independently by region. Figure 11 shows how market activity in the Tigray region came to a standstill in April and initially slowly recovered, while the neighboring Amhara region saw a much smaller decline and recovered to pre-crisis levels within a month. When violent fighting broke out with the declaration of a state of emergency on November 4th, 2020, activity again decreased markedly in Tigray only. Figure 12 further exploits the data’s geographic and temporal detail, plotting changes in activity relative to the 2017-19 average by Ethiopian administrative zone. It shows how in early 2022, market activity remained below pre-2020 levels in Tigray, while markets elsewhere appear relatively busy.

4 Discussion & conclusion

High-frequency satellite imagery allows policymakers and researchers to fill a data gap affecting the millions of people worldwide whose livelihoods are linked to agriculture and small-scale trading. The method presented above allows for the detection of periodic marketplaces and their monitoring over time. Periodic markets are important places of exchange and income generation in rural areas of developing countries, and their monitoring allows for the collection of localized indicators in real-time at a low unit cost.

A back-of-the-envelope calculation suggests that with the imagery that can currently be freely accessed over the course of one year from Planet Labs, around 200 markets can be mapped and their activity for the three preceding years measured. With larger, commercially available licenses, activity data for one year at one market can cost as little as 0.1\$, if the method is implemented at scale over large sets of markets (1,000+) or over long time frames.

While in principle the approach outlined above can be transferred to any setting with weekly markets, there are some practical limitations to it. Firstly, frequent cloud cover in humid regions leads to relatively few observations per weekday which complicates market detection as periodic reflection changes from markets are more difficult to distinguish from idiosyncratically bright cloud patches. This problem can partly be mitigated by using longer time series, but high-frequency activity measuring will still be more difficult in these cloudy than in drier contexts. Secondly, smaller, dispersed markets or those whose activity peaks in the afternoon are generally harder to detect. The satellites capture images at around 11am local time for each location, which may or may not be the time when markets are busy. Thirdly, the approach may, if at all, only detect the surrounding areas of covered markets. If one's interest is in also measuring their size, a manual post-processing step may be necessary where one expands the detected shape outside of the covered market to also include the covered market.

Keeping these limitations in mind and while market activity of course captures only part of rural economic activity, there is potential in using it for the evaluation of social policies, such as large-scale cash transfer or infrastructure programs. Furthermore, the data will permit insights during times when traditional data collection becomes infeasible, e.g. during times of violent conflict. The approach presented here relies on accessible imagery and can be scaled across contexts using similar methodologies.

References

- Addison, D., Stewart, B., 2015. Nighttime Lights Revisited The Use of Nighttime Lights Data as a Proxy for Economic Variables. World Bank Policy Research Working Paper 7496.
- Balineau, G., Bauer, A., Kessler, M., Madariaga, N., 2021. Food Systems in Africa: Rethinking the Role of Markets. Africa Development Forum Series. A copublication of the Agence française de développement and The World Bank, Paris/Washington DC.
- Bergquist, L.F., Dinerstein, M., 2020. Competition and Entry in Agricultural Markets: Experimental Evidence from Kenya. *American Economic Review* 110, 3705–3747.
- Björkman-Nyqvist, M., 2013. Income shocks and gender gaps in education: Evidence from Uganda. *Journal of Development Economics* 105, 237–253. URL: <http://dx.doi.org/10.1016/j.jdeveco.2013.07.013>, doi:10.1016/j.jdeveco.2013.07.013.
- Bromley, R.J., Symanski, R., Good, C.M., 1975. The Rationale of Periodic Markets. *Annals of the Association of American Geographers* 65, 530–537. doi:10.1111/j.1467-8306.1975.tb01060.x.
- Didan, K., Barreto, A., 2018. VIIRS/NPP Vegetation Indices 16-Day L3 Global 500m SIN Grid V001 [Data set]. NASA EOSDIS Land Processes DAAC doi:<https://doi.org/10.5067/VIIRS/VNP13A1.001>.
- Donaldson, D., Storeygard, A., 2016. The View from Above: Applications of Satellite Data in Economics. *Journal of Economic Perspectives* 30, 171–198. doi:10.1257/jep.30.4.171.
- Egger, D., Miguel, E., Warren, S.S., Shenoy, A., Collins, E., Karlan, D., Parkerson, D., Mobarak, A.M., Fink, G., Udry, C., Walker, M., Haushofer, J., Larrebourg, M., Athey, S., Lopez-Pena, P., Benhachmi, S., Humphreys, M., Lowe, L., Meriggi, N.F., Wabwire, A., Davis, C.A., Pape, U.J., Graff, T., Voors, M., Nekesa, C., Vernot, C., 2021. Falling living standards during the COVID-19 crisis: Quantitative evidence from nine developing countries. *Science Advances* 7. doi:10.1126/sciadv.abe0997.
- Gibson, J., Olivia, S., Boe-Gibson, G., Li, C., 2021. Which night lights data should we use in economics , and where ? *Journal of Development Economics* 149. URL: <https://doi.org/10.1016/j.jdeveco.2020.102602>, doi:10.1016/j.jdeveco.2020.102602.
- Good, C.M., 1975. Periodic Markets and Traveling Traders in Uganda. *Geographical Review* 65, 49. doi:10.2307/213833.
- Google LLC, 2021. Google COVID-19 Community Mobility Reports. URL: <https://www.google.com/covid19/mobility/>.
- Hale, T., Angrist, N., Goldszmidt, R., Kira, B., Petherick, A., Phillips, T., Webster, S., Cameron-Blake, E., Hallas, L., Majumdar, S., Tatlow, H., 2021. A global panel database of pandemic policies

- (Oxford COVID-19 Government Response Tracker). *Nature Human Behaviour* 5, 529–538. URL: <http://dx.doi.org/10.1038/s41562-021-01079-8>, doi:10.1038/s41562-021-01079-8.
- Henderson, J.V., Storeygard, A., Weil, D.N., 2012. Measuring economic growth from outer space. *American Economic Review* 102, 994–1028. doi:10.1257/aer.102.2.994.
- Hill, P., 1963. Markets in Africa. *The Journal of Modern African Studies* 1, 441–453. doi:10.1017/S0022278X00001361.
- Huang, L.Y., Hsiang, S., Gonzalez-Navarro, M., 2021. Using Satellite Imagery and Deep Learning to Evaluate the Impact of Anti-Poverty Programs, 39.
- ILO, 2018. Women and men in the informal economy: a statistical picture. volume 3. doi:10.1179/bac.2003.28.1.018.
- Jean, N., Burke, M., Xie, M., Davis, W.M., Lobell, D., Ermon, S., 2016. Machine Learning To Predict Poverty. *Science* 353, 790–794. URL: <http://science.sciencemag.org/www2.lib.ku.edu/content/sci/353/6301/790.full.pdf>.
- Kithuka, D.M., Irandu, E.M., Moronge, J.M., 2020. The role of periodic markets in provision of manufactured goods and creation of income in rural areas: a case study of selected markets in Mukaa subcounty, Kenya. *International Journal of Scientific Research and Innovative Technology* 7, 65–79.
- Lobell, D., Azzari, G., Burke, M., Gourlay, S., Jin, Z., Kilic, T., Murray, S., 2019. Eyes in the Sky, Boots on the Ground: Assessing Satellite- and Ground-Based Approaches to Crop Yield Measurement and Analysis in Uganda. *American Journal of Agricultural Economics* doi:<https://doi.org/10.1093/ajae/aaz051>.
- Maidment, R.I., Grimes, D., Black, E., Tarnavsky, E., Young, M., Greatrex, H., Allan, R.P., Stein, T., Nkonde, E., Senkunda, S., Alcántara, E.M.U., 2017. A new, long-term daily satellite-based rainfall dataset for operational monitoring in Africa. *Scientific Data* 4, 170063. URL: <http://10.0.4.14/sdata.2017.63>, doi:<https://doi.org/10.1038/sdata.2017.63>.
- Mukwaya, P., 2016. The dynamics and resilience of periodic markets in the Mt Elgon Region Uganda: a geographical perspective, in: *Rural-Urban Dynamics in the East African Mountains*, pp. 235–271.
- Rolf, E., Proctor, J., Carleton, T., Bolliger, I., Shankar, V., Ishihara, M., Recht, B., Hsiang, S., 2021. A generalizable and accessible approach to machine learning with global satellite imagery. *Nature Communications* 12. URL: <http://dx.doi.org/10.1038/s41467-021-24638-z>, doi:10.1038/s41467-021-24638-z, arXiv:2010.08168.
- Sarsons, H., 2015. Rainfall and conflict: A cautionary tale. *Journal of Development Economics* 115, 62–72. doi:10.1016/j.jdeveco.2014.12.007.

Skinner, G.W., 1965. Marketing and Social Structure in Rural China: Part I. *The Journal of Asian Studies* 24, 363–399. doi:10.2307/2050342.

Wood, L.J., 1973. The temporal efficiency of the rural market system in Kenya. *East African Geography Review* .

Wood, L.J., 1975. The Functional Structure of a Rural Market System. *Geografiska Annaler. Series B, Human Geography* 57, 109. doi:10.2307/490545.

World Bank, 2021. *World Development Report 2021 - Data for better lives.*

Appendix A - detected market shapes from validation sample

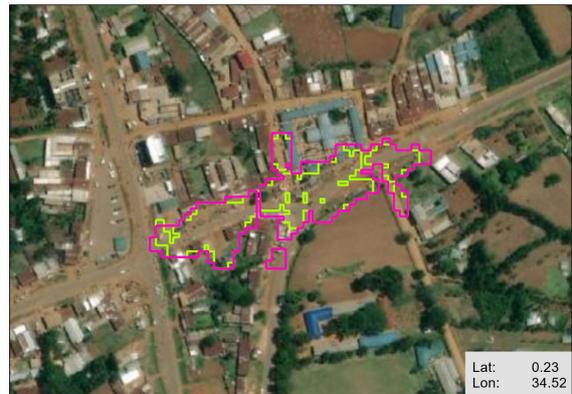
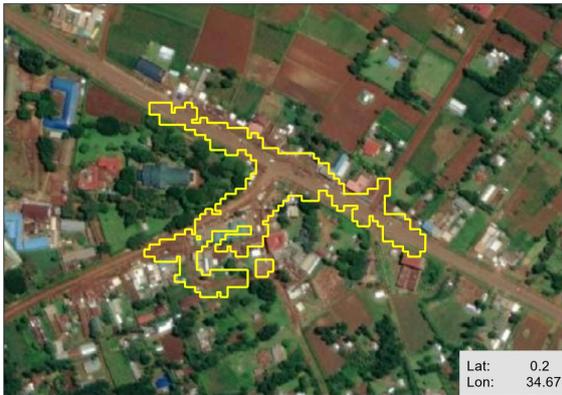
Extent of market by market day

Legend: Sunday (red), Monday (orange), Tuesday (yellow), Wednesday (light green), Thursday (cyan), Friday (blue), Saturday (magenta)



Extent of market by market day

Sunday Monday Tuesday Wednesday Thursday Friday Saturday



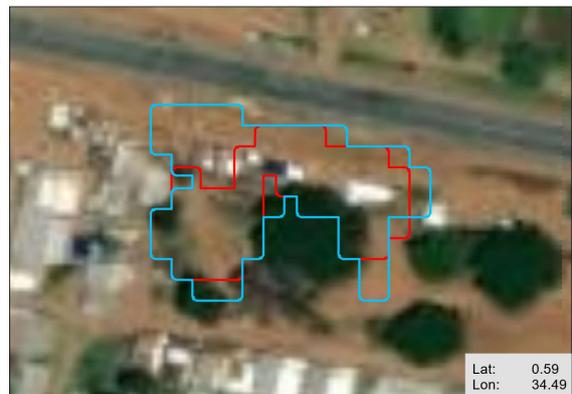
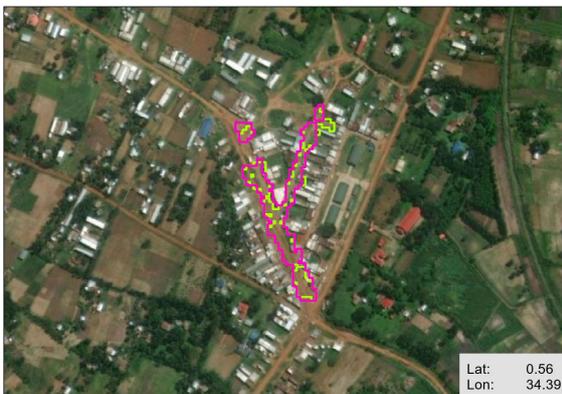
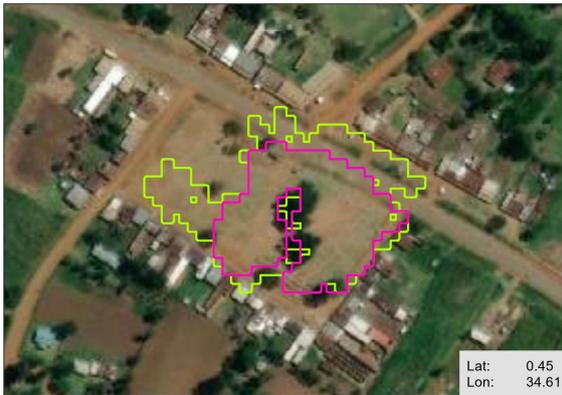
Extent of market by market day

Sunday Monday Tuesday Wednesday Thursday Friday Saturday



Extent of market by market day

Sunday Monday Tuesday Wednesday Thursday Friday Saturday



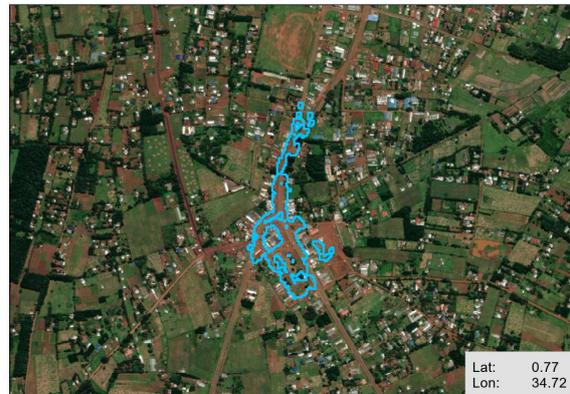
Extent of market by market day

Monday Tuesday Wednesday Thursday Friday Saturday



Extent of market by market day

Sunday Monday Tuesday Wednesday Thursday Friday Saturday



Extent of market by market day

Sunday Monday Tuesday Wednesday Thursday Friday Saturday



Extent of market by market day

☐ Sunday ☐ Monday ☐ Tuesday ☐ Wednesday ☐ Thursday ☐ Friday ☐ Saturday

