

SURFING A WAVE OF ECONOMIC GROWTH

Thomas McGregor* and Samuel Wills^{†*}

March 2017

Abstract

We investigate whether the geographic determinants of growth extend to natural amenities. We combine data on spatial and temporal variation in the quality of over 5000 surf breaks globally with data on local economic performance, proxied by night-time lights. We document a strong association between natural amenity quality and local economic development. Economic activity grows faster near good surf breaks; following the discovery of new breaks, or the technology needed to ride them; and during El Niño events that generate high-quality waves. The effects are concentrated in nearby towns and emerging economies, and population changes are consistent with tourism.

JEL codes: O13, O44, O47, Q26, Q51, Q56, R11, R12

Key words: Natural amenities, economic growth, new economic geography, natural advantages, tourism, surfing, night-time lights.

*Department of Economics and Oxford Center for the Analysis of Resource Rich Economies, University of Oxford, UK. [†]School of Economics, University of Sydney and CAMA, ANU, Australia. The authors would like to thank Tiho Ancev, Julia Bird, Tony Butt, Damian Clarke, Matthew Collin, Mark Hemer, Sophocles Mavroeidis, Steven Poelhekke, Simon Quinn, John Romalis, Brock Smith, Gerhard Toews, Rick van der Ploeg, Tony Venables, Pierre-Louis Vezina, Tim Willems, Justin Wolfers, and Dean Yang for their comments, though all errors are the authors' own. Samuel Wills would like to thank the Economic and Social Research Council for financial support [grant number ES/K009303/1]. Corresponding email: samuel.wills@sydney.edu.au

1 Introduction

Geography has long been thought to play a role in economic growth, because some places enjoy natural advantages over others. These advantages may be direct, as rivers facilitate trade and rich soil makes farms more productive; or indirect, as nice environments make places more desirable to live. The indirect effects are “natural amenities”, and there is no consensus on whether they are important for growth. This paper addresses that gap.

The difficulties with studying natural amenities are in measurement and identification. Natural amenities like a nice view and clean air do not have a market price, so their quality is typically inferred from wages (Roback 1982, 1988) or house prices (Rosen, 1974, Chay and Greenstone, 2005). This is infeasible on a global scale. Even if an amenity’s quality can be measured, it is difficult to identify its contribution to economic growth. For example, economies with a clement climate grow faster (e.g. Deller et al., 2001; Cheshire and Magrini, 2006; Wu and Gopinath, 2008; Glaeser and Gottlieb, 2009), but it is difficult to identify the role of climatic amenity versus other channels, like agriculture.¹ As a result a recent review finds that “the evidence on any positive role of landscape amenities for local economic development... remains limited” (Waltert and Shlapfer, 2010).

This paper uses three clean natural experiments to identify how a particular natural amenity contributes to economic growth... surf breaks.² Surf breaks are well-suited to this type of study because their quality is exogenously determined by a finely balanced combination of weather and sea-floor geography (bathymetry), both locally and at distance. Two locations may be metres apart but of vastly different appeal to surfers, which we can exploit for identification. We measure quality directly through an independently verified rating system from a unique global database of 5000+ locations, combined with data on El Niño events and global wave heights.³ As well as their value as an identification strategy, surf breaks are also a scarce and valuable amenity to more than 35 million surfers (The Economist, 2012); they continue to be discovered (Noble, 2017) and built (Dean, 2016); and will continue to grow in popularity as populous, wave-rich economies like Brazil and Indonesia consume more leisure.⁴ To our knowledge this provides the first

¹In contrast, there is some evidence that endogenous amenities affect growth. For example, Cullen and Levitt (1999) find that crime drives urban flight, and Diamond (2016) finds that endogenous amenities like shops, transport, and the quality of schools and jobs have fuelled the sorting of skills in US cities.

²A note on terminology. A “surf break” is a location. The quality of a break is determined by geography and is generally stable over time, unless there has been some human intervention (see Section 4). A “wave” is an individual pulse of energy that propagates across the surface of the ocean, and which may be the fleeting subject of a surfer’s sport. Groups of waves are “swells”. A wave’s quality depends on the wind at its source and end, and the break where it is ridden (see Section 3.1).

³Previous studies have estimated the economic impact of surfing at individual locations using travel costs (Coffman and Burnett, 2009), surveys (Lazarow, 2009), and hedonic pricing (Scorse et al., 2015).

⁴This trend is illustrated by Brazil winning the men’s World Championship Tour for the first time in 2014, and again in 2015. Surfing will also appear at the Olympics for the first time in Tokyo, 2020.

global, spatially detailed study of natural amenities.

We measure economic activity using two detailed and geographically disaggregated datasets on night-time lights and population. The first records the amount of light emitted at night around the globe at approximately a 1km^2 resolution annually, from 1992-2013. This has been used by many recent studies as a geographically disaggregated proxy for economic activity (Chen and Nordhaus, 2011; Henderson et al., 2011; 2012; Donaldson and Storeygard, 2016); covering institutions (Michalopoulos and Papaioannou, 2013; 2014), political favouritism (Hodler and Raschky, 2014), infrastructure investment (Jedwab et al., 2015; Jedwab and Moradi, 2016) and poverty (Smith and Wills, 2016) amongst others. The second, from LandScan, records global population at a 1km^2 resolution annually, by interpolating sub-national population counts using satellite images of buildings, roads and land-cover, amongst other things.

The first experiment asks the question, “do places with high quality surf breaks grow faster than those with low quality breaks”? To do this we exploit finely-grained cross-sectional variation in the quality of surf breaks in a difference-in-difference specification. We find that the lights in the 5km around 3- and 4-star breaks (out of 5), grew 0.4 percentage points faster per annum than the area around 1-star breaks from 1992-2013; and there were positive and diminishing spillovers out to at least 50km^5 . 4-star breaks have the largest effect on growth, as 5-star breaks require particular expertise to surf. This is worth up to an extra \$2.45 million in output per break per year for the surrounding 10km, or \$4.00 billion in local effects globally.

Experiment I also finds that areas of existing activity benefit the most, with the towns closest to good surf breaks growing 0.5-0.6 percentage points per annum faster than those near bad breaks. This suggests that natural amenities exacerbate path dependence in the location of economic growth. The effects on growth were also strongest in emerging markets. We find that the permanent population falls near good breaks, which is consistent with tourists driving out locals, though this masks a reallocation of people from rural to urban areas. The results are robust to a variety of controls, including shortcomings in the lights data, omitted geographic variables, selection effects and spatial correlation.

The second experiment asks, “does discovering a new surf break (or seeing one disappear) alter nearby economic growth”? We do this using two approaches. The first conducts event studies on three recent discoveries of surf breaks, due to competitions organised by Surfer Magazine and the World Championship Tour; and two disappearances, due

⁵An earlier working paper version of this paper, publicised in The Economist (August 2016), BBC World (August 2016) and the Financial Times (January 2017) amongst others, used a less conservative polynomial model (see Section 3.2). In this version of the paper, we include three natural experiments instead of one, and focus on the most conservative results.

to construction of a coastal road and dredging of a river-mouth. We find that lights grew (or fell) 2.2 percentage points per annum faster after discovery (or disappearance), relative to the global average. The second conducts an event study around a technological discovery: Rip Curl's 2007 invention of battery-heated wetsuits. We find that lights grew 2.7 percentage points faster than the global average in the 5km surrounding 83 cold-water breaks above 55 degrees latitude. This suggests that discovering a new natural amenity, or destroying an existing one, may have a large short-term effect on growth.

Finally, the third experiment asks, "what happens when the surf is good"? Generally, bigger waves are better, but they must be generated at long range to provide space for the swell to organize as it travels, and for the weather at the surf break to not be affected by the originating storm. To capture this we combine monthly data on wave heights with information on El Niño events, which produce famously good long-range swells in the Pacific. Using a triple-interaction approach we estimate the effect of large waves, during El Niño years, at high-quality breaks. We find that, during El Niño events, a one standard deviation increase in wave height increases light growth in the 50km surrounding 4-star and 5-star breaks by 5.6 and 3.9 percentage points, relative to 1-star breaks.

Collectively these three experiments show that natural amenities raise economic growth. Natural amenities, like high-quality surf breaks, are usually not paid for directly. Instead, they augment the productivity of physical capital, by providing a basis for recreational industries and tourism; and the productivity of labour, by attracting migrants with footloose incomes, workers willing to accept lower wages, and entrepreneurs willing to accept lower profits (see review by Waltert and Shlapfer, 2010). In a standard Ramsey growth model, natural amenities will increase the stock of natural capital, and in turn the steady-state level of output. This encourages investment and raises the growth rate along the transition path. Experiment II shows that the increment to growth is larger when the amenity is initially discovered, and Experiment I shows that the effect on growth can persist for decades as surfing originally became popular in the 1960s.

High quality natural amenities may also help solve the coordination problem behind locating sectoral clusters (Rodrik, 1996; 2004). For example, along a stretch of coast a number of beaches may be suitable for a tourism cluster. However, for the cluster to form many independent public and private agents must decide to invest in the same location at the same time. A high quality surf break, patronised by intrepid surfers, provides the focus for this investment. It can then grow to support a broader, non-surfing tourism industry, like Byron Bay in Australia, Jeffreys Bay in South Africa, Taghazout in Morocco, and Arugam Bay in Sri Lanka, which all started as small surfing towns. Future work might disentangle whether natural amenities primarily augment capital, labour, or solve an investment coordination problem; though the focus of the present work is just to establish whether they affect growth at all.

This work contributes to the “geography hypothesis”, which has a long history in explaining the pace of economic growth. This includes channels like suitability for agriculture (Marshall, 1890), exposure to disease (Diamond, 1999; Sachs 2000, 2001), coastal access (Gallup and Sachs, 2000), and frequency of natural disasters (Hsiang and Meng, 2015). It is typically contrasted with the “institutions hypothesis” (North, 1989; 1990; Acemoglu et al., 2002), which argues that institutional technology is the major driver of growth. By using spatial and temporal variation in the quality of natural amenities at a fine scale we find that geography does play a role in explaining local economic growth. This is solely due to spillovers as surfing does not contribute directly to economic activity. Faber and Gaubert (2015) find a similar result, by studying the effects of tourism in Mexico.⁶ We also find that the importance of geography varies by institutional quality, measured using political stability and ease of doing business indices, which provides some evidence to reconcile these competing hypotheses. At a local level both geography and institutions matter, and so does the way they interact.

It is also well-established that natural geography plays a role in determining the location of economic activity. The New Economic Geography literature draws a distinction between “first-nature geography”, concerning natural advantages, and “second-nature geography”, concerning endogenous forces of agglomeration and dispersion (see Fujita et al., 2001; Redding 2009, 2010). Natural advantages, like rivers, ports and resource endowments, have been shown to cause agglomeration as they directly reduce trade and input costs (Ellison and Glaeser, 1997; 1999; Ellison et al., 2010; Redding, 2010; Redding and Rossi-Hansberg, 2016). Natural amenities, like coastlines, mountains and lakes, have also been found to anchor the location of high-income suburbs (Lee and Lin, 2015). Path dependence is important, as natural advantages have persistent effect after they become obsolete, like portage sites near rivers (Bleakley and Lin, 2012; 2015), and existing agglomeration can prevent natural advantages being exploited (Michaels and Rauch, 2016). This study finds that the benefits from surf breaks accrue to nearby towns, suggesting that natural advantages benefit existing areas of economic activity. We also find that when breaks are discovered or disappear, economic growth hastens or slows accordingly.

Finally, this work also contributes to the recent literature on the local effects of natural resources (see reviews by Cust and Poelhekke, 2015 and van der Ploeg and Poelhekke, 2016). These tend to focus on the effects of non-renewable resources with market prices, like oil, gas and minerals (both short run, e.g. Aragon and Rud, 2013; Caselli and Michaels, 2013; Allcott and Kenniston, 2014; and very long run, e.g. Dell, 2010). Complementing those studies we find evidence that non-market, renewable resources can also have a positive impact on the local economy, in both the short and long run.

⁶That study exploits spatial variation in beach quality in Mexico, instrumented by sand colour and offshore islands amongst other measures. In contrast we directly measure the quality of surf breaks around the world, and exploit both spatial and temporal variation in quality across three separate experiments.

The paper proceeds as follows. Section 2 introduces the data and outlines how this is used to measure economic activity near surf breaks. Section 3 exploits cross-sectional variation in the quality of surf breaks by asking the question “do good surf breaks contribute more to economic growth than bad breaks?”. Section 4 exploits temporal variation in surf break quality by asking, “does discovering a new break alter economic growth?”. It also considers breaks that have disappeared, and the effects of a discovery in wetsuit technology. Section 5 interacts both cross-sectional and temporal variation in quality by asking, “what happens when the surf is good?”. Section 6 concludes by discussing policy implications and extensions.

2 Data and Measurement

We use two main datasets to study how surf breaks affect local economic growth: on the location and characteristics of over 5000 surf breaks and on night time light emissions at a 1km² resolution. We use local population estimates to identify nearby towns and cities and study population changes; measures of urban areas and political/economic activity to add granularity to results; and data on atmospheric pressure and wave heights to study how good, large waves interact with our headline results.

2.1 Surf breaks

WannaSurf (www.wannasurf.com) is an online “world surf spot atlas”. It records the location, quality, type, accessibility, coastal and oceanic characteristics of 5,288 surf breaks around the world, which we collect using Python (Figure 2.1). Of these we drop 137 for which the data on quality is either missing or rated 0 stars (“choss”), leaving 5,151 breaks in our dataset.

The main characteristics used in this study are surf break location and quality. The GPS coordinates of each break are provided using the WGS84 datum. Accuracy ranges from 3m-200m depending on the original datum used to record the location. Each break is assigned one of five quality ratings, which are verified by independent experts and range from “sloppy” (1-star) to “totally epic” (5-star) (see Figure 2.1). Quality describes the average physical quality of the waves at the location, and in most cases is fixed over time (we study rare examples of changing qualities in Experiment II). It does not capture crowds, ease of access, surrounding views, etc., which are recorded separately by WannaSurf. In the first instance data on a new break is nominated by one of the website’s 78,000 registered users. The data is then checked and monitored for accuracy by one or

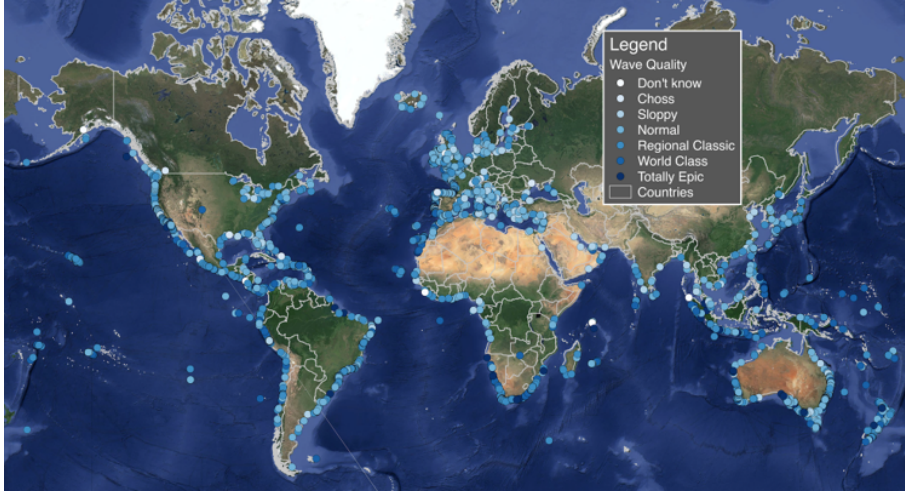


Figure 2.1: Overview of surf break locations in WannaSurf database.

Star Rating	Description	Frequency	Share
1	Sloppy	384	7.4%
2	Normal	2,041	39.4%
3	Regional Classic	2,141	41.1%
4	World Class	464	8.8%
5	Totally Epic	161	3.1%
Total		5,205	100%

Table 2.1: Breakdown of WannaSurf surf breaks by quality.

more WannaSurf Regional Correspondents. For reference, of the eleven locations on the surfing World Championship Tour seven are 5-star, three are in popular surfing areas with a dense concentration of 4-star breaks, one is 3-star (in central Rio de Janeiro, to promote the sport), and crowds range from “empty” to “ultra-crowded”.

Measurement error may be a problem if the quality ratings on WannaSurf are inaccurate. Surfers may bias ratings down to discourage crowding at their favourite breaks. This is not a concern for four reasons. First, we use quality ratings that are verified by independent WannaSurf experts, and ignore the poll-based “user ratings” (see <http://www.wannasurf.com/help/faq/index.html>). Second, WannaSurf removes legitimately secret locations. Third, a downward bias would overstate the growth potential of low-quality breaks, making our estimates more conservative. Fourth, while wave quality is crucial, so too is shared knowledge of it, which is directly measured by WannaSurf.⁷

The data may also suffer some selection bias. Ideally we would have data on wave quality

⁷Oceanologists are not yet been able to predict break quality using geographic and climactic models, due to the complex interactions that determine it at a metre-by-metre scale (see Appendix A; though they do forecast wave heights relatively successfully, for example www.swellnet.com.au). Rather than using geographic variables to instrument for break quality, this paper has the advantage of measuring quality directly.




Figure 2.2: WannaSurf data covers 146 countries, though data on surf breaks is concentrated in Australia and the US.

at regular intervals across the world’s ~620,000km of coastline. Instead our breaks are distributed across 146 countries, but are concentrated in Australia (888 breaks) and the US (878) (see Figure 2.2). Break quality varies across countries due either to exogenous variation or to selection.⁸ Low-quality breaks also tend to have more light than high-quality breaks, which we attribute to surfers entering accessible low-quality breaks into the database, but ignoring isolated ones. We discuss robustness to selection bias in Section 3.3.3 and Appendix D.3.

WannaSurf records a variety of other characteristics for each break. These include variables on accessibility (“Distance”, “Easy to find?”, “Public access?”, “Crowd”), difficulty (“Experience”), the type of wave (“Frequency”, “Type”, “Direction”, “Length”, “Bottom”, “Power”) and oceanic conditions (“Good swell direction”, “Good wind direction”, “Swell size”, “Best tide”). Of these we use the “Type” variable, which indicates whether the shoreline is a beach, a reef, a river-mouth, a headland (point-break) or a breakwater, to test whether other omitted geographic variables influence our results (see Figure 2.3 and Table D.2).

2.2 Night-time lights

The Defence Meteorological Satellite Program’s Operational Linescan System (DMSP-OLS) uses satellites to record the average annual night-time light intensity around the world, from 1992-2013 (Figure 2.4i.). The data is provided at a resolution of 30x30

⁸Namibia, Western Sahara and the Maldives have the highest average quality; Ukraine, Qatar and Kuwait the lowest.

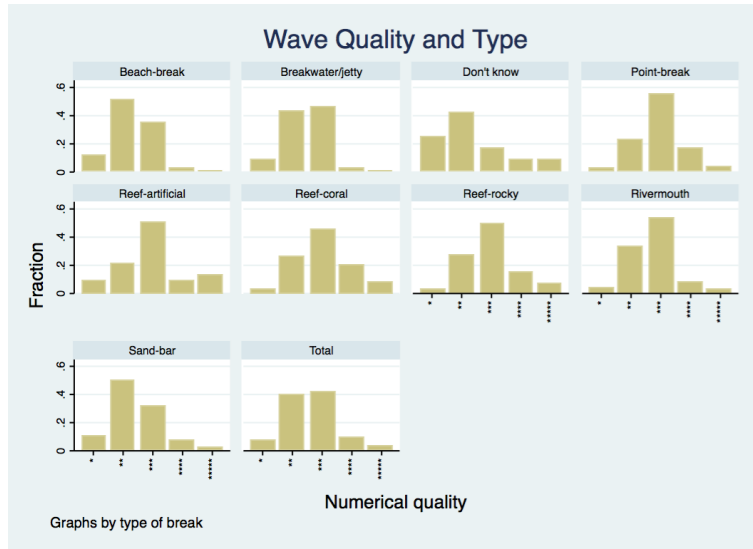


Figure 2.3: WannaSurf data on surf breaks is broken into 9 geographical types.

arcseconds (approximately 1 square kilometre near the equator), and ranges from 0 to 63. The data is constructed by overlaying all daily images over the course of a year, discarding those that are obfuscated by cloud cover, lightning, aurora, etc. for a given pixel.

There is a strong link at the national level between the growth of GDP and mean light intensity (Doll et al., 2006; Henderson et al., 2012; Michalopoulos and Papaioannou, 2014). This is illustrated in Figure 2.5, which plots the log of the sum of light readings by country against two measures of log PPP-adjusted GDP: based on expenditure and production (Penn World Tables 8.1). The associated regressions yield an adjusted r-squared of .82 and .80 respectively. We make use of the high spatial resolution of the data to study economic activity at a sub-national level, as has been done in a number of other studies (Chen and Nordhaus, 2011; Michalopoulos and Papaioannou, 2013; Hodler and Raschky, 2014; Jedwab et al., 2015; Jedwab and Moradi, 2016; Smith and Wills, 2016).

We use the night-time light data to measure real economic in two ways: illumination in the immediate vicinity of a surf break and illumination in nearby towns.

Illumination in the immediate vicinity of a break is measured using luminosity in concentric circles of various radii. We draw these circles at 5km, 5-10km and 10-50km around each break, clipped along the coastline (see Figure 2.6), and take the total illumination on the remaining land in each circle for each year. Coastlines are defined by a shape file from VDS Technologies, a private mapping firm, which may be subject to some small measurement error but it will be uncorrelated with break quality. Circles may overlap, so we include all overlapping areas in all breaks they are close to. This is the most conservative approach because high quality breaks will raise light growth around nearby low quality breaks, making the null hypothesis more difficult to reject. An alternative ap-

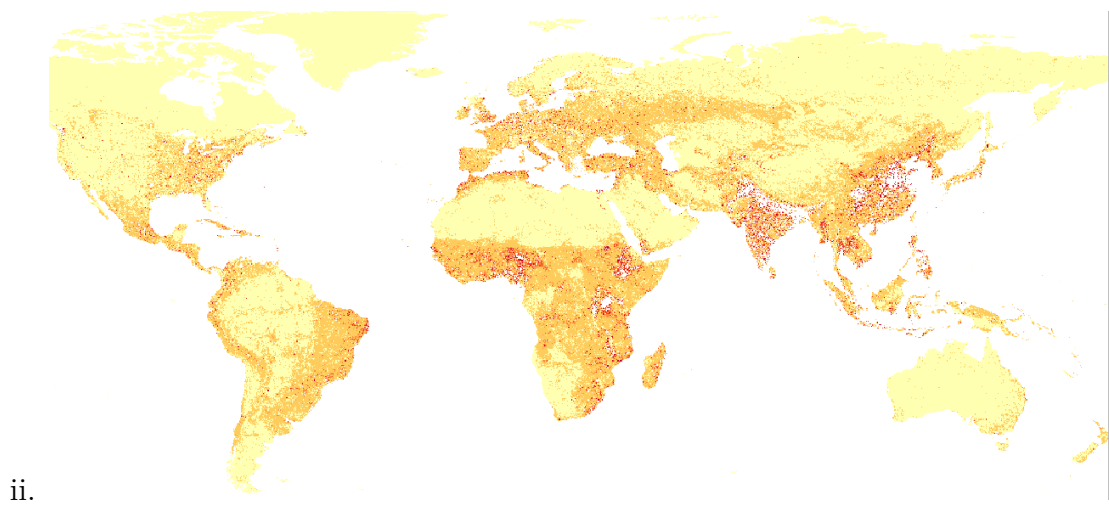



Figure 2.4: Data on i. night-time lights (DMSP-OLS) and ii. population (LandScan)




Figure 2.5: PPP-adjusted GDP vs Night-time lights (in logs), 2003 (see Smith and Wills, 2016).

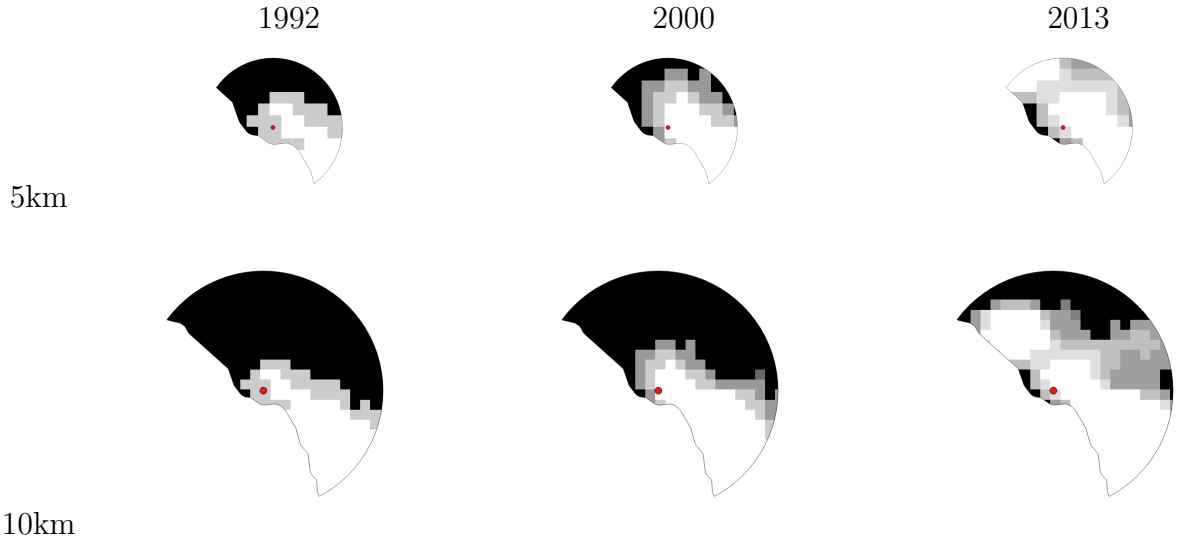


Figure 2.6: Example of illumination growth in the 5km and 10km surrounding Anchor Point, a “World Class” (4-star) break in southern Morocco, clipped along the coastline.

proach, based on treating individual pixels with the quality of nearby waves, is described in Appendix D.1.

Illumination in nearby towns is measured by endogenously locating towns by their population density, using data outlined below. A town is defined by a perimeter enclosing cells with a population density of 300 persons per square kilometre or more (see Figure C.2).⁹ Each break is linked to two towns: the closest town and the largest town within 50km radius (based on total population).

Night-time light data is subject to some issues that are relevant for this study. First, “top-coding” refers to pixels with a light reading of 63, beyond which we cannot distinguish levels of economic activity. This is not important in our data as it occurs for less than five “novelty” breaks in the centre of cities, like a wave park in Kuala Lumpur and a river-wave in central Munich. Second, light data includes significant luminosity readings from gas flares, which do not reflect economic activity. To control for this we drop all cells with gas flare activity according to the provider of the lights data (the NOAA Earth Observation Group), and trim any observations over water. Third, light data is affected by overglow (or “blooming”), where light is recorded in pixels away from its origin, and is magnified over terrain like water and snow (Doll, 2008). Small et al. (2005) find that overglow is linearly proportional to lit area, which is consistent with a physical model for atmospheric scattering. We clip our observations around the coastline, and in Appendix B confirm that overglow changes linearly with light, so will not bias our study of light growth. Fourth, the satellites used to construct the data change in 1994 (F11), 1995 (F13), 1997 (F14), 2000 (F15), 2003 (F16), 2007 (F17), and 2010 (F18) - which we average in

⁹We also use a cut-off of 600 persons per square km as a robustness check.

years with multiple satellites, and the effectiveness of the sensors diminishes over time. To control for this we use year fixed effects.

2.3 Population

Oak Ridge National Laboratory produces the LandScan dataset which provides annual mid-year spatial population counts at a 30x30 arcsecond resolution from 2000-2013 (Figure 2.4). It is similar to the Gridded Population of the World data from NASA’s Socioeconomic Data and Applications Center (SEDAC), which measures population at a 30 × 30 arc second resolution in 1990, 1995 and 2000, and has been used by Dell (2010) and Alesina et al. (2015) amongst others. LandScan records estimates of the “ambient”, or daily average, population in each grid cell. Importantly, this excludes intermittent population such as tourists or temporary relief workers, and may not reflect things like seasonal migrations or refugee movements. This allows us to focus explicitly on spillovers to the permanent population. The estimates are based on aggregate data for second order administrative units compiled by the International Programs Center of the US Bureau of Census. The aggregate data is distributed throughout the grid according to a likelihood model that uses inputs including data on elevation, land cover, roads, night lights, coastlines, settlements, and high resolution satellite imagery, among other sources.¹⁰ It pays special attention to coastal features. To account for the dynamics of coastal change the LandScan model extends all coastal boundaries several kilometres seaward. This ensures that all shore and small island features are included within an administrative unit boundary.

This population data also has some shortcomings. Aggregate population is allocated spatially using a likelihood model. This is subject to model error. The census data is also not collected every year. Between census years it is based on annual mid-year national population estimates from the Geographic Studies Branch of the US Bureau of Census, so year-to-year population comparisons suffer from interpolation error. LandScan cautions against using the data for annual cell-by-cell migration comparisons. Therefore, we present all population changes in 5km to 50km circles, over a 13 year horizon which incorporates at least two censuses in most cases.

2.4 Urban and rural classification

SEDAC also provides an “Urban Extents Grid”, from its Global Rural-Urban Mapping Project (GRUMP), which uses 1995 population estimates to classify each square of a

¹⁰For further detail http://web.ornl.gov/sci/landscan/landscan_documentation.shtml




Figure 2.7: The NOAA’s monthly Southern Oscillation Index.

30x30 arc second global grid as either urban or non-urban. The classification is based on contiguous lighted squares (as of 1995) and settlements known to hold at least 5000 people, and agrees with urban extents based on DHS surveys (Dorelien et al., 2013).

2.5 Political stability and ease of doing business

The World Bank provides data on political stability in its Worldwide Governance Index, and on the ease of doing business in its Doing Business Survey. We collect countries into four groups with similar numbers of breaks based on their 2014 scores on each. Table C.4 groups countries into bins by political stability, and table C.3 groups them by the quality of their business environment, with bin 1 being “very low” and 4 being “high”.

2.6 El Niño and Wave Heights

The National Oceanic and Atmospheric Administration records the monthly Southern Oscillation Index (SOI), based on the difference in the standardized sea level pressure between Tahiti and Darwin, Australia (see Figure 2.7). An El Niño event is defined as any where $SOI \leq -0.7$ for three or more consecutive months, which gives events in 1992, 1993, 1994, 1997, 1998, 2002, 2006, 2009, and 2010.

Australia’s Commonwealth Scientific and Industrial Research Organisation (CSIRO) provided data on the monthly mean significant wave height from its Centre for Australian Weather and Climate Research (CAWCR) Wave Hindcast on a global 24 arcminute grid, from

1992-2013 (Hemer et al., 2011). The mean significant wave height is defined as the average height (trough to crest) of the upper third of the waves in the wave-field. They also provided data on the monthly wave height anomaly, which is the deviation from the annual cycle at each grid point.

3 Experiment I: Do good surf breaks contribute more to economic growth than bad breaks?

“Surfers are the trendsetters, then the other tourists follow”

- Tarik Senhaji, Director General, Moroccan Sovereign Wealth Fund (2016)

3.1 Identification

Surfing offers a clean natural experiment because the quality of surf breaks is exogenously determined by a careful calibration of climatic, bathymetric and geographic conditions, which we exploit in the following three experiments.

Waves are created by wind acting on the surface of the ocean far out to sea. If the wind is strong and consistent enough the waves can travel up to 15,000km across the ocean, until they break against the shallow sea floor. The quality of a surfing wave depends on a complex and finely-calibrated interaction of time-variant factors, like the strength, direction and duration of winds where the wave originates and breaks; and time-invariant factors, like the sea floor bathymetry over the distance the wave travels, and the shape of the coastline where it breaks. These factors determine the quality of a surf break, which will be zero for the vast majority of the world’s coastline (for more detail see Appendix A).

Therefore, while a surf break’s existence may be related to economic activity (due to trade advantages of coastlines, etc.) the quality of the break is essentially random. Experiment I exploits this randomness to determine whether areas near high-quality waves grow faster than those near low-quality waves. The control group is areas near 1-star waves, which is a relatively high hurdle as they are also coastal and of sufficient interest to surfers to appear in the WannaSurf database. This identification strategy may still be subject to issues like endogeneity, selection bias and measurement error, which we consider in our robustness tests in Section 3.3.3.

An example may be helpful. In less than a mile along the north shore of Oahu, Hawaii, lie three of the world’s most famous, 5-star reef-breaks: Backdoor, Pipeline and Sunset




Figure 3.1: Satellite image of the North Shore of Oahu, Hawaii, showing how a single mile of beach can contain three 5-star breaks and one 2-star break, due to subtly detailed variation in bathymetry (Google Maps, 2016).

Beach (see Figure 3.1). Every December the world surfing tour concludes there, as champion surfers test themselves against long-range north Pacific swells that break against a succession of shallow reefs. In between Pipeline and Sunset lies Kammieland: another reef break on exactly the same beach, facing the same direction, and receiving the same swell, which is only rated 2-stars due to a complex interaction of water flowing over that reef and others. The rest of the breaks on this uninterrupted stretch of sand do not even warrant a name.

3.2 Estimating Equations

We exploit cross-sectional variation in the quality of surf breaks to estimate their effect on spatial outcomes using the following difference-in-difference model:

$$\ln(Y_{i,2013}^d) - \ln(Y_{i,1992}^d) = \alpha + \beta Q_i + F_i + Z_t + \epsilon_{i,t} \quad (3.1)$$

where $Y_{i,t}^d$ is the light/population within d kilometres of break i in year t , $Q_i \in [1, 5]$ is the star quality rating of break i which we treat as both continuous and an indicator, F_i is zone or country fixed effects, and Z_t is year fixed effects. The counterfactual is the change in $\ln(Y_{i,t}^d)$ from 1992-2013 for a 1-star break. This is a relatively high hurdle because 1-star breaks are on the coast and are sufficiently known by surfers to appear in WannaSurf, so our estimates are conservative. We deal with potential spatial correlation by clustering

the standard errors by country and zone,¹¹ and using spatially-robust heteroskedastic- and autocorrelation consistent (HAC) standard errors following Conley (1999 and 2010; implemented using Hsiang, 2010 and Fetzer, 2014).¹²

To visually represent the effects of surf break quality on growth we also use a polynomial time-trend model:

$$\ln(Y_{i,t}^d) = \alpha + \beta(t)Q_i + \gamma(t) + W_i + Z_t + \epsilon_{i,t} \quad (3.2)$$

where

$$\beta(t) = \beta_1 t + \beta_2 t^2 + \beta_3 t^3 + \beta_4 t^4$$

$$\gamma(t) = \gamma_1 t + \gamma_2 t^2 + \gamma_3 t^3 + \gamma_4 t^4$$

where $Y_{i,t}^d$ is as defined above for time $t = [1992, \dots, 2013]$, Q_i is an indicator equal to zero if the break is of poor quality (1-star) and one if the break is of some higher quality (2-5 stars), W_i is break fixed effects and Z_t is year fixed effects. The polynomial structure is imposed on $\beta(t)$ and $\gamma(t)$ to reduce the effects of collinearity in the data, though for the pixel-level analysis we replace $\beta(t)$ and $\gamma(t)$ with year dummies as we have more observations. We again report spatially-robust HAC standard errors following Conley (1999 and 2010).

3.3 Results

Economic activity near high quality breaks grew faster than near low quality breaks from 1992-2013. Activity (proxied by night-time lights) increased overall, rather than simply being reallocated from nearby areas, and grew particularly faster in existing towns and in emerging economies. The permanent population fell near good breaks, which is consistent with tourism driving up property prices, but this was accompanied by urbanisation. These results are robust to a variety of controls including omitted geographic characteristics, overlapping breaks, wave-rich countries, and selection bias.

3.3.1 Surfing and economic activity

Economic activity near high-quality breaks grew significantly faster than near low-quality breaks from 1992-2013. Table 3.1 (columns 1-3) shows that a 1-star increase in break

¹¹Zone is the largest subnational unit (e.g. the United States is comprised of eight zones).

¹²This applies a spatial weighting matrix to the standard errors that decays linearly out to 100km, and accounts for serial correlation at each location out to three lags non-parametrically.

Total difference in $\ln(\text{lights}^{5km})$ from 1992-2013						
	(1)	(2)	(3)	(4)	(5)	(6)
Quality (cts)	0.0358*** (0.00155)	0.0377*** (0.00369)	0.0358*** (0.00922)			
2 star				0.0196 (0.531)	0.0297 (0.388)	0.0196 (0.553)
3 star				0.0808** (0.0109)	0.101*** (0.00460)	0.0808** (0.0165)
4 star				0.0878** (0.0188)	0.0857** (0.0485)	0.0878* (0.0713)
5 star				0.0924 (0.306)	0.0948 (0.127)	0.0924 (0.186)
Observations	4,289	4,289	4,289	4,289	4,289	4,289
R-squared	0.389	0.522	0.003	0.390	0.523	0.003
Sample	All breaks	All breaks	All breaks	All breaks	All breaks	All breaks
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
FE	Country	Zone	Country	Country	Zone	Country
SE	Country	Zone	Conley	Country	Zone	Conley

Robust p-values in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 3.1: The effect of surf break quality on the change in $\ln(\text{lights})$ in the surrounding 5km, from 1992-2013. Quality is measured both continuously, (1)-(3), and using indicator variables, (4)-(6). Fixed effects are at country and zone level. Standard errors are clustered at the country and zone level, and allow for spatial correlation within 100km, and autocorrelation to 3 periods (Conley).

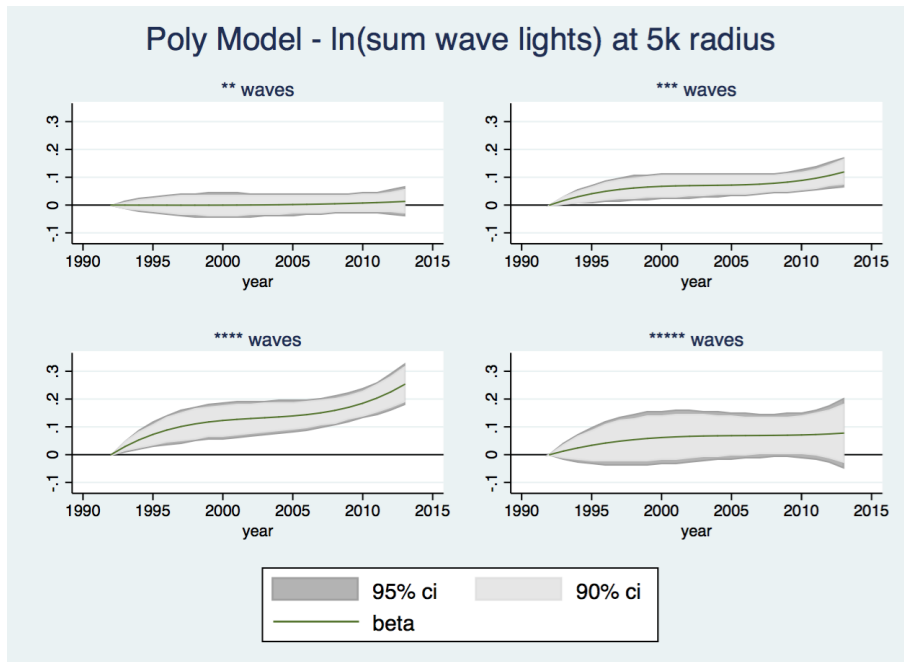


Figure 3.2: Polynomial model of the difference in difference of light levels from 1992-2013, using 1-star breaks as the baseline. Includes year and break fixed effects, and Conley standard errors.

quality increased lights in the surrounding 5km by 3.6-3.8 log points over our sample (~0.17 percentage points per annum). This was primarily driven by 3- and 4-star breaks, which grew by 8.6-8.8 log points more than 1-star breaks (~0.41-0.42 percentage points per annum; columns 4-6). 5-star breaks saw a large but imprecisely estimated increase. Higher economic activity near good breaks was not just redistributed from surrounding areas, with positive and diminishing spillovers found at least 50km away (see Tables C.1 and C.2).

Figure 3.2 supports these results by showing a consistently positive trend in the level of lights near 3- and 4-star breaks, relative to 1-star breaks, from 1992-2013 (equation 3.2). The effects in this figure are even larger because it applies break fixed effects to light levels, rather than country fixed effects to light growth; though we focus on the table to be conservative.

Together these results indicate that natural amenities can have a persistent effect on economic growth, because modern surfing has been popular since the 1960s. As with other natural amenities, people do not pay to enjoy surf breaks so additional growth must come through spillovers in the markets for physical capital and labour. Surf-breaks can augment physical capital by providing a basis for tourism and recreational industries. They also augment labour by attracting workers willing to accept lower wages, and entrepreneurs willing to accept lower profits, if they can regularly surf good waves (see review of the literature on natural amenities by Waltert and Shlapfer, 2010). In a standard Ramsey

growth framework this would encourage investment, as places with more natural capital have a higher steady-state level of output. Additional investment causes faster output growth, and the results above suggest that this happens on a timescale of decades.

Surfing also boosts economic growth in nearby towns (see Table 3.3). Defining towns by population density (see Figure C.2), and controlling for convergence by including the initial level of lights,¹³ we find that the closest town to 3-star and 4-star breaks grew by 11 and 13 log points more respectively from 1992-2013, relative to 1-star breaks (~ 0.54 and 0.63 percentage points faster per annum). The largest town within 50 km of 3- and 4-star breaks grew an additional 7.1 log points respectively over the period (~ 0.34 percentage points faster per annum).¹⁴ 4-star breaks had the largest effect, which is consistent with 5-star breaks requiring a lot of experience to surf (see Figure C.3). These results suggest that natural amenities particularly benefit areas of existing economic activity, so that path dependence matters for the location of economic growth.

Surfing's effect on economic growth is concentrated in developing countries (see Table 3.2). We proxy the supply of tourist services by the country's ease of doing business, and demand for tourism by its political stability (with and without Australia and the USA). Growth near breaks is strongest when the ease of doing business is moderately high. Such countries may have a sufficient business environment that allows tourism to meet demand from surfers (unlike areas with the poorest business environment), without being too mature. In contrast, growth is strongest when political stability is low or very low. Surfers therefore seem more willing to bear the costs of political instability than other tourists to enjoy good waves. As geography's effect on growth appears to be "turned on" at lower institutional levels, it provides further evidence that growth is determined by the interaction between geographic and institutional factors.

These results can be placed in perspective by converting them into approximate dollar values. The average value of a unit of light can be found by taking global GDP (US\$ 2013 PPP, World Bank WDI) and dividing by the global sum of lights in that year. Taking the areas within 5km and 10km of each break, and correcting for overlap, we find that 3-star breaks contributed on average an extra \$1.01 million per break per year to the surrounding 5km (2013 PPP) and \$2.45 million to the surrounding 10km, relative to 1-star breaks from 1992-2013 (based on the results in Table 3.1). Similarly, 4-star breaks contributed \$0.53 million and \$1.2 million respectively, which is less as these breaks are on average more rural. In total 3- and 4-star surf breaks add \$1.36 billion to their surrounding 5km globally per year, and \$4.00 billion to their surrounding 10km, relative to 1-star breaks.

¹³This uses annual growth, $\Delta \ln(Y_{i,t})$, rather than total growth, $\ln(Y_{i,2013}) - \ln(Y_{i,1992})$, to avoid having $\ln(Y_{i,1992})$ on both sides of the regression for every observation.

¹⁴The results are robust to replacing the initial level of town lights with country fixed effects for the largest nearby town but not the closest. The results are no longer significant if both the initial level of town lights and country fixed effects are included.

Total difference in $\ln(\text{lights}^{5km})$ from 1992-2013				
	Doing Business Index (DBI)		Political Stability Index (PSI)	
	All countries	Excl. Aus/USA	All countries	Excl. Aus/USA
Quality (continuous) interacted with index bin:				
Very Low	-0.0741** (0.0378)	-0.0741** (0.0394)	0.0786** (0.0398)	0.0786** (0.0415)
Low	0.0553 (0.173)	0.0553 (0.177)	0.0715** (0.0106)	0.0715** (0.0113)
Medium	0.0805** (0.0106)	0.0805** (0.0113)	0.0370*** (0.00434)	0.0300 (0.181)
High	0.0293** (0.0177)	0.0367** (0.0416)	0.00929 (0.464)	0.0277 (0.411)
Constant	0.446*** (0)	0.551*** (0)	0.442*** (0)	0.545*** (0)
Observations	4,052	2,554	4,105	2,607
R-squared	0.392	0.404	0.392	0.404
Year FE	Yes	Yes	Yes	Yes
FE	Country	Country	Country	Country
SE	Country	Country	Country	Country

Robust p-values in parentheses

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

Table 3.2: The effect of break quality on the change in $\ln(\text{lights})$ within 5km of each break, from 1992-2013, broken down into bins based on the World Bank's Doing Business and Political Stability indices, with and without Australia/USA.

Annual difference in $\ln(\text{lights}^{\text{town}})$ from 1992-2013								
VARIABLES	Closest town				Largest town in 50km			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Quality (cts)	0.00165** (0.0401)	0.00165* (0.0670)			0.00106** (0.0176)	0.00106** (0.0181)		
2 star			0.00222 (0.320)	0.00222 (0.316)			0.00174 (0.126)	0.00174* (0.0974)
3 star			0.00537*** (0.00780)	0.00537*** (0.00754)			0.00341*** (0.00367)	0.00341*** (0.00170)
4 star			0.00626* (0.0559)	0.00626** (0.0374)			0.00341* (0.0791)	0.00341** (0.0367)
5 star			0.000410 (0.899)	0.000410 (0.920)			0.00300 (0.145)	0.00300 (0.180)
$\ln(\text{lights}^{\text{town}})$ in 1992	-0.00402*** (9.82e-07)	-0.00402*** (1.53e-10)	-0.00401*** (1.11e-06)	-0.00401*** (1.64e-10)	-0.00550*** (0)	-0.00550*** (0)	-0.00550*** (0)	-0.00550*** (0)
Constant	0.163*** (3.95e-08)	0.163*** (1.77e-10)	0.164*** (2.46e-08)	0.164*** (1.18e-10)	0.0559*** (8.51e-06)	0.0559*** (5.17e-11)	0.0561*** (4.43e-06)	0.0561*** (0)
Observations	68,279	68,279	68,279	68,279	81,455	81,455	81,455	81,455
R-squared	0.598	0.598	0.598	0.598	0.858	0.858	0.858	0.858
Sample	All breaks	All breaks	All breaks	All breaks	All breaks	All breaks	All breaks	All breaks
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	No	No	No	No	No	No	No	No
SE	Country	Zone	Country	Zone	Country	Zone	Country	Zone

Robust p-values in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Table 3.3: The effect of break quality on the annual change in $\ln(\text{lights})$ at the closest town (1-4) and largest town within 50km (5-8), from 1992-2013. Quality is measured both continuously, (1,2,5,6), and using indicator variables, (3,4,7,8). Regression includes the initial level of town lights in 1992 to control for convergence. Standard errors are clustered at the country and zone level.

Total difference in $\ln(\text{population})$ from 1992-2013												
	5km				5 to 10km				10 to 50km			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Quality (cts)	-0.0272 (0.255)	-0.0235 (0.396)			-0.0636*** (0.000216)	-0.0551** (0.0202)			-0.0348*** (0.00371)	-0.0255*** (0.00580)		
2 star			-0.107* (0.0782)	-0.112 (0.122)			-0.0225 (0.565)	-0.0141 (0.782)			0.0284 (0.355)	0.0270 (0.436)
3 star			-0.137** (0.0266)	-0.125** (0.0466)			-0.0751** (0.0422)	-0.0697 (0.215)			0.0293 (0.416)	0.0246 (0.461)
4 star			-0.172* (0.0866)	-0.169* (0.0900)			-0.171** (0.0145)	-0.122 (0.158)			-0.106** (0.0299)	-0.0622* (0.0514)
5 star			-0.0446 (0.766)	-0.0565 (0.674)			-0.267** (0.0485)	-0.243* (0.0667)			-0.146** (0.0152)	-0.118* (0.0984)
Constant	0.656*** (0)	0.646*** (0)	0.700*** (0)	0.697*** (0)	0.631*** (0)	0.609*** (0)	0.529*** (0)	0.518*** (0)	0.413*** (0)	0.388*** (0)	0.313*** (0)	0.310*** (0)
Obs.	5,054	5,054	5,054	5,054	5,083	5,083	5,083	5,083	5,155	5,155	5,155	5,155
R-sq.	0.296	0.376	0.296	0.376	0.293	0.394	0.293	0.394	0.322	0.469	0.324	0.471
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FE	Country	Zone	Country	Zone	Country	Zone	Country	Zone	Country	Zone	Country	Zone
SE	Country	Zone	Country	Zone	Country	Zone	Country	Zone	Country	Zone	Country	Zone

Robust p-values in parentheses

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

Table 3.4: Total change in permanent population within 5km, 5 to 10km and 10 to 50km of surf breaks from 1992-2013. Quality is measured both continuously, (1,2,5,6,9,10), and using indicator variables, (3,4,7,8,11,12). Fixed effects and standard error clusters are at the country and zone level.

Total difference in $\ln(\text{population})$ from 1992-2013												
	Closest town				Largest town in 50km				Rural areas in 50km			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Quality (cts)	0.0909*	0.104**			0.0421**	0.0432**			-0.0372***	-0.0358***		
	(0.0512)	(0.0291)			(0.0164)	(0.0467)			(0.00280)	(0.00164)		
2 star			0.130	0.0504			0.0775	0.0535			0.0138	-0.00130
			(0.247)	(0.744)			(0.167)	(0.439)			(0.411)	(0.968)
3 star			0.233*	0.217			0.114**	0.129**			-0.0178	-0.0380
			(0.0747)	(0.136)			(0.0161)	(0.0253)			(0.343)	(0.301)
4 star			0.374**	0.314			0.104	0.0718			-0.0796**	-0.0821*
			(0.0475)	(0.126)			(0.181)	(0.196)			(0.0481)	(0.0519)
5 star			0.205	0.208			0.252	0.201**			-0.153**	-0.143**
			(0.328)	(0.473)			(0.139)	(0.0235)			(0.0475)	(0.0417)
Constant	1.197***	1.163***	1.246***	1.288***	0.356***	0.353***	0.371***	0.378***	0.466***	0.463***	0.383***	0.397***
	(0)	(0)	(0)	(0)	(0)	(1.55e-09)	(0)	(0)	(0)	(0)	(0)	(0)
Observations	3,663	3,663	3,663	3,663	4,005	4,005	4,005	4,005	5,179	5,179	5,179	5,179
R-squared	0.245	0.350	0.245	0.351	0.314	0.532	0.315	0.533	0.522	0.679	0.523	0.679
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FE	Country	Zone	Country	Zone	Country	Zone	Country	Zone	Country	Zone	Country	Zone
SE	Country	Zone	Country	Zone	Country	Zone	Country	Zone	Country	Zone	Country	Zone

Robust p-values in parentheses

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

Table 3.5: Total change in population in the closest town, largest town within 50km, and rural areas within 50km, from 1992-2013. Quality is measured both continuously (1,2,5,6,9,10) and using indicator variables (3,4,7,8,11,12). Fixed effects and standard error clusters are at the country and zone level.

3.3.2 Surfing and population

The permanent population near high-quality surf breaks fell relative to low-quality breaks from 1992-2013, up to 50km away (Table 3.4). The effect was strongest out to 10km from 3- and 4-star breaks, which also saw the largest relative increase in economic activity. As the LandScan data explicitly excludes tourists, this result is consistent with tourists moving towards areas with good surf breaks, raising property prices and driving away locals.

The aggregate fall in the permanent population appears to mask a reallocation of people from rural to urban areas near high-quality breaks. The population in the closest town, and the largest town within 50km of a high-quality break, rose from 1992-2013 (Table 3.5). This was particularly true for the 3- and 4-star breaks. At the same time the rural population within 50km of a high-quality break fell. This suggests that the increase in economic activity near high-quality surfing breaks was accompanied by urbanisation.

3.3.3 Robustness

There is a possibility that omitted geographic characteristics both increase the quality of surfing breaks and drive economic growth. We use quality as the treatment because this will be less correlated with economic activity than the existence of a wave. However, there may still be omitted characteristics like openings in reefs and river-mouths, which are favourable for surfing and may also facilitate fishing and trade through boating lanes. To address this we re-run the analysis excluding the breaks that are of disproportionately high quality: reefs (30% of sample), point-breaks (or headlands, 13%) and river mouths (2%) (see Figure 2.3 and Table D.2). Table D.3 shows that 3- and 4-star breaks see similarly-sized effects to the main specification, though the estimate for 4-star breaks is no longer significant due to the smaller sample size.

Circles around surf breaks may overlap. As noted in Section 2.2 we address this by conservatively including overlapping areas in all nearby breaks making it more difficult to reject the null (as good breaks would increase lights at nearby bad breaks). In addition, we re-run our analysis using pixels, rather than breaks, as the unit of observation (see Appendix D.1). This allows us to treat each pixel individually with the maximum quality break within 5km, and we find similar results to the baseline specification. When treating each pixel with a count of breaks of each quality within 5km, we find that more low-quality breaks reduce light growth, which disappears for higher-quality breaks.

The sample is dominated by Australia and the USA, which are large, developed countries and account for 17% of breaks each (see Figure 2.2). This is a problem if idiosyncratic

trends in Australia and the USA bias the overall results, or if lights are a poor proxy for growth in large developed countries. Excluding Australia and the USA from the sample confirms, and in some cases strengthens, the results in the original specification (see Table D.3).

Finally, the WannaSurf database exhibits some selection bias, where the lowest quality breaks are more likely to appear in the data when they are in urban areas (see Figure D.3). This may bias our results down if urbanisation drives faster growth in lit areas, or up if lights in rural areas converge to those in towns. We control for this in three ways. First, we re-run the analysis using 2-star breaks as the baseline, which suffer less from selection because they have a similar urban share (49%) as the 4-star breaks that drive our main results (48%). We find that the areas near 3- and 4-star breaks also grow faster than those near 2-star breaks (see Table D.4). Second, Table 3.3 focuses only on towns to avoid any selection bias in rural areas, and controls for convergence using the initial level of lights. The results confirm that lights grow faster in towns near high-quality breaks. Third, we also conduct the pixel level analysis controlling for the initial level of lights in Table D.5. The results again suggest that 3- and 4-star breaks have a positive and significant effect on light growth, though the effect for 5-star breaks becomes negative.

4 Experiment II: Does discovering a new surf break alter economic growth?

“Four years ago, American surfer Brian Gable submitted Skeleton Bay as an entry to Surfing Magazine’s ‘Google Earth Challenge’ competition, which he subsequently won. Shortly thereafter Cory Lopez was filmed in a ridiculously long barrel...Since then it’s been documented by an increasing number of surfers to the point where each swell sees the beach lined with 4WDs and cameramen who’ve travelled from South Africa or even further afar.” - Brokensha (2012)

4.1 Identification

Experiment I compared how quickly lights grew near high-quality surf breaks, relative to low-quality breaks. Experiment II uses a different counterfactual: areas with no surf breaks at all. To do this we use a series of event studies to understand how economic activity is affected by discovering a new surf break, seeing an existing one disappear, or gaining the ability to surf a new break because of technological innovation.

Break	Country	Discov/Disapp	Date	Quality	Source
La Jolla	Mexico	Discovery	2006	4 star	Rip Curl Search
El Gringo	Chile	Discovery	2007	4 star	Rip Curl Search
Skeleton Bay	Namibia	Discovery	2008	5 star	Surfer Magazine
Jardim do Mar	Madeira, Portugal	Disappearance	2005	4 star	www.savethewaves.org
Mundaka	Spain	Disappearance	2005	5 star	www.savethewaves.org

Table 4.1: Discoveries and disappearances of surf breaks.

The first approach in Experiment II uses a sample of three discovery and two disappearance events, drawn from a survey of the surfing literature (see Table 4.1).¹⁵ We require these events to have four characteristics. First, they must involve high-quality surf breaks. Second, discovered breaks must become known, and disappearances must have been known, to the global surfing community via mainstream media. Third, this must happen at a clearly defined time during our sample period, 1992-2013. Fourth, the event must be exogenous to local economic activity.

The first discovery comes from Surfer magazine’s 2008 “Google Earth Challenge”. Contestants were asked to pore over Google Earth and identify a break that had never been surfed before. The winner was Skeleton Bay, a 1500m long point break off the coast of the Namibian desert. Since the break was found the surfing world - amateur and professional alike - has descended into the desert to ride the relentlessly grinding waves. We study whether economic growth followed.

The other two discoveries come from the first times the surfing World Championship Tour was held at a “secret spot”: in 2006 at La Jolla, Oaxaca, Mexico;¹⁶ and in 2007 at El Gringo, Arica, Chile.¹⁷ Conversely, breaks disappeared in 2005 at Jardim do Mar, Madeira, Portugal, when a rock wall was built to protect a new coastal road; and in Mundaka, Biscay, Spain, when a rivermouth was dredged for boats leading to the cancellation of a 2005 world tour event. Both disappearances were by-products of local investment, making any estimates of economic decline very conservative.

The second approach in Experiment II uses Rip Curl’s release of the first battery-heated surfing wetsuit in 2007, which made surfing in extremely cold waters much more accessible

¹⁵The majority of breaks were discovered during the 1960s and 70s, when surfing first became a global phenomenon. Unfortunately our data does not go back that far and there is no official surfing body or archive that stores and maintains this kind of information.

¹⁶The effect is captured by this quote, “...the ASP [Association of Surfing Professionals world tour] arrived [at La Jolla] in full regalia and scored it at its peak, and the best surfers in the world rode the best waves they’ve ever seen. And now, the blitz is on. The world watched everything unfold as live footage, beamed via satellite, flooded the web. The photos are everywhere. And even though the location remains, for now, a secret, every surfer in his right mind is frothing to find this quietly reeling point.” (Surfer Magazine, 2006)

¹⁷Rip Curl hosted “Search” events on the world tour from 2005-2011, but the locations were only secret in 2006 and 2007.

(Longman, 2016). We study whether this affected economic activity near cold-water surf breaks by treating 2007 as a discovery event for the 83 locations in our database at latitudes above 55 degrees north (none are below 55 degrees south; see Table E.2).

4.2 Estimating Equations

For both approaches we study the difference between illumination near the breaks in question and the global mean using the model:

$$\ln(Y_{i,t}^{50}) - \overline{\ln(Y_t^{50})} = \alpha_i + \beta_i t + \alpha_i^D I_i^D + \beta^D I_i^D D_i t + \epsilon_{i,t} \quad (4.1)$$

where $\ln(Y_{i,t}^{50})$ is log illumination in the 50km surrounding break i in year t ; and $\overline{\ln(Y_t^{50})}$ is the mean of $\ln(Y_{i,t}^{50})$ over all 5000+ breaks, capturing global trends and changes in satellite sensitivity. We focus on 50km rather than 5km circles because of the relative isolation of the breaks of interest, and the small number of observations. On the right-hand side t is a linear time trend, I_i^D is an indicator equal to 1 after the discovery/disappearance event and 0 otherwise, D_i is equal to +1 for discoveries and -1 for disappearances, and the standard errors are clustered at the break level. This allows each break to have an individual intercept and time trend before the event, and an individual intercept after the event. It then jointly estimates the average marginal effect on the time trend of the break being discovered or disappearing, β^D .

4.3 Results - Discoveries and Disappearances

Discovering a new 4- or 5-star break increased (and a disappearance decreased) annual light growth in the surrounding 50km by an average of 2.2 percentage points, relative to the global average (see Table 4.2). For robustness we draw alternative event dates from a uniform $U(1992, 2013)$ distribution and find no effect (column 2). To illustrate the effects for each individual break we can compare the growth in de-trended lights, $(\ln(Y_{i,t}^{50}) - \overline{\ln(Y_t^{50})})$, before and after each event. Figure 4.1 shows that de-trended light growth more than doubled after each discovery, increasing by 4.7 percentage points at Skeleton Bay, 1.8 at La Jolla and 2.1 at El Gringo (though the effects are imprecisely estimated as each break has seven or less post-discovery observations). Figure 4.2 shows that both disappearances caused light growth to slow by ~2 percentage points (significant at the 5% level). This is particularly notable because in both cases the road-building and river-dredging that caused the disappearance was intended to raise economic growth. This shows that depleting the quality of natural amenities can be an economically important side-effect of infrastructure investment.

	Discov. and Disapp.		Cold Water Breaks			
	(1)	(2)	(3)	(4)	(5)	(6)
	Main	Robust	Main	Main	Robust	Robust
Year trend: post-event	0.0222*** (0.003)	-0.00495 (0.655)	0.0267*** 0.0000	0.0267** 0.0328	0.0032 0.8260	0.0032 0.8897
Individual breaks interacted with:						
Year trend: pre-event	Yes	Yes	Yes	Yes	Yes	Yes
Constant: post-event	Yes	Yes	Yes	Yes	Yes	Yes
Constant: pre-event	Yes	Yes	Yes	Yes	Yes	Yes
Observations	110	110	1,804	1,804	1,804	1,804
R-squared	0.996	0.995	0.032	0.032	0.038	0.038
No. of Breaks	5	5	83	83	83	83
Sample	Discov/Disapp		Lat \notin [-55, 55]			
SE Cluster	Break	Break	Break	Zone	Break	Zone

Robust p-values in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Table 4.2: Difference-in-difference in average log lights around (1) break discoveries and disappearances using actual event dates (for details see Table E.1); (2) break discoveries and disappearances using dates drawn from a uniform U(1992,2013) distribution; (3)-(4) cold water breaks after battery-heated wetsuits were invented in 2007; and (5)-(6) cold water breaks using 1997 as the invention date.

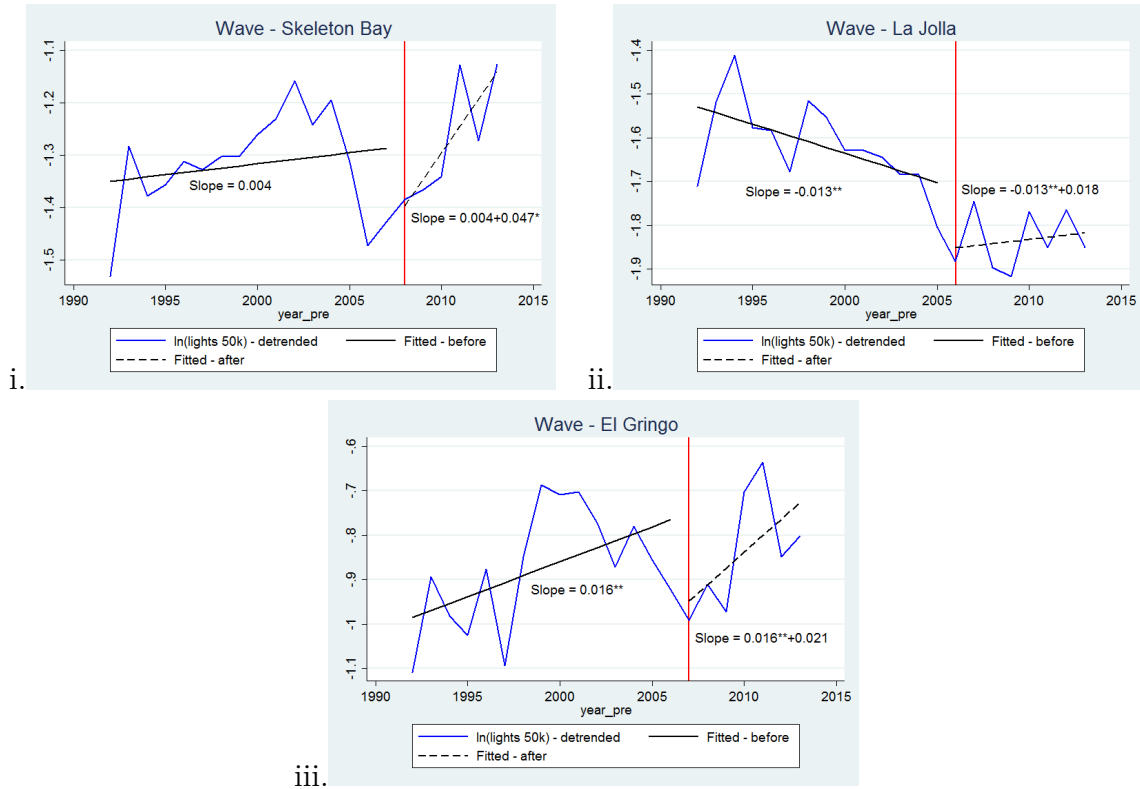


Figure 4.1: Effect of break discoveries on de-trended average log lights in the surrounding 50km for i. Skeleton Bay, Namibia, ii. La Jolla, Mexico, and El Gringo, Chile.




Figure 4.2: Effect of surf break disappearances on de-trended average log lights in the surrounding 50km for i. Jardim do Mar, Portugal, ii. Mundaka, Spain.

4.4 Results - Cold-Water Breaks

The invention of battery-heated wetsuits in 2007 increased annual light growth in the 50km surrounding cold-water surf breaks by 2.67 percentage points, relative to the global average (see Table 4.2). This result is significant at more than the 1% level. For robustness we test if the invention had hypothetically occurred a decade earlier, in 1997, and find no effect. In Table E.3 we re-run the analysis in equation 4.1 but estimate only one $\beta_i = \beta$, $\alpha_i^D = \alpha^D$, and $\beta_i^D = \beta^D$ for all i . This shows that prior to the wetsuit invention, light near cold water breaks was growing 1.6 percentage points p.a. slower than the global average. After the invention it grew by 1.25 percentage points faster.

Experiment II shows that being able to exploit a new, high-quality natural amenity for the first time can raise growth by over 2 percentage points per annum in the short term. This can be reconciled to the 0.4 percentage point p.a. increment in growth from Experiment I, which focused on longer term effects and used 1-star waves as the counterfactual (rather than no wave at all).

5 Experiment III: What happens when the surf is good?

“Getting ready for swells is one thing, but getting ready for a season of swells, like El Niño, is a whole different thing. It’s exciting to think that we might possibly score this year.”

-Timmy Reyes, professional surfer (Surfer Magazine, 3 November 2015)

“Timmy Reyes’ girlfriend even went as far as to say she hates El Niño after Timmy spent five months travelling to four countries on three continents, racking up more than 20,000 miles in an airplane and 10,000 miles in a car.”
- (Surfing Magazine, 2 June 2016)

5.1 Identification

Experiment III studies whether economic growth near high-quality surf breaks is particularly strong in years with good waves. To do this we exploit unanticipated changes in the El Niño weather pattern which generate large, well-ordered (long period), long-range swells. In contrast to the previous experiments, this is a panel study that interacts temporal variation in wave heights (size) and El Niño patterns (period) with cross-sectional variation in break quality to isolate the effects of surf breaks on economic activity.

The El Niño Southern Oscillation (ENSO) fluctuates between El Niño and La Niña states every three to seven years (Butt, 2009).¹⁸ In a typical year trade winds blow across the Central Pacific from east to west, pushing warm surface water towards Australasia. Here warm air rises, creating low air pressure and precipitation. The surface water is replaced with cool upwellings near Chile and Peru, which creates high pressure and dry conditions there. During a La Niña episode these trade winds are stronger than average, reinforcing the effects. During El Niño events the trade winds weaken, if not reverse. Warm surface water stays in the east, causing low-pressure systems in the far North Pacific to track much further south. Broadly, the result is larger waves in the North and Eastern Pacific (including California and Hawaii); and smaller waves in Australia and the North-West Atlantic (Butt, 2009; Housman, 2016). The advantage of El Niño swells, particularly for surfers in the North-East Pacific, is that they are generated far away and are often coupled with clement local conditions (Butt, 2009; Housman, 2016).

We use a triple-difference approach to estimate the effect of large waves, during El Niño years, at high-quality breaks. The interaction of all three variables is needed. Large waves alone are not sufficient for good surfing conditions, as they will be unruly if generated nearby and accompanied by strong winds and storms. El Niño events produce long-range swells and good weather in some parts of the world, but reduce wave heights in others (see Figure 5.1ii.). They also affect the economy through other channels, like weather and climate (Cashin et al., 2015). Finally, neither large waves nor good weather will matter if the break itself is low quality and cannot accommodate large swells.

¹⁸Exogenous ENSO changes have been used as natural experiments in a range of economic studies on a variety of countries, sectors and commodities (e.g. Handler and Handler 1983; Brunner, 2002; Hsiang et al., 2011; Ubilava, 2012; Iizumi et al., 2014; Hsiang and Meng, 2015). Cashin et al. (2015) find El Niño events cause economic activity to fall in Australia, Chile, Indonesia, India, Japan, New Zealand and South Africa, while the US and Europe benefit.

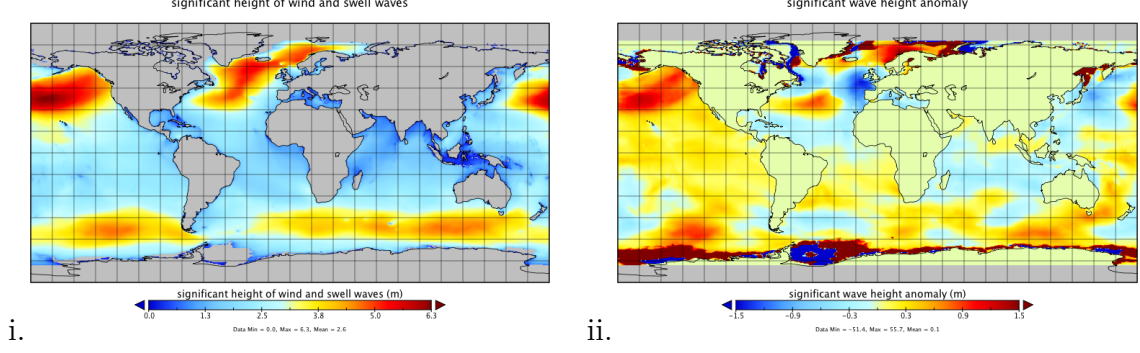


Figure 5.1: i. Significant wave heights (m) and ii. Significant wave height anomalies (m) during the January 1992 El Niño event.

5.2 Estimating Equations

We conduct this experiment using a triple-difference specification:

$$\begin{aligned}
\Delta \ln(Y_{i,t}^d) &= \alpha + \beta wha_{i,t} + \sum_{j=2}^5 \beta_j wha_{i,t} Q_j \\
&+ \gamma ENyr_t + \sum_{j=2}^5 \gamma_j ENyr_t Q_j \\
&+ \delta wha_{i,t} ENyr_t + \sum_{j=2}^5 \delta_j wha_{i,t} ENyr_t Q_j \\
&+ W_i + Z_t + \epsilon_{it}
\end{aligned} \tag{5.1}$$

where $\Delta \ln(Y_{i,t}^d)$ is the growth in log lights in the d km surrounding break i in year t , $wha_{i,t}$ is the mean wave height anomaly (a measure of wave height in each 24-arcminute grid relative to its average) at break i in year t , standardised to have a mean of 0 and a standard deviation of 1, $Q_j \in [1, 5]$ is the quality of break i , $ENyr_t$ is an indicator taking the value of 1 if year t was an El Niño year and 0 otherwise, W_i is break fixed effects and Z_t is year fixed effects. Standard errors are clustered at the zone level to account for spatial correlation.

The coefficients of interest are β_j and δ_j , which estimate the marginal effect of larger waves, and of larger waves during El Niño years, on breaks of different qualities. We first run the regression excluding any terms involving $ENyr_t$. This estimates the effect of wave heights on light growth, regardless of whether El Niño is responsible. We then run the full specification, to estimate the marginal impact of larger waves during El Niño events.

5.3 Results and discussion

We find that larger waves tend to have a negligible, or negative, impact on economic growth on average, but the effect becomes positive for good-quality breaks during El Niño years.

First we run the regression in equation 5.1 excluding any terms involving $ENyr_t$, to find the average effect of wave heights on light growth (columns (1) and (2) of Table 5.1). We find that, on average, unusually large waves had no effect on light growth for most break qualities, and a negative effect for “Normal” (2 star) breaks. This result may be attributed to a “stormy seas” effect: years with larger than average waves may also experience worse than average weather. As noted above, for large waves to be attractive to surfers, they must be accompanied by calm local conditions.

To isolate the marginal effect of waves generated by El Niño events we run the full specification in equation 5.1 at radii of 5km and 50km around each break (see columns (3) and (4) of Table 5.1). We find that light growth was considerably higher during El Niño years for all locations. This is consistent with better weather and improved agricultural yields in North America and Europe, as found by Cashin et al. (2015). However, satellite sensitivity also significantly improved in 1994, 1997 and 2010 when new satellites were launched, which also happen to be El Niño years. Therefore these improvements in sensitivity may incorrectly have been attributed to $ENyr_t$ rather than year fixed effects, overestimating the true effect. However, this is not our focus.

Our focus is the interaction of larger waves, during El Niño years, at high quality breaks. A one standard deviation increase in the wave height anomaly during El Niño years led to a marginal increase in light growth of 5.6 and 3.9 percentage points for 4- and 5-star breaks respectively. This suggests that the average stormy seas effect is offset during El Niño years, especially for good quality surf breaks.

These results are robust to excluding the US, which is a major beneficiary of El Niño swells. They are also robust to replacing the $ENyr_t$ indicator with the continuous Southern Oscillation Index (SOI), which also includes La Niña events that produce larger waves in the western Pacific. When repeating the analysis including an additional two lags for $ENyr_t$ and $wha_{i,t}$ we find no persistence in the $Quality * ENyr * wha$ interaction. This suggests El Niño events’ effect on light growth is constrained to that particular year, so picks up direct effects of surfers travelling to ride waves.¹⁹ Any persistent effects may be picked up in the long-term growth rates in Experiment I.

¹⁹Night-time lights can pick up short-term changes in activity through higher building occupancy, car headlights, etc.

Large El Niño Waves				
VARIABLES	(1)	(2)	(3)	(4)
	$D. \ln(\text{lights}^5)$	$D. \ln(\text{lights}^{50})$	$D. \ln(\text{lights}^5)$	$D. \ln(\text{lights}^{50})$
Wave Height Anomaly (cts) interacted with:				
Constant	-0.001 (0.254)	-0.001 (0.306)	-0.000 (0.706)	-0.001 (0.533)
2 star	-0.011** (0.018)	-0.007* (0.095)	-0.011* (0.066)	-0.008 (0.120)
3 star	-0.001 (0.490)	-0.001 (0.567)	-0.001 (0.432)	-0.003 (0.210)
4 star	0.001 (0.942)	0.006 (0.361)	-0.004 (0.622)	-0.014** (0.035)
5 star	0.006 (0.606)	-0.002 (0.740)	0.016 (0.373)	-0.014** (0.031)
El Nino Year (indicator) interacted with:				
Constant			0.177*** (0.000)	0.077*** (0.000)
2 star			0.012* (0.100)	-0.003 (0.500)
3 star			0.010 (0.140)	0.001 (0.840)
4 star			0.020 (0.102)	0.004 (0.649)
5 star			0.024* (0.102)	0.010 (0.649)
El Nino Year interacted with Wave Height Anomaly interacted with:				
Constant			-0.024*** (0.006)	-0.010 (0.138)
2 star			0.021** (0.027)	0.011* (0.082)
3 star			0.019*** (0.007)	0.014*** (0.004)
4 star			0.030 (0.143)	0.056*** (0.001)
5 star			-0.007 (0.835)	0.039** (0.023)
Observations	97,547	106,271	97,547	106,271
R-squared	0.473	0.758	0.473	0.758
Year FE	Yes	Yes	Yes	Yes
Break FE	Yes	Yes	Yes	Yes
Cluster	Zone	Zone	Zone	Zone

Robust p-values in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 5.1: The effect of a 1 standard deviation increase in the wave height anomaly on lights in the surrounding 5 and 50km, during El Niño years.

6 Conclusion

This paper estimates the impact of natural amenities on the location and pace of economic growth, by exploiting exogenous variation in the quality of surf breaks. To do this we combine four high-resolution spatial datasets, on the quality and location of 5000+ surf breaks, wave heights, night-time light emissions and population, to conduct three natural experiments.

These experiments find that high quality surf breaks significantly raise economic growth in the surrounding area, relative to low quality breaks, over both the short and the long run. The first experiment exploits cross-sectional variation and finds that the effect is concentrated in the 5km surrounding breaks, with spillovers up to 50km away. Surf breaks have a particularly large effect on nearby towns and in emerging economies; and tend to reduce the nearby permanent population in a way that is consistent with tourism. The second experiment exploits temporal variation and finds that discovering a high-quality break, or technology that makes cold-water breaks more accessible, increases growth in the surrounding areas. Conversely, destroying a break reduces growth, even if it is replaced by a new road or a dredged river. The third experiment uses a panel approach that exploits both cross-sectional and temporal variation, and finds that the area around good quality breaks grows particularly quickly when they have large waves during El Niño years.

Collectively these results show that natural amenities play an important role in economic growth. As noted in the introduction, there is extensive evidence that geography is important for growth. However, most of these studies focus on natural capital that directly affects the costs of production, like access to waterways, fertile soil and mineral resources. In contrast this paper studies natural amenities, which indirectly affect production by augmenting physical capital and labour. Existing work has found inconclusive evidence that natural amenities are important for growth, due to difficulties with identifying and measuring their effect. This paper fills that gap by using three unique natural experiments and a novel dataset to estimate how one particular natural amenity affects local growth.

The paper also has implications for policy. The first is that policymakers can use natural amenities like surf breaks as engines for growth over a range of time horizons, especially in developing countries. To do this they can promote the public and private investment needed to enjoy these amenities while protecting their environmental quality. This could include investment in hotel capacity, as Morocco's sovereign wealth fund is doing in Taghazout; shark-detection programs, as the New South Wales government is doing in Australia; or artificial inland surf breaks as in Wales, California and Dubai (see for example www.kswaveco.com). The estimates in this paper may be useful for comparing the costs and benefits of such projects.

The second is that policymakers must take into account the effects on natural amenities when evaluating other projects. The fall in growth around Jardim do Mar, Portugal and Mundaka, Spain, after their breaks disappeared should be a cautionary tale for policymakers in places like Doughmore, Ireland, where Trump International Golf Links seeks to build a 2.8km seawall; and Jeffreys Bay, South Africa, where a new power plant may lead to 6.3 million cubic meters of sand being pumped offshore (www.savethewaves.org).

The paper also suggests further avenues for research. One is to study the local impacts of natural amenities at a firm level, to disentangle the mechanisms through which they affect growth. This would distinguish between tourists, permanent employees accepting lower wages, and entrepreneurs accepting lower profits to enjoy the amenities, which is beyond the scope of the present study. Another extension is to estimate the economic contribution of other natural amenities in a spatially disaggregated model, with obvious candidates including scuba-diving, rock-climbing, and UNESCO natural heritage sites. Future work may also directly study the feedback effects of economic growth onto the quality of natural amenities, through pollution and overcrowding. Finally, amenities like surf breaks may be a useful as instruments when studying the impact of economic growth on other variables at a local level.

References

- [1] Acemoglu, D., S. Johnson, and J. A. Robinson (2002). Reversal of fortune: Geography and institutions in the making of the modern world income distribution. *The Quarterly Journal of Economics* 117(4), 1231–1294.
- [2] Allcott, H. and D. Keniston (2014). Dutch disease or agglomeration? the local economic effects of natural resource booms in modern america. Technical report, National Bureau of Economic Research.
- [3] Aragón, F. M. and J. P. Rud (2013). Natural resources and local communities: evidence from a peruvian gold mine. *American Economic Journal: Economic Policy* 5(2), 1–25.
- [4] Bleakley, H. and J. Lin (2012). Portage and path dependence. *The quarterly journal of economics* 127(2), 587.
- [5] Bleakley, H. and J. Lin (2015). History and the sizes of cities. *The American Economic Review* 105(5), 558–563.
- [6] Brokensha, C. (2012). Skeleton bay: The death of a wave foretold. *Swellnet.com*.

- [7] Brunner, A. D. (2002). El nino and world primary commodity prices: warm water or hot air? *Review of Economics and statistics* 84(1), 176–183.
- [8] Butt, T. (2009). *The Surfer's Guide to Waves, Coasts and Climates*. Alison Hodge Publishers.
- [9] Butt, T. and P. Russell (2004). *Surf science: An introduction to waves for surfing*. University of Hawaii Press.
- [10] Caselli, F. and G. Michaels (2013). Do oil windfalls improve living standards? evidence from brazil. *American Economic Journal: Applied Economics* 5(1), 208–238.
- [11] Chay, K. and M. Greenstone (2005). Does air quality matter? evidence from the housing market. *Journal of Political Economy* 113(2), 376–424.
- [12] Chen, X. and W. D. Nordhaus (2011). Using luminosity data as a proxy for economic statistics. *Proceedings of the National Academy of Sciences* 108(21), 8589–8594.
- [13] Cheshire, P. C. and S. Magrini (2006). Population growth in european cities: weather matters—but only nationally. *Regional studies* 40(1), 23–37.
- [14] Coffman, M. and K. Burnett (2009). The value of a wave: An analysis of the mavericks region, half moon bay, california. *Study commission for Save The Waves Coalition*.
- [15] Conley, T. G. (1999). Gmm estimation with cross sectional dependence. *Journal of econometrics* 92(1), 1–45.
- [16] Conley, T. G. (2010). Spatial econometrics. In *Microeconometrics*, pp. 303–313. Springer.
- [17] Cust, J. and S. Poelhekke (2015). The local economic impacts of natural resource extraction. *Annu. Rev. Resour. Econ.* 7(1), 251–268.
- [18] Dean, J. (2016). Kelly slater built the perfect wave. can he sell it to the world? *Bloomberg Businessweek* May 24.
- [19] Dell, M. (2010). The persistent effects of peru's mining mita. *Econometrica* 78(6), 1863–1903.
- [20] Deller, S. C., T.-H. Tsai, D. W. Marcouiller, and D. B. K. English (2001). The role of amenities and quality of life in rural economic growth. *American Journal of Agricultural Economics* 83(2), 352–365.
- [21] Diamond, J. (1999). *Guns, germs, and steel: The fates of human societies*. WW Norton & Company.

- [22] Diamond, R. (2016). The determinants and welfare implications of us workers' diverging location choices by skill: 1980–2000. *The American Economic Review* 106(3), 479–524.
- [23] Donaldson, D. and A. Storeygard (2016). Big grids: Applications of remote sensing in economics. *forthcoming, JEP*.
- [24] Ellison, G. and E. L. Glaeser (1997). Geographic concentration in us manufacturing industries: A dartboard approach. *Journal of Political Economy* 105(5).
- [25] Ellison, G. and E. L. Glaeser (1999). The geographic concentration of industry: does natural advantage explain agglomeration? *The American Economic Review* 89(2), 311–316.
- [26] Ellison, G., E. L. Glaeser, and W. R. Kerr (2010). What causes industry agglomeration? evidence from coagglomeration patterns. *The American Economic Review* 100(3), 1195–1213.
- [27] Faber, B. and C. Gaubert (2015). Tourism and economic development: Evidence from mexicos coastline. *UC Berkeley Economics Department working paper*.
- [28] Fetzer, T. (2014). Can workfare programs moderate violence? evidence from india.
- [29] Fujita, M., P. R. Krugman, and A. Venables (2001). *The spatial economy: Cities, regions, and international trade*. MIT press.
- [30] Gallup, J. L. and J. D. Sachs (2000). Agriculture, climate, and technology: why are the tropics falling behind? *American Journal of Agricultural Economics* 82(3), 731–737.
- [31] Glaeser, E. L. and J. D. Gottlieb (2009). The wealth of cities: Agglomeration economies and spatial equilibrium in the united states. *Journal of Economic Literature* 47(4), 983–1028.
- [32] Handler, P. and E. Handler (1983). Climatic anomalies in the tropical pacific ocean and corn yields in the united states. *Science* 220(4602), 1155–1156.
- [33] Hemer, M., J. Jack Katzfey, and C. Hotan (2011). The wind-wave climate of the pacific ocean. *Report for the Pacific Adaptation Strategy Assistance Program Department of Climate Change and Energy Efficiency*.
- [34] Henderson, J. V., A. Storeygard, and D. N. Weil (2011). A bright idea for measuring economic growth. *The American economic review* 101(3), 194.

- [35] Henderson, J. V., A. Storeygard, and D. N. Weil (2012). Measuring economic growth from outer space. *The American Economic Review*, 994–1028.
- [36] Hodler, R. and P. A. Raschky (2014). Regional favoritism. *The Quarterly Journal of Economics*, qju004.
- [37] Housman, J. (2016). El niño’s relentless perfection. *Surfer Magazine*.
- [38] Hsiang, S. M. (2010). Temperatures and cyclones strongly associated with economic production in the caribbean and central america. *Proceedings of the National Academy of sciences* 107(35), 15367–15372.
- [39] Hsiang, S. M. and K. C. Meng (2015). Tropical economics. *The American Economic Review* 105(5), 257–261.
- [40] Hsiang, S. M., K. C. Meng, and M. A. Cane (2011). Civil conflicts are associated with the global climate. *Nature* 476(7361), 438–441.
- [41] Iizumi, T., J.-J. Luo, A. J. Challinor, G. Sakurai, M. Yokozawa, H. Sakuma, M. E. Brown, and T. Yamagata (2014). Impacts of el niño southern oscillation on the global yields of major crops. *Nature communications* 5.
- [42] Jedwab, R., E. Kerby, and A. Moradi (2015). History, path dependence and development: Evidence from colonial railroads, settlers and cities in kenya. *The Economic Journal*.
- [43] Jedwab, R. and A. Moradi (2016). The permanent effects of transportation revolutions in poor countries: evidence from africa. *Review of economics and statistics* 98(2), 268–284.
- [44] Lee, S. and J. Lin (2015). Natural amenities, neighborhood dynamics, and persistence in the spatial distribution of income.
- [45] Longman, J. (2016). Surfing under northern lights. *The New York Times* 18 October.
- [46] Marshall, A. (1890). Principles of economics. *London: Macmillan*.
- [47] Michaels, G. and F. Rauch (2016). Resetting the urban network: 117-2012. *The Economic Journal* forthcoming.
- [48] Michalopoulos, S. and E. Papaioannou (2013). Pre-colonial ethnic institutions and contemporary african development. *Econometrica* 81(1), 113–152.
- [49] Michalopoulos, S. and E. Papaioannou (2014). National institutions and subnational development in africa. *The Quarterly Journal of Economics* 129(1), 151–213.

- [50] North, D. C. (1989). Institutions and economic growth: An historical introduction. *World development* 17(9), 1319–1332.
- [51] North, D. C. (1990). *Institutions, institutional change and economic performance*. Cambridge university press.
- [52] Redding, S. J. (2009). Economic geography: a review of the theoretical and empirical literature.
- [53] Redding, S. J. (2010). The empirics of new economic geography. *Journal of Regional Science* 50(1), 297–311.
- [54] Redding, S. J. and E. Rossi-Hansberg (2016). Quantitative spatial economics. Technical report, National Bureau of Economic Research.
- [55] Roback, J. (1982). Wages, rents, and the quality of life. *The Journal of Political Economy*, 1257–1278.
- [56] Roback, J. (1988). Wages, rents, and amenities: differences among workers and regions. *Economic Inquiry* 26(1), 23–41.
- [57] Rodrik, D. (1996). Coordination failures and government policy: A model with applications to east asia and eastern europe. *Journal of international economics* 40(1), 1–22.
- [58] Rodrik, D. (2004). Industrial policy for the twenty-first century.
- [59] Rosen, S. (1974). Hedonic prices and implicit markets: product differentiation in pure competition. *Journal of political economy* 82(1), 34–55.
- [60] Sachs, J. (2000). Notes on a new sociology of economic development. *Culture matters: How values shape human progress*, 29–43.
- [61] Sachs, J. D. (2001). Tropical underdevelopment. Technical report, National Bureau of Economic Research.
- [62] Scorse, J., I. Reynolds, A. Sackett, et al. (2015). Impact of surf breaks on home prices in santa cruz, ca. *Tourism Economics* 21(2), 409–418.
- [63] Smith, B. and S. Wills (2016). Left in the dark? oil and rural poverty. Technical report, Oxford Centre for the Analysis of Resource Rich Economies, University of Oxford.
- [64] Surfer Magazine (2006). The best surfers in the world land somewhere in mexico for the rip curl pro. *Surfer Magazine*.

- [65] The Economist (2012). Surf economics: Beach rush. *The Economist March 17*.
- [66] Ubilava, D. (2012). El nino, la nina, and world coffee price dynamics. *Agricultural Economics 43*(1), 17–26.
- [67] Van Der Ploeg, F. and S. Poelhekke (2016). The impact of natural resources: Survey of recent quantitative evidence. *The Journal of Development Studies*, 1–12.
- [68] Waltert, F. and F. Schläpfer (2010). Landscape amenities and local development: A review of migration, regional economic and hedonic pricing studies. *Ecological Economics 70*(2), 141–152.
- [69] Wu, J. and M. Gopinath (2008). What causes spatial variations in economic development in the united states? *American Journal of Agricultural Economics 90*(2), 392–408.

Appendix

A What makes a good wave?

Waves are created by wind acting on the surface of the ocean. The wind creates lines of swell which propagate up to 15,000km away from the originating low pressure system. During their travels the waves disperse, as longer wavelengths speed to the front; and group, as different wavelengths cancel and amplify one another. On their journey the waves warp and bend as they interact with the underwater topography (bathymetry), in a process called refraction. This can make the waves bigger, smaller, longer, shorter, faster, slower, fatter or hollower, depending on the bathymetry along the way. As waves approach the coast they break: the base slows as it interacts with the sea floor while the peak continues at speed. It is at the point of breaking that surfers draw off some of the waves’ energy for their sport.²⁰ After the wave breaks it imparts energy to the shore, shaping sand bars which in turn affect breaking in a “self-organising system” (Butt and Russell, 2004).²¹

The quality of a surfing wave depends, broadly, on three characteristics: size, length and shape. The size of a wave is dictated by the strength, direction and duration of the

²⁰Surfers ride waves by balancing the force of water pushing the surfboard up the wave’s face against gravity pulling it down.

²¹Butt and Russell (2004) provide an introduction to oceanography and coastal engineering with a focus on surfing.




Figure A.1: A deep trench in the Bay of Biscay allows waves generated in the North Atlantic to travel uninterrupted to the Basque region’s coast, creating good surf breaks.

originating winds, the area over which they act, the distance to where they are ridden, and the bathymetry in between (see Figure A.1). Wave length is determined more locally by the angle at which the swell meets the shore, which in turn depends on the direction of swell, and how it refracts and the shape of the coast. Wave shape depends on local bathymetry, swell period and wind direction. If the sea floor rises sharply then the base of the wave will slow abruptly, relative to the peak, and the wave will pitch to create a “barrel”. This is exacerbated for long period swells, which travel faster; and offshore (from land to sea) winds, which hold up the wave and cancel out short-period windswell.

From this we see that the quality of a surf break is essentially random. Not only does it depend on the direction and strength of swell at its source, and wind and coastal characteristics locally, but it also depends on the bathymetry over the entire intervening distance, down to a very fine scale. Sand-bars, which dictate bathymetry for the sandy locations that account for over 50% of our sample, are formed “in a chaotic system, where imperceptibly small changes in input produce vastly differing outputs” (Butt and Russell, 2004). While some of these individual characteristics may be correlated with economic outcomes at a local level, the fine balance of many factors needed to produce good waves is unlikely to be.

B Overglow

Overglow is when light emitted in one pixel is also recorded in surrounding pixels. It has been documented as an issue when light is used to measure the level of economic activity (Doll, 2008). Here we test how overglow changes with emitted light in two ways: comparing light either side of a wasteland boundary (following Pinkovski, 2013), and light emitted from isolated pixels. Both tests find that overglow varies linearly with emitted light (as do Small et al., 2005), so should not bias our study of light growth.

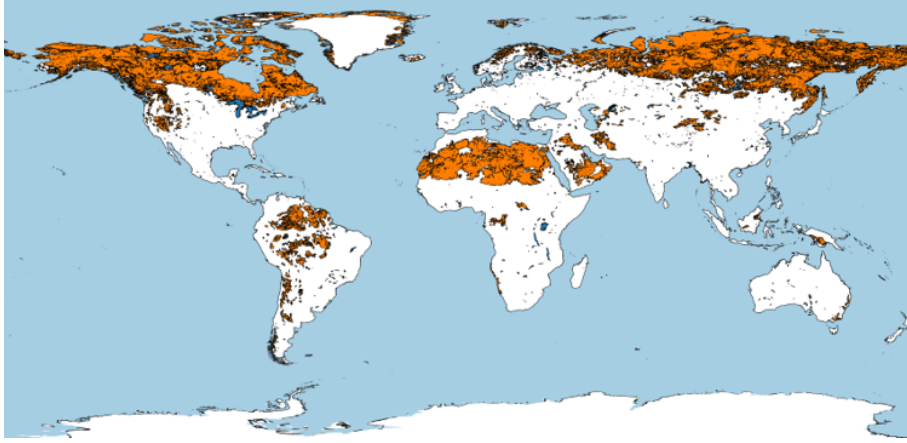


Figure B.1: 1000 largest wasteland areas in the Anthropogenic Biomes v2 dataset.

	$\Delta Y_{j,2006}^{in}$
$\Delta Y_{j,2006}^{out}$	0.870*** (0.187)
Constant	0.121* (0.0669)
Observations	457
R-squared	0.045
Standard errors in parentheses	
*** p<0.01, ** p<0.05, * p<0.1	

Table B.1: Results of regressing light growth 10km inside wasteland borders on that 10km outside.

B.1 Test 1: Wasteland boundaries

We study overglow by comparing light either side of a wasteland border, following Pinkovskiy (2013). Wastelands are unpopulated areas classified as “wild woodlands, wild treelands and barren lands” in 2000, in the Anthropogenic Biomes Version 2 datasets (Ellis et al., 2010). Light may be emitted outside the wasteland, but any light inside should be overglow. We study lights 10km either side of the border for the 1000 largest wasteland areas (Figure B.1), using the following specification,

$$\Delta Y_{j,2006}^{in} = \alpha + \beta \Delta Y_{j,2006}^{out} + \epsilon_t \quad (\text{B.1})$$

where $\Delta Y_{j,2006}^{in}$ and $\Delta Y_{j,2006}^{out}$ is the change in light from 2005-2006, for 10km bands inside and outside the border of wasteland j respectively. The results in Table B.1 estimate $\beta = 0.87$, which is not statistically different from 1 at more than the 1% level. This suggests that overglow varies proportionally with emitted light.

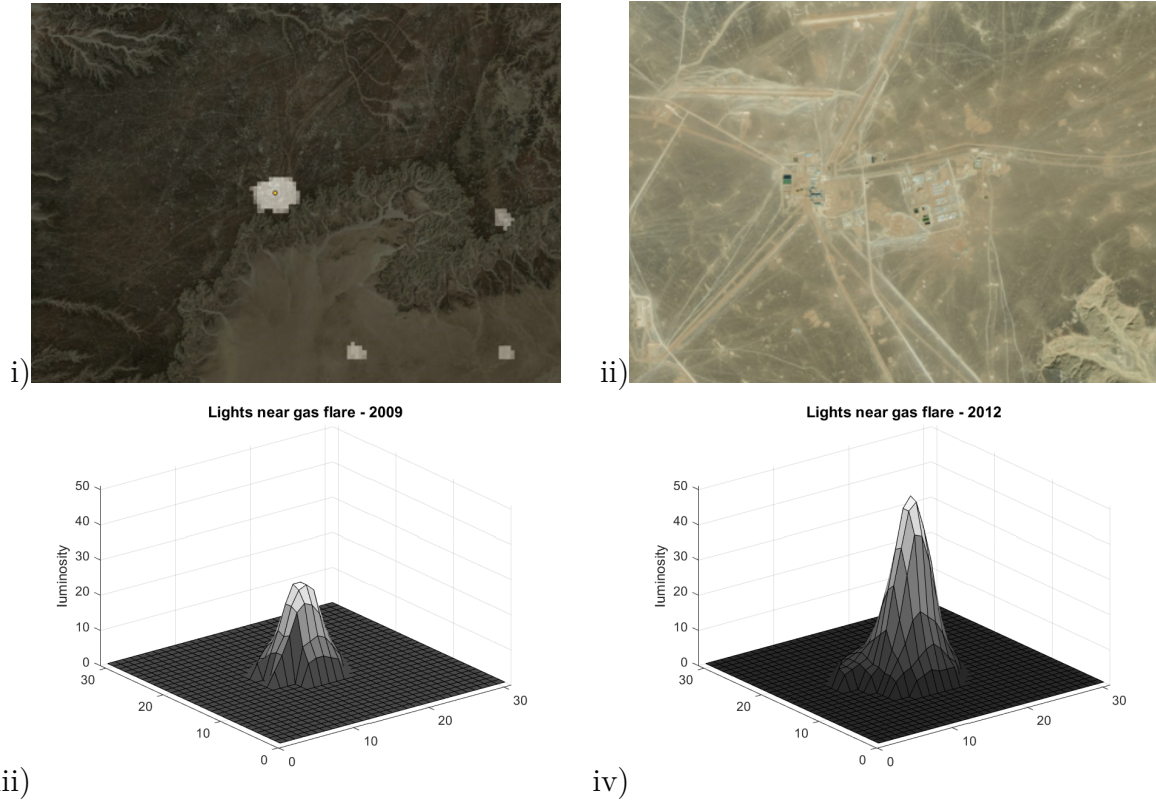


Figure B.2: Example isolated gas flare site in Algerian wasteland: i) night-time lights, ii) google earth (zoomed), iii) 2009 light histogram, iv) 2012 light histogram.

B.2 Test 2: Isolated light sources

We also study overflow from a small number of isolated light sources in Algeria, using gas flare data from the NOAA Earth Observation Group, and wasteland data described above. We use gas flares that i) are in wastelands with no habitation or physical economic activity within 15km, ii) come from a $< 1\text{km}^2$ site as identified on Google Earth, and iii) have the brightest light in the central pixel (Figure B.2). This gives us an insight into how overflow is recorded by the DMSP-OLS satellites.

Figure B.3 illustrates the change in light between 2009 and 2012 at various distances from the flare. It shows that overflow grows proportionally to light in the central pixel, for non-trivial levels of light. This may be due to differences in the intensity of emitted light or the sensitivity of the satellite. Figure B.3ii. illustrates the difference in light between the original site, and another in the Algerian desert, both in 2012. Overflow differs from 15-30% but not systematically. We conclude that overflow varies proportionally to emitted light, and so will not bias our study of changes in light emissions.

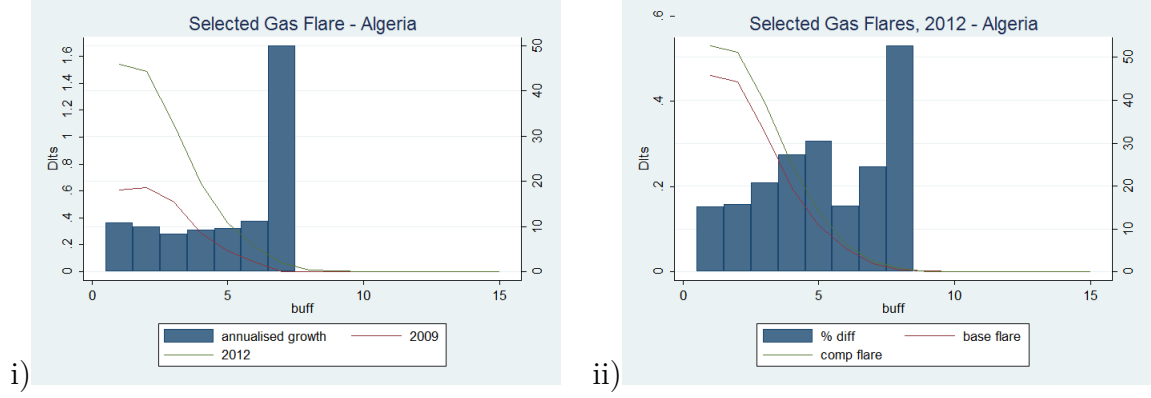


Figure B.3: Overglow comparison: i) same site, 2009 vs 2012, ii) different sites, 2012.

C Experiment I: Additional results

C.1 Surfing and economic activity

	Total difference in $\ln(\text{lights})$ in 5 to 10km band from 1992-2013					
	(1)	(2)	(3)	(4)	(5)	(6)
Quality (cts)	0.0196 (0.141)	0.0201* (0.0918)	0.0196 (0.162)			
2 star				0.0336 (0.351)	0.0257 (0.493)	0.0336 (0.326)
3 star				0.0588 (0.139)	0.0706* (0.0541)	0.0588 (0.117)
4 star				0.0982** (0.0406)	0.0649 (0.136)	0.0982* (0.0766)
5 star				0.0151 (0.870)	0.0151 (0.830)	0.0151 (0.835)
Observations	4,480	4,480	4,480	4,480	4,480	4,480
R-squared	0.401	0.531	0.001	0.402	0.532	0.001
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
FE	Country	Zone	Country	Country	Zone	Country
SE	Country	Zone	Conley	Country	Zone	Conley

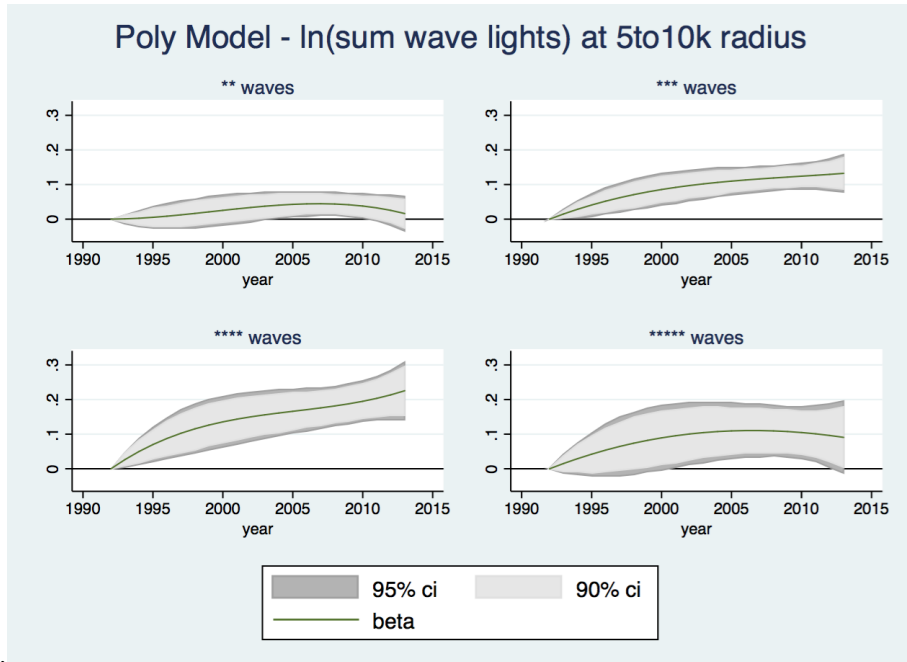
Robust p-values in parentheses
 *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table C.1: The effect of break quality on the change in $\ln(\text{lights})$ from 5-10km of each break, from 1992-2013. Fixed effects are at country and zone level. Standard errors are clustered at the country and zone level, and allow for spatial correlation within 100km, and autocorrelation to 3 periods (Conley).

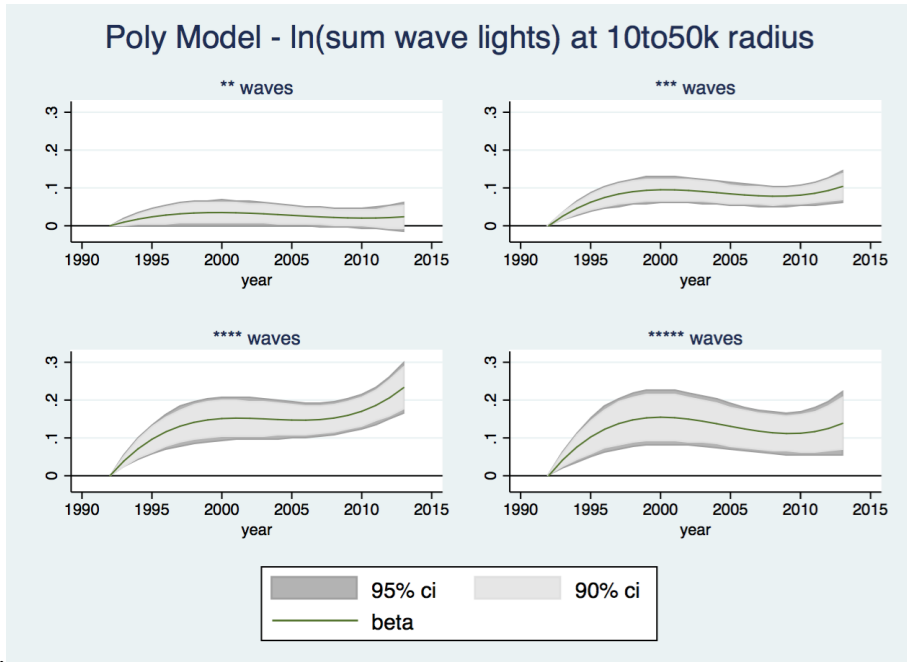
Total difference in $\ln(\text{lights})$ in 10 to 50km band from 1992-2013						
	(1)	(2)	(3)	(4)	(5)	(6)
Quality (cts)	0.0166** (0.0278)	0.0126 (0.387)	0.0166 (0.137)			
2 star				0.0225 (0.250)	0.0225 (0.215)	0.0225 (0.273)
3 star				0.0344 (0.106)	0.0466** (0.0434)	0.0344 (0.149)
4 star				0.0639* (0.0761)	0.0279 (0.516)	0.0639 (0.121)
5 star				0.0623* (0.0522)	0.0489 (0.232)	0.0623 (0.202)
Observations	5,004	5,004	5,004	5,004	5,004	5,004
R-squared	0.582	0.728	0.001	0.582	0.728	0.001
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
FE	Country	Zone	Country	Country	Zone	Country
SE	Country	Zone	Conley	Country	Zone	Conley

Robust p-values in parentheses
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table C.2: The effect of break quality on the change in $\ln(\text{lights})$ from 10-50km of each break, from 1992-2013. Fixed effects are at country and zone level. Standard errors are clustered at the country and zone level, and allow for spatial correlation within 100km, and autocorrelation to 3 periods (Conley).



i.



ii.

Figure C.1: Polynomial model of the difference in difference of light levels at i. 5 to 10km, and ii. 10 to 50km from each break, using 1-star breaks as the baseline. Includes year and break fixed effects, and Conley standard errors.

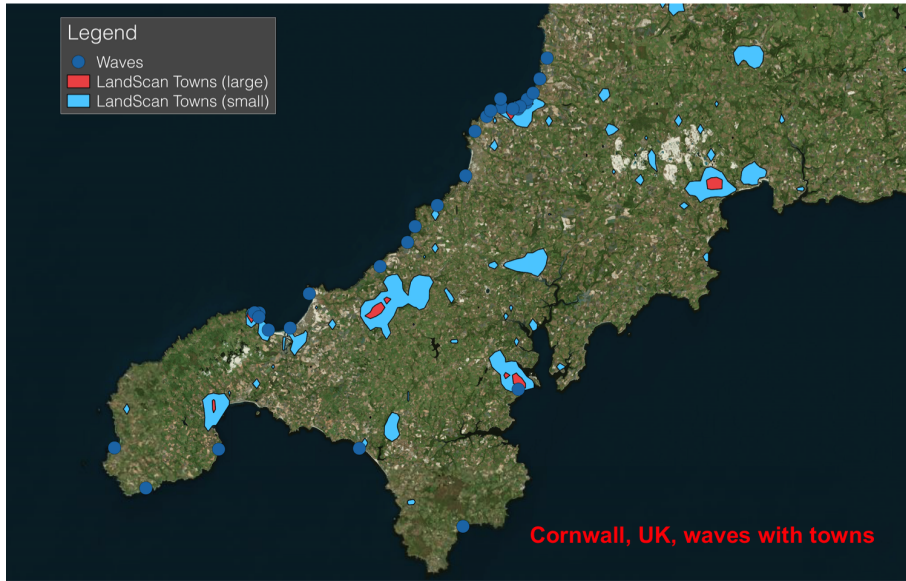


Figure C.2: Example of endogenously defining the location of towns by their population density, Cornwall, UK.




Figure C.3: Distribution of surfing experience required to surf breaks of each quality. 5-star waves are disproportionately rated “3: Pros or kamikazes only”.

C.2 Ease of doing business and political stability categories

World Bank Doing Business Index							
Very Low		Low		Moderate		High	
Country	Breaks	Country	Breaks	Country	Breaks	Country	Breaks
Brazil	291	South Africa	207	Spain	182	France	290
Indonesia	136	Puerto Rico	53	Japan	118	Portugal	162
Ecuador	47	Chile	50	Italy	113	UK	149
Argentina	34	Costa Rica	50	Mexico	92	New Zealand	114
Venezuela	24	Greece	50	Peru	76	Ireland	58
Nicaragua	20	Morocco	43	Mauritius	14	Canada	38
Senegal	20	Philippines	29	Belgium	12	Netherlands	32
Barbados	19	Panama	27	UAE	9	Germany	22
Sri Lanka	18	Israel	25	Bulgaria	8	Taiwan	18
India	13	Uruguay	24	Poland	4	Denmark	15
Micronesia	11	Turkey	18	Croatia	3	Iceland	14
Bahamas	10	Namibia	17	Malaysia	14		
PNG	10	Dominican Rep.	16	Sweden	11		
Angola	9	Seychelles	15	South Korea	7		
Maldives	9	Russia	14	Hong Kong	5		
Verde	9	Thailand	14	Switzerland	5		
Ghana	8	El Salvador	12	Lithuania	3		
Madagascar	8	China	11	Estonia	2		
Egypt	7	Colombia	11	Finland	2		
Mozambique	7	Tunisia	11	Latvia	2		
Lebanon	6	Samoa	10	Austria	1		
Guinea	5	Vietnam	10				
Algeria	4	Cyprus	8				
Liberia	4	Guatemala	7				
Sao Tome and P.	4	Dominica	6				
Togo	4	Fiji	6				
Cameroon	3	Malta	5				
Cote d'Ivoire	3	Oman	4				
Gambia	3	Brunei	3				
Kenya	3	St Lucia	3				
Myanmar	3	Albania	2				
Nigeria	3	Tonga	2				
Sierra Leone	3	Vanuatu	2				
Tanzania	3	Jamaica	1				
Benin	2	Kuwait	1				
Gabon	2	Qatar	1				
Grenada	2	Trinidad and Tob.	1				
Haiti	2	Ukraine	1				
Kiribati	2						
Palau	2						
Rep Congo	2						
Bangladesh	1						
Belize	1						
Cambodia	1						
Honduras	1						
Iran	1						
St Kitts And Nev.	1						
St Vincent And T.	1						
Solomon Islands	1						
Timor-Leste	1						
Zimbabwe	1						
Total	785			770	631		964

Table C.3: Countries and break count by World Bank Doing Business index categories.

World Bank Worldwide Governance Indicators, Political Stability category							
Very Low		Low		Moderate		High	
Country	Breaks	Country	Breaks	Country	Breaks	Country	Breaks
South Africa	207	Brazil	291	France	290	Portugal	162
Indonesia	136	Spain	182	UK	149	Japan	118
Mexico	92	Greece	50	Italy	113	New Zealand	114
Peru	76	Ecuador	47	Puerto Rico	53	Ireland	58
Morocco	43	Argentina	34	Chile	50	Canada	38
Philippines	29	Panama	27	Costa Rica	50	Netherlands	32
Israel	25	Reunion	21	Namibia	17	Uruguay	24
Venezuela	24	Nicaragua	20	Seychelles	15	Germany	22
Senegal	20	Dominican Rep.	16	Malaysia	14	Barbados	19
Sri Lanka	18	Vietnam	10	Belgium	12	Taiwan	18
Turkey	18	Bulgaria	8	Verde	9	Denmark	15
Russia	14	South Korea	7	Cyprus	8	Saint Martin	15
Thailand	14	Sao Tome and P.	4	Fiji	6	Iceland	14
India	13	Gabon	2	Oman	4	Mauritius	14
El Salvador	12	Benin	2	Croatia	3	Micronesia	11
China	11	Belize	1	Albania	2	Sweden	11
Colombia	11	Trinidad and Tob.	1	French Guiana	2	Bahamas	10
Tunisia	11	Kuwait	1	Kiribati	2	Samoa	10
PNG	10	Jamaica	1	Latvia	2	Maldives	9
Angola	9	Cambodia	1	Vanuatu	2	UAE	9
Ghana	8	St. Kitts and Nev.	1	Aruba	7		
Madagascar	8	Solomon Islands	1	Dominica	6		
Egypt	7	Hong Kong	5				
Guatemala	7	Malta	5				
Mozambique	7	Switzerland	5				
Lebanon	6	Poland	4				
Guinea	5	Virgin Islands	4				
Algeria	4	Brunei	3				
Liberia	4	Lithuania	3				
Togo	4	Saint Lucia	3				
Cameroon	3	Anguilla	2				
Cote d'Ivoire	3	Bermuda	2				
Gambia	3	Estonia	2				
Kenya	3	Finland	2				
Myanmar	3	Grenada	2				
Nigeria	3	Tonga	2				
Sierra Leone	3	Austria	1				
Tanzania	3	Qatar	1				
Haiti	2	St Vincent and T.	1				
Rep Congo	2						
Bangladesh	1						
Honduras	1						
Iran	1						
Somalia	1						
Timor-Leste	1						
Ukraine	1						
Zimbabwe	1						
Total	888		726		805		783

Table C.4: Countries and break count by World Bank Political Stability index categories

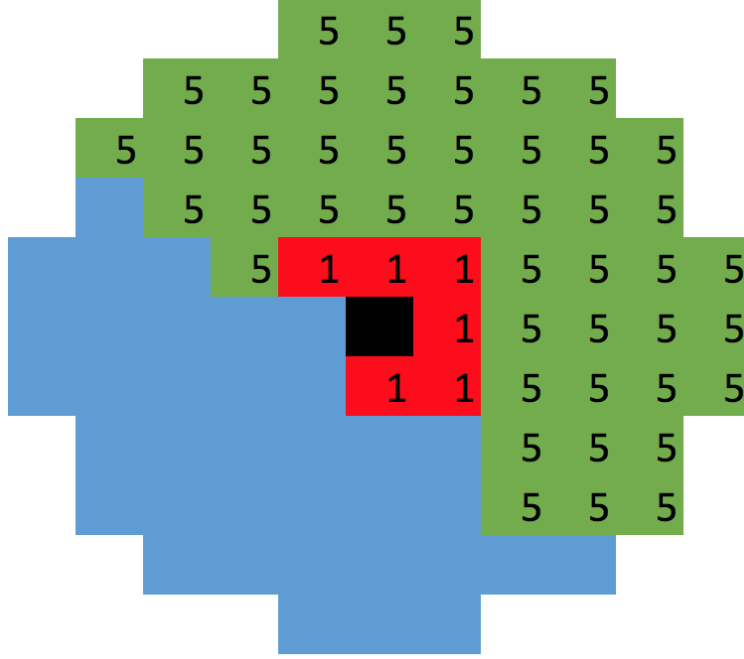


Figure D.1: Pixels within 1km (red) and 5km (green) of a break (black), excluding the area over water (blue).

D Experiment I: Robustness

D.1 Pixel-level analysis

The analysis in the main text uses the area within 5km of a surf break as the unit of analysis. These areas may overlap. Here we use individual pixels as the unit of analysis, which are “treated” by surf breaks to give an estimate of the marginal contribution of each break quality to illumination. Every pixel within 5km of at least one break is included in our sample, excluding those wholly over water, which increases the sample size considerably (see Figure D.1). Each pixel is treated in two ways using the specification,

$$\ln(Y_{i,2013}) - \ln(Y_{i,1992}) = \alpha + \beta Q_i + W_i + Z_t + \epsilon_{i,t} \quad (\text{D.1})$$

where $Y_{i,t}$ is light intensity in pixel i at time t , $Q_i \in [M_i, N_{i,k}]$, M_i is the maximum quality of breaks within 5km of pixel i (treated as both a continuous and a categorical variable), $N_{i,k}$ is the number of breaks of quality $k \in [1, 5]$ within 5km of pixel i , W_i is country fixed effects, and Z_t is year fixed effects. The results are given in Table D.1 and discussed in Section 3.3.3.

To give a visual representation of results we study a pixel-level analogue of equation 3.2,

- i. Sum of breaks of quality k within 5km ii. Maximum break quality within 5km

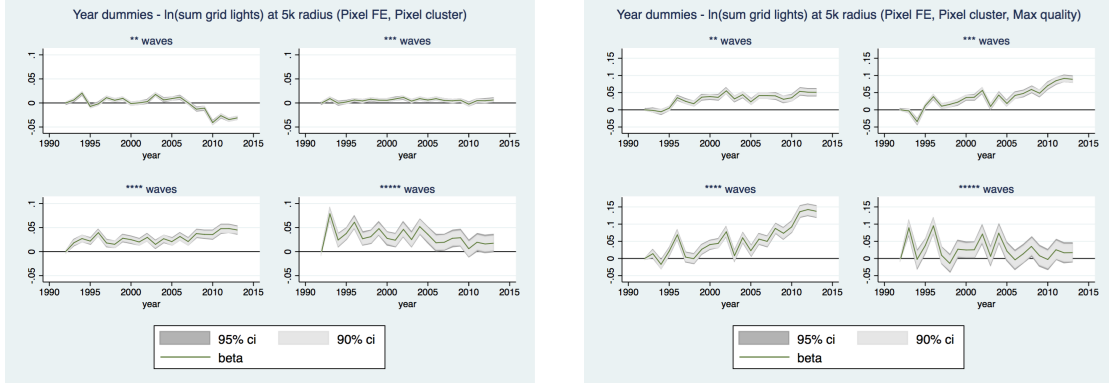


Figure D.2: Results of pixel-level analysis, treating each pixel with i. the number of breaks of quality k within 5km, and ii. the maximum break quality within 5km.

$$\ln(Y_{i,t}) = \alpha + \sum_{t=1992}^{2013} \beta_t T_t Q_i + F_i + T_t + \epsilon_{i,t} \quad (\text{D.2})$$

where $Y_{i,t}$ is light intensity in pixel i at time t , T_t is a year dummy, $Q_i \in [M_i, N_{i,k}]$ is defined above, and F_i are pixel fixed effects. Standard errors are clustered at the break and country level. The control group is 1-star breaks. The results are given in Figure D.2 and discussed in Section 3.3.3.

To control for possible effects of convergence between pixels with low and high we also control for the initial level of lights using the specification,

$$\Delta \ln(Y_{i,t}) = \alpha + \beta Q_i + \ln(Y_{i,1992}) + W_i + Z_t + \epsilon_{i,t} \quad (\text{D.3})$$

where $\Delta \ln(Y_{i,t})$ is the annual change in log lights (so $\ln(Y_{i,1992})$ doesn't appear on the RHS in all observations), and the other variables are defined above. These results are given in Table D.5 and discussed in Section 3.3.3.

Pixel-level Analysis						
Total difference in ln(pixel lights) from 1992-2013						
	(1)	(2)	(3)	(4)	(5)	(6)
Maximum quality:						
Continuous	0.0240*	0.0240**				
	(0.0550)	(0.0465)				
2 star			0.0420	0.0420		
			(0.146)	(0.229)		
3 star			0.0640**	0.0640*		
			(0.0340)	(0.0561)		
4 star			0.0659	0.0659*		
			(0.178)	(0.0703)		
5 star			0.100	0.100		
			(0.349)	(0.430)		
Quality count:						
1 star					-0.102***	-0.102***
					(7.49e-06)	(2.65e-05)
2 star					-0.0505***	-0.0505***
					(4.55e-08)	(0.000194)
3 star					-0.0232**	-0.0232***
					(0.0143)	(0.00886)
4 star					-0.00397	-0.00397
					(0.862)	(0.793)
5 star					0.0279	0.0279
					(0.530)	(0.666)
Constant	0.642***	0.642***	0.629***	0.629***	0.744***	0.744***
	(0)	(0)	(0)	(0)	(0)	(0)
Observations	91,353	91,353	91,353	91,353	91,353	91,353
R-squared	0.358	0.358	0.358	0.358	0.363	0.363
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
SE	Break	Country	Break	Country	Break	Country

Robust p-values in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table D.1: The effect of break quality on the change in ln(lights) for each pixel within 5km of a surf break, from 1992-2013. Columns (1)-(2) treat each pixel with the maximum break quality within 5km. Columns (3)-(4) do the same using max-quality indicators. Columns (5)-(6) treat each pixel with the number of breaks of each quality within 5km.

D.2 Omitted Geographic Characteristics

Break Type	Frequency	Percent
Beach-break	2,070	41%
Reef-rocky	1,057	21%
Point-break	670	13%
Sand-bar	570	11%
Reef-coral	453	9%
Breakwater/jetty	132	3%
River-mouth	120	2%
Reef-artificial	24	0%
Don't know	12	0%
Total	5,108	100%

Table D.2: Breakdown of surf breaks by type.

Total difference in $\ln(\text{lights}^{5km})$ from 1992-2013								
	excluding Reefs, Point-breaks and River-mouths				excluding Australia and USA			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Quality (cts)	0.0305** (0.0475)	0.0338* (0.0697)			0.0441*** (0.00193)	0.0401*** (0.00818)		
2 star			0.0224 (0.509)	0.0220 (0.563)			0.0483 (0.304)	0.0429 (0.369)
3 star			0.0671* (0.0567)	0.0862** (0.0300)			0.101** (0.0246)	0.103** (0.0253)
4 star			0.0535 (0.536)	0.0389 (0.689)			0.126*** (0.00975)	0.0826 (0.197)
5 star			0.0960 (0.447)	0.0165 (0.873)			0.167* (0.0697)	0.172* (0.0629)
Constant	0.445*** (0)	0.437*** (0)	0.479*** (0)	0.474*** (0)	0.560*** (0)	0.570*** (0)	0.598*** (0)	0.603*** (0)
Observations	2,512	2,512	2,512	2,512	2,791	2,791	2,791	2,791
R-squared	0.413	0.542	0.413	0.543	0.396	0.557	0.396	0.557
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FE	Country	Zone	Country	Zone	Country	Zone	Country	Zone
SE	Country	Zone	Country	Zone	Country	Zone	Country	Zone

Robust p-values in parentheses

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

Table D.3: The effect of break quality on the change in $\ln(\text{lights})$ within 5km of each break, from 1992-2013. Columns (1)-(4) exclude reefs, pointbreaks and rivermouths, and columns (5)-(8) exclude Australia and the USA.

D.3 Selection




Figure D.3: Breakdown of the 5km surrounding breaks of each quality

2 star baseline		
Total difference in $\ln(\text{lights}^{5km})$ from 1992-2013		
	(1)	(2)
1 star	-0.0196 (0.531)	-0.0297 (0.388)
3 star	0.0611*** (0.00461)	0.0716** (0.0201)
4 star	0.0682** (0.0174)	0.0560 (0.153)
5 star	0.0728 (0.413)	0.0651 (0.237)
Constant	0.521*** (0)	0.518*** (0)
Observations	4,289	4,289
R-squared	0.390	0.523
Year FE	Yes	Yes
FE	Country	Zone
SE	Country	Zone

Robust p-values in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Table D.4: The effect of break quality on the change in $\ln(\text{lights})$ within 5km of each break, from 1992-2013, using 2-star breaks as the baseline.

	Annual difference in $\ln(\text{pixellights})$ from 1992-2013			
	(1)	(2)	(3)	(4)
Maximum quality:				
2 star	0.00136*** (5.87e-08)	0.00136 (0.240)		
3 star	0.00143*** (3.93e-08)	0.00143 (0.226)		
4 star	0.00340*** (0)	0.00340* (0.0938)		
5 star	-0.00666*** (0)	-0.00666** (0.0307)		
Quality count:				
1 star			-0.00124*** (0)	-0.00124 (0.150)
2 star			0.000117 (0.139)	0.000117 (0.740)
3 star			0.000200** (0.0156)	0.000200 (0.561)
4 star			0.00173*** (0)	0.00173* (0.0559)
5 star			-0.00227*** (6.72e-07)	-0.00227 (0.223)
$\ln(\text{pixellights})$ in 1992	-0.00857*** (0)	-0.00857*** (0)	-0.00857*** (0)	-0.00857*** (0)
Constant	0.0402*** (0)	0.0402*** (0)	0.0413*** (0)	0.0413*** (0)
Observations	1,877,278	1,877,278	1,877,278	1,877,278
R-squared	0.139	0.139	0.139	0.139
Year FE	Yes	Yes	Yes	Yes
SE	Pixel	Break	Pixel	Break

Robust p-values in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table D.5: The effect of break quality on the annual change in $\ln(\text{lights})$ for each pixel within 5km of a surf break, from 1992-2013. These results include the initial level of lights in 1992 to control for convergence.

E Experiment II: Additional Results

Detail: Discoveries and Disappearances		
	(1)	(2)
Year trend: post-event	0.0222*** (0.003)	-0.00495 (0.655)
Constant: post-event, interacted with:		
Skeleton Bay	-0.0660*** (0.001)	-0.172*** (4.74e-06)
Jardim do Mar	0.0504*** (0.000)	0.0839*** (8.19e-05)
Mundaka	0.120*** (1.05e-05)	0
La Jolla	-0.143*** (2.04e-05)	-0.0529 (0.138)
El Gringo	-0.201*** (8.28e-06)	0.129*** (0.006)
Year trend: pre-event, interacted with:		
Skeleton Bay	0.00540*** (5.67e-06)	0.0193** (0.0197)
Jardim do Mar	0.00522*** (0.004)	-0.00685 (0.253)
Mundaka	-0.00867*** (0.001)	-0.0130 (0.273)
La Jolla	-0.0140*** (1.30e-05)	-0.0104 (0.297)
El Gringo	0.0156*** (9.66e-07)	0.00753 (0.500)
Constant: pre-event, interacted with:		
Skeleton Bay	-12.11*** (3.58e-06)	-39.88** (-3.891)
Jardim do Mar	1.671 (0.459)	53.53* (0.0593)
Mundaka	30.99*** (0.00)	67.45* (0.0929)
LaJolla	38.40*** (8.05e-07)	59.03*** (0.00111)
ElGringo	-19.98*** (2.61e-07)	23.86* (0.0756)
Observations	110	110
R-squared	0.996	0.995
SE Cluster	Break	Break
(Robust p-values): *** p<0.01, ** p<0.05, * p<0.1		

Table E.1: Detail of Table 4.2. Difference-in-difference in average log lights around break discoveries and disappearances on (1) actual discovery dates, and (2) discovery dates drawn from a uniform U(1992,2013) distribution.

Cold-Water Breaks	
Country	Break Count
Denmark	15
Estonia	2
Faeroe Islands	1
Finland	2
Iceland	14
Ireland	9
Latvia	2
Lithuania	3
Russia	3
Sweden	11
United Kingdom	16
United States	5
Total	83

Table E.2: List of countries with surf breaks above +55 or below -55 degrees latitude.

Alternative specification: Cold Water Breaks				
	(1)	(2)	(3)	(4)
	Main	Main	Robust	Robust
Year trend: post-event	0.0285*** (2.95e-06)	0.0285** (0.0139)	0.00120 (0.931)	0.00120 (0.952)
Constant: post-event	0.103*** (0.000829)	0.103 (0.200)	0.00492 (0.720)	0.00492 (0.820)
Year trend: pre-event	-0.0160*** (1.92e-06)	-0.0160** (0.0283)	-0.192*** (0.000137)	-0.192* (0.0695)
Constant: pre-event	31.51*** (2.48e-06)	31.51** (0.0301)	-2.673 (0.924)	-2.673 (0.946)
Observations	1,804	1,804	1,804	1,804
R-squared	0.032	0.032	0.038	0.038
Number of wid	83	83	83	83
Sample	Lat \notin [-55, 55]	Lat \notin [-55, 55]	Lat \notin [-55, 55]	Lat \notin [-55, 55]
FE	Break	Break	Break	Break
SE Cluster	Break	Zone	Break	Zone

Robust p-values in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table E.3: Difference-in-difference in average log lights around the invention of battery-heated wetsuits in 2007 (columns 1-2), and using 1997 as a robustness test (3-4).