

Essays in Urban Development Economics

by

Alexei Sisulu Abrahams

B.Sc. Mathematics, Syracuse University, USA, 2008

B.Sc. Economics, Syracuse University, USA, 2008

M.A. Economics, Brown University, USA, 2010

A dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy in the
Department of Economics at Brown University

Providence, Rhode Island

May 2015

© Copyright 2015 by Alexei Sisulu Abrahams

This dissertation by Alexei Sisulu Abrahams is accepted in its present form
by the Department of Economics as satisfying the
dissertation requirements for the degree of Doctor of Philosophy.

Date _____
Nathaniel Baum-Snow, Advisor

Recommended to the Graduate Council

Date _____
Andrew Foster, Reader

Date _____
David Weil, Reader

Approved by the Graduate Council

Date _____
Peter Weber, Dean of the Graduate School

Vita

Alexei Sisulu Abrahams was born on August 14th, 1986, in Sherbrook, Québec, Canada. He graduated with a B.S. in Economics (Honors, Summa Cum Laude) and a B.S. in Mathematics (Distinction, Summa Cum Laude) from Syracuse University in Syracuse, USA. At the time of graduation, his GPA was 3.94 and he was awarded the Archimedes Prize as the top graduating Mathematics student. After taking graduate courses and providing teaching assistance in Mathematics at Syracuse during the 2008-2009 academic year, he began his doctoral studies in Economics at Brown University in Providence, Rhode Island, in the Fall of 2009. En route to his doctorate, he earned his M.A. in Economics in 2010. In 2011 he received funding from Brown's Population Studies & Training Center (PSTC) and Graduate School to attend Columbia University's Arabic summer immersion program in Jordan. Thereafter the PSTC awarded him NIH fellowships for the 2011-2012 and 2012-2013 academic years, allowing him to live abroad in the Occupied Palestinian Territories of the West Bank as a guest scholar at Bir Zeit University. Some of his travels were also funded by a Watson Institute fellowship in Spring of 2013. During the academic years 2013-2015 he worked as a short-term consultant to the World Bank Group in Washington, D.C. He completed his Ph.D. in 2015.

Acknowledgements

I thank my advisor, Nathaniel Baum-Snow, for remaining patient and even-keeled throughout this long and winding process, and for advising and supporting me despite large differences in our research interests. And I thank Andrew Foster for taking a chance on me, providing both moral and financial support for my decision to learn Arabic and to live in the Middle East.

I thank David Weil for his excellent comments and enthusiasm for my nighttime lights work, and I am grateful to Sriniketh Nagavarapu for being so generous with his time and advice.

Without the support and patience of Rita Giacaman and Majdoleen Jibril, I would not have gained access to the Palestinian census data integral to my dissertation's first chapter. I thank them, the ladies at ICPH, and the staff at the PCBS for their humor and professionalism.

I thank Nancy Lozano-Gracia for all our conversations on nighttime lights, and for sedulously guarding my time against other tasks – we could not have written my second chapter otherwise. I am likewise grateful to Chris Oram for his tireless enthusiasm for our work, and for teaching me the value of well-drawn diagrams.

I am very grateful to Roy van der Weide, Bob Rijkers, and Brian Blankespoor, who generously invited me to join their project (my third chapter) as a co-author. I thank them for freely sharing their advice and data, and I hope to emulate their collaborative research habits.

The doctoral experience was often emotionally trying, so I am lucky to have made so many

good friends along the way – too many to name here, but among them there was Sanjay, who stayed late researching with me at the PSTC so many evenings, and who taught me how to ride a bicycle; Shiva, who taught me to be a gentleman, and always saw the best in me; Philipp and his diligence and thoroughness; Alexander and his heart for needy people; Yana, who laughed at my jokes and always had a clever repartee; Morgan, who encouraged me to write songs; one girl, somewhere out there, about whom I suppose all the songs were ultimately written; and finally Ana Karine, Jake, and Jiries, about whom I shall simply say that the experiences of the last half-decade, and the person I have become over that time, are not imaginable in their absence.

Finally, I thank my parents, whose lifelong passion for social justice and solidarity with politically oppressed peoples led me to embark on this journey in the first place. Thank you both for leading by example, and for believing in me throughout.

‘What is now proved was once only imagined.’ - William Blake

Preface

This dissertation lies at the intersection of two broad questions: what are the short-run economic consequences of travel costs in urban and peri-urban environments of the developing world? And in what manner can satellite imagery be used to support urban research in developing and conflicted parts of the world? Chapters 1 and 3 focus on Israeli army road obstacles inside of the Occupied Palestinian Territories of the West Bank, exploiting their deployment during the 2002-2012 period as an arguably exogenous, quasi-experimental positive shock to travel costs that provides a chance to unbiasedly estimate the short-run economic consequences of travel costs. Both chapters draw on nighttime lights satellite imagery (NTL), and along the way certain systematic errors with the data are discovered. Chapter 2 explains and ameliorates a pervasive spatially autocorrelated error in these data, increasing their usefulness globally.

In Chapter 1, I model and test the idea that commuting costs in a developing urban environment redistribute economic welfare, obstructing some laborers from jobs but, in so doing, leaving vacancies open for others to seize opportunistically. I focus on the West Bank during the Second Palestinian Uprising (2000-2007), when Israeli army road obstacles were deployed along the West Bank's internal road network in an effort to defend Israeli civilian settlements from attacks by Palestinian militants. The obstacles, which took the form of either manned checkpoints or unmanned boulders, earthmounds, or gates, had the unintended consequence of disrupting Palestinian civilian traffic. Census data from 1997 and 2007, pre- and post-dating the Uprising, indicate that 61% of employed Palestinians commuted to work at census locations different from where they lived. Palestinians therefore enjoyed an active spatial economy and were vulnerable to disruption of commuting patterns. Negotiating access to confidential census data disaggregated down to the 'census location' level, I am able to observe employment rates at 485 Palestinian West Bank towns before and 5 years into

the deployment of obstacles. Digitizing a series of UN maps showing the geolocations of obstacles, I quantify each location's obstruction from jobs by a function of the count of obstacles along the primary road to each other location, weighted by that location's count of firms. I then quantify each location's protection from competing laborers by a function of the count of obstacles along the primary road to each other location, weighted by that location's count of laborers. I write down a commuting model to predict that obstruction from jobs should reduce a location's employment rate, but protection from labor inflows should increase a location's employment rate. Using Israeli settlements' proximity to commuter routes as an instrument, I 2SLS-regress post-uprising employment rates on *obstruction* and *protection*, controlling for cross-border commuting to Israel. I find that, insofar as they obstructed laborers from reaching jobs, obstacles caused a 4.28 percentage-point decline in employment throughout the territory. But insofar as these same obstacles reduced labor inflows, they caused a 3.84 percentage-point rise in employment. In aggregate, therefore, the obstacles were responsible for a rather negligible -0.44 percentage-point reduction to employment. Employment changes were by no means uniform across space, however: while inflowing laborers from peripheral areas were shut out of the labor market, laborers dwelling in core commercial areas were largely net-beneficiaries of obstacle deployment, seizing job vacancies in the absence of their unlucky countrymen. Various robustness checks, including the use of nighttime lights as an alternative to firm counts, suggest this result is remarkably robust. The result is also supported by somewhat weaker results suggesting that peripheral locations lost population to core areas.

Chapter 3 is co-authored work with Roy van der Weide, Bob Rijkers, and Brian Blankespoor. Whereas Chapter 1 focused on the commuting channel and emphasized redistributive consequences of travel costs, Chapter 3 focuses on aggregate consequences and remains agnostic about the channel, speaking in broader terms of 'market access'. Moreover, Chapter 3 deals with the time period 2004-2012, where the key source of temporal variation in location-level

market accessibility derives from the standing down of many army road obstacles in 2009-2010. In that sense, Chapter 3 looks at the easing of access, whereas Chapter 1 studies its constriction. Chapter 3 takes nighttime lights data as the left-hand-side variable, using year-to-year growth in lights as a proxy for year-to-year growth in GDP (a la Henderson et al 2012). Since NTL data are available for every year 1992-2012, we can eliminate location fixed effects in a panel regression framework. Another advantage over Chapter 1 is that Chapter 3 draws on field interviews with the UN and PCBS to estimate travel times for all roads, and traversal times for all obstacles. These estimates allow us to calculate optimal routes and minimum travel times between Palestinian locations or to border crossings, and then *recalculate* these whenever obstacles are deployed. This provides a much more nuanced quantification of accessibility changes over time, allowing locations to adapt to obstacle deployment by re-routing their exported goods along secondary routes. Regression analysis suggests doubling a location's market access causes on average a 10% increase to that location's light (GDP) growth. The inclusion of conflict-related and international-market related covariates appears not to perturb this estimate by more than 1 standard deviation.

The use of nighttime lights in Chapters 1 and 3 reveals a systematic source of error: lights emitted at any location blur onto neighboring locations, causing an underestimate of economic activity in source areas and an overestimate of economic activity in receiving areas. The problem is particularly acute in contexts like the West Bank, where separate locations are geographically close together. Chapter 2, which is co-authored work with Nancy Lozano-Gracia and Christopher Oram, explores the physical origins of this problem, arguing that the satellite's optical system is to blame. Using metadata on the satellite's altitude and angles of view, we provide an estimate of the blurring function, then derive and code an inverse filtration algorithm to de-blur nighttime lights imagery. Results over sub-Saharan African and South Asian cities suggest that blurring is a massive source of error in urban economic statistics. For example, cities appear 9.5 times larger than they actually are, with 72% of

their light spilt beyond city boundaries. Such error necessitates application of our deblurring algorithm before these data can be credibly used to answer urban research questions.

Contents

1	Hard Traveling: Redistributive Effects of Commuting Costs in the Second Palestinian Uprising	1
1.1	Introduction	1
1.2	Context & Data	7
1.3	Model	22
1.4	Identification Strategy	31
1.5	Regression Analysis	39
1.6	Conclusion	56
1.7	Bibliography	57
2	Correcting Overflow in Nighttime Lights Data	61
2.1	Introduction	61
2.2	Physical origins of overflow	66
2.3	Modeling and removing overflow	77
2.4	Empirical Results	84
2.5	Conclusion	91
2.6	References	92
3	How Valuable is Market Access? Evidence from the West Bank	95
3.1	Introduction	95
3.2	Nighttime Lights Data	99
3.3	Accessibility data	103
3.4	Regression Analysis	110
3.5	Conclusion	113
3.6	Bibliography	114

List of Tables

1.1	Palestinian West Bank: summary of census data 1997, 2007	9
1.2	Do lights measure production or consumption?	15
1.3	First-stage regression results	40
1.4	Main results: effects of obstacles on 2007 employment rates	41
1.5	Main results over different samples	42
1.6	Main results, weighting obstruction by radiance	42
1.7	Results with PCBS/World Bank spatial clusters	53
1.8	Effect of obstacles on locations of firms and laborers	54
2.1	Ellipses at various displacements off nadir	71
2.2	Urban extent and light over MODIS polygons in 2000; raw vs. de-blurred images	85
2.3	Intensive vs. Extensive Growth, 2000-2010; raw vs. de-blurred images	87
3.1	Effects of accessibility on per-capita lights, 2004-2012; 100 largest towns . . .	111
3.2	Effects of accessibility on per-capita lights, 2004-2012; all locations, weighted	111

List of Figures

1.1	The West Bank is <i>small</i>	8
1.2	Blurred image, F16-2006	17
1.3	De-blurred image, F16-2006	17
1.4	Author's digitization of UN-OCHA map (2007)	19
1.5	Transference of employment: A's loss is B's gain	28
1.6	Spatial histograms of <i>ob_firms</i> and <i>protection</i>	34
1.7	Road segments where roadblocks and checkpoints are likely	35
1.8	Defensive buffer zones around Israeli settlements	36
1.9	Spatial histograms of <i>ob_firms_seg</i> and <i>protection_seg</i>	37
1.10	%-point change in employment due to obstacles (Table 1.4, Column 1) . . .	44
2.1	DMSP-OLS sensor data collection	66
2.2	Misattribution of light in consecutive scans	67
2.3	Misattribution of light in consecutive scans	70
2.4	Conversion of fine to coarse pixels	70
2.5	Trigonometry of off-nadir ground scanning areas	72
2.6	Overflow frequency distribution (pixel size .867km)	72
2.7	Overflow distribution for 100 units of light (pixel size .867km)	74
2.8	The internal function of a photomultiplier tube	75
2.9	Malfunction of a photomultiplier tube	76
2.10	Stylized linear city <i>without</i> overflow	79
2.11	Stylized linear city <i>with</i> overflow	80
2.12	Elliptical bands	82
2.13	Addis Ababa, Ethiopia; F15-2000 radiance-calibrated data	86
2.14	Addis Ababa, Ethiopia; F16-2010 radiance-calibrated data	86
2.15	Addis Ababa, Ethiopia; F15-2000 rad-cal and MODIS urban extent	87

2.16	Addis Ababa, Ethiopia; F15-2000 radiance-calibrated; <i>de-blurred</i>	88
2.17	Addis Ababa, Ethiopia; F16-2010 radiance-calibrated; <i>de-blurred</i>	89
3.1	Light over Palestinian West Bank locations, 1992-2012	100
3.2	West Bank F162006, blurred image	101
3.3	West Bank F162006, de-blurred image	102
3.4	West Bank roads and Israeli army obstacles	104
3.5	Nablus: before and after the lifting of the blockade	105
3.6	Travel times to governorates & border crossings	107
3.7	Ease of access to Palestinian governorate capitals	108
3.8	Ease of access to Israeli border crossings	109

Chapter 1

Hard Traveling: Redistributive Effects of Commuting Costs in the Second Palestinian Uprising

1.1 Introduction

How do commuting costs affect laborers' welfare in the short run? In the long run, reduced commuting costs allow a city's population to suburbanize,¹ but economists have little to say about the welfare consequences of suburbanization, which may allow households to inhabit larger properties but may also reduce knowledge spillovers and lead to urban encroachment on natural areas. But in the short run, while land use regulations, housing supply inelasticity, and various kinds of moving costs delay the spatial redistribution of firms and laborers, what are the welfare impacts of commuting costs? To date, empirics show an inverse relationship: a real estate subliteration finds that homeowners' property values rise when commuting costs decline between their neighborhood and commercial areas;² and the spa-

¹See Baum-Snow (2007), Baum-Snow et al (2014), Baum-Snow & Turner (2014)

²See Debrezion et al (2007) for a review of papers focused on metro rail expansions in the developed world. For the developing world, see Gonzalez-Navarro & Quintana-Domeque (2014).

tial mismatch literature (Kain (1968), Ihlanfeldt & Sjoquist (1998)) contends that in large American metropolitan areas, commuting costs harm the labor market outcomes of low-skill, minority, downtown residents by reducing their access to job vacancies in suburbs.³

Both of these literatures, however, have ignored the fact that like trade barriers, commuting costs have a protectionist quality that benefits a subpopulation of laborers even while harming others. A rise in commuting costs harms the welfare of some laborers by reducing their access to jobs, but their subsequent absence from the market leaves jobs vacant, creating local labor supply shortages, and putting upward pressure on wages. If, as in many developing cities, laborers are very substitutable and there are initially high levels of unemployment across the city, the absence of any laborer from work leaves open a job vacancy that can quickly be seized by another unemployed laborer. Under such circumstances, a municipality's decision to improve or expand commuter infrastructure will not be purely welfare-enhancing; some laborers will benefit, but others will suffer. Likewise, a municipality's decision to neglect or destroy commuter infrastructure will not be purely welfare-reducing; some laborers will benefit, while others will suffer. In the absence of Pareto effects, any evaluation of aggregate welfare consequences will depend on how each laborer's welfare is weighted by the social planner.⁴

In this paper I model commuting costs as being both obstructive and protective, then test for dual welfare effects empirically by focusing on a commuter economy in the developing world: the Occupied Palestinian Territories of the West Bank. I negotiate access to confidential versions of Palestinian population censuses (1997, 2007) and a firm census (2004) to plot the positions of firms and laborers across 485 census locations of the West Bank, inter-connected

³In the long run, these laborers can relocate from the inner city to the suburbs and avoid these commuting costs altogether, but while low-income suburban housing options remain limited and racial discrimination in suburban housing markets persists, commuting costs reduce job access and welfare.

⁴This has obvious political economy implications. In municipalities whose constituents stand to lose jobs to inflowing commuters, the municipal government will tend to weigh its constituents' welfare more heavily and may oppose improvements to commuter infrastructure.

by a road network. The West Bank is only 130km North-South and 50km East-West at its widest, roughly one quarter the size of New Jersey, and Palestinians' low overall participation in agriculture (6.8%) and high commuting rates (61%) means the territory resembles more the features of a sprawled city with multiple commercial centers, than a system of specialized cities trading goods. Consistent with the model's assumptions, West Bank laborers are highly substitutable, possessing low education levels (15% high school completion rate) and competing among each other for low-skill jobs focused around construction, retail, and small-time manufacturing. Unemployment rates in both 1997 and 2007 were also very high (19.1% and 14.9%), further contributing to the labor market's fluidity.

During a violent uprising (the Second Intifada, 2000-2007), the Israeli army built a 500km wall to block West Bank Palestinian militants from entering Israel, and deployed hundreds of checkpoints, roadblocks, and other types of obstacles along the West Bank's internal road network in an effort to deter or intercept militant traffic along roads leading to or passing near Israeli settlements. Obstacles remained deployed throughout 2002-2007, and are widely reported to have had the unintended consequence of raising commuting costs for Palestinian civilian travel inside of the West Bank (World Bank 2007, B'Tselem 2007). To derive an index of commuting cost changes at the census location level, I obtain UN road network data and calculate bilateral commuting paths between every pair of census locations using ArcGIS Network Analyst. I georeference and digitize a series of UN poster-maps from 2003-2007 geolocating obstacles at 11 different instants during the uprising. I quantify each location's total exposure to obstacles by a function of the count of obstacles along each of the location's bilateral paths to other locations or border crossings with Israel.

My identification strategy centers on the widely supported claim (World Bank 2007, B'Tselem 2007) that obstacles were deployed to defend Israeli settlements, without any particular reference to the economic conditions of Palestinian locations. If obstacles were deployed more

often around Palestinian census locations with, say, a proclivity for violent confrontation, then the obstacle exposure index would be an endogenous regressor. I instrument for obstacle exposure with settlement proximity to commuter routes. I obtain polygon data of all Israeli settlements lying inside the West Bank during the uprising, and I draw defensive buffer zones around each settlement to identify segments of the road network that lay in close proximity to settlements. If indeed the army deployed obstacles to defend settlements, then these road segments would have been the likeliest places. I show that indeed these segment lengths are excellent predictors of obstacle counts, so that if a Palestinian’s commuter route to work unluckily passed often within close proximity of Israeli settlements, that commuter faced a higher total increase to commuting costs.

The commuting model suggests a simple empirical architecture to tease apart the dual effects of commuting costs on post-uprising (2007) employment rates. For each location j , I count up the obstacles along j ’s path to k , weighting this count on the one hand by the share of total labor force dwelling at k (to quantify j ’s protection from k), and on the other hand by the share of total firms located at k in 2004 (to quantify j ’s obstruction from k). The summed products result in two treatment variables, *protection* and *obstruction*, quantifying the labor-protective and job-obstructive capacities of obstacles. For *obstruction*, I would preferably use pre-uprising shares of firms, since perhaps the shares of firms in 2004 are themselves responding to obstacle deployment during 2002-2004. In the absence of pre-uprising firm data, I present evidence that firms were neither opening nor closing in a manner responsive to obstacle deployment. As a further robustness check, I establish that nighttime light emissions as measured by DMSP⁵ satellites are a credible proxy for locations’ job counts, then impute pre-uprising job shares using pre-uprising nighttime lights emissions over Palestinian locations, adopting these as an alternative to the firm shares. I use radiance-calibrated nighttime lights imagery to avoid top-censoring of data in bright areas around Jerusalem,

⁵The United States Defense Meteorological Satellite Program <http://ngdc.noaa.gov/eog/>

and I correct for blurring (‘overflow’) by applying a novel de-blurring method developed in Abrahams et al (2014).

I perform 2SLS regressions to estimate the dual effects of obstacles on 2007 (post-uprising) employment rates of 485 census locations. Integrating over location labor force sizes, the net effect of obstacles was to reduce total employment of the economy by just .44 percentage points, with the protective effect (3.84 percentage points) importantly mitigating the obstructive effect (-4.28 percentage points). Indeed, projecting the fitted values, some 180 of 485 Palestinian census locations enjoyed net gains to employment as a result of obstacles. The correlation of net employment effects of obstacles and pre-uprising proximity to jobs is 20.1%, indicating that residents of locations that were initially closer to jobs tended to fare better when their more distant counterparts struggled to reach work.

Were those who lost jobs fundamentally different from those who gained them? I add control variables for educational attainment, age, acreage of household agricultural holdings, private car ownership, but find that my point-estimates of interest are virtually unchanged. Are my results driven by increased costs of inputs imported from international markets? I use firms’ industrial codes to calculate each census location’s composition of industrial sectors; controlling for these again has no effect on my point-estimates of interest. Most convincingly, controlling for locations’ pre-uprising employment rates likewise has no effect on the point-estimates of interest, suggesting my instruments for obstacle exposure truly generate quasi-random variation.

Why is commuting, rather than trade (of final goods), the relevant channel here? Disruption of trade generates dual economic effects, too, but it turns out these effects run *opposite* to my empirical results: in the final goods market, obstacles should protect local industry and raise local employment by making imports from out of town costlier; but obstacles should

harm local industry and reduce local employment by making local products less competitive in out-of-town markets. In other words, the trade channel predicts that remotification from firms will raise employment, but remotification from people (consumers) will reduce employment. But my regression analysis reveals precisely the opposite result: remotification from firms reduces employment, while remotification from people (laborers) raises employment. I conclude, therefore, that the commuting channel dominates the trade channel in this context.

The short-run consequences of commuting costs are not particularly interesting if the short run is really short, i.e. if laborers and firms can quickly relocate to avoid these costs. In the spatial mismatch literature, the short-run consequences of commuting costs for minority residents of American inner cities are prolonged by housing market discrimination in the suburbs, which prevents these laborers from easily relocating to suburban areas. In the developing world, there are likewise many deterrents to relocation, including the importance of local family networks. Topalova (2010), Bryan et al (2013), and Munshi & Rosenzweig (2013) all find persistent welfare gaps between urban and rural areas in India and Bangladesh, and advance various theories explaining factor immobility. I present evidence that the West Bank likewise exhibits considerable factor immobility. Only 1.6% of 2007 labor force census respondents reported moving between 2000-2007 for job-related reasons. I show that among the few laborers who did end up relocating, their relocation patterns do indeed respond to obstacle deployment, with obstruction from jobs tending to induce depopulation, while protected locations gain labor immigrants. Relatedly, I show that incidence of renting falls in obstructed locations and rises in protected locations, suggesting that some laborers may have adapted their commuting habits in response to obstacles, returning to their hometowns on a less frequent basis, while renting apartments near their places of work in the meantime. Owing to the small fraction of laborers and households that reported moving or changing their renting behavior, however, these results are statistically unstable. Regressions using

firm census and nighttime lights data further suggest very weak statistical evidence that firms redistributed across locations in response to obstacles.

In addition to making a novel contribution to the literature on welfare and commuting costs, this paper also contributes to the debate over the effects of Israeli movement restrictions on the West Bank Palestinian economy (Cali & Miaari (2014), van der Weide et al (2014), World Bank (2007)). Whereas research has so far tended to argue that movement restrictions were purely harmful to Palestinian economic outcomes (one notable exception is Angrist (1996)), this paper finds robust evidence that the effect of obstacles was to redistribute welfare across subpopulations of the West Bank, generating economic inequality that adopted a clear spatial pattern: core locations benefited while peripheral locations suffered. Over time, this inequality will tend to encourage the spatial redistribution of firms and laborers. My results suggest firms are not moving, but there is some evidence that laborers are relocating to be closer to firms, abandoning peripheral areas in favor of core cities like Ramallah or Nablus. Given that the borders of a Palestinian state have yet to be determined, the depopulation of peripheral areas would seem to invite annexation, whereafter it will become even harder to reach a sustainable agreement on the geography of a two-state solution.

1.2 Context & Data

The West Bank is surrounded by Israel on its north, west, and south sides, while sharing an eastern border with Jordan. It is very small: about 56km (34.8 miles) at its widest and about 133km (82.6 miles) at its lengthiest, with roughly 1/4 the area of New Jersey. As depicted in Figure 1.1, the Israeli capital and historical city of Jerusalem, located inside Israel near the mid-section of the West Bank, is not much more than 75km (46.6 miles) from any West Bank location. Without obstacles or restricted roads, a driver could reach Jerusalem in under an hour from any part of the West Bank. The West Bank is also within an easy commuting




Figure 1.1: The West Bank is *small*

distance of most Israeli cities. The Palestinian town of Qalqiliya, for example, is just 15km from the coastal Israeli city of Netanya. More than 2 million Palestinians live in the West Bank, as do more than 300,000 Israeli civilian settlers.⁶

The West Bank has been under Israeli military occupation since 1967. Palestinian aspirations to achieve national self-determination have remained unrealized for over 47 years. Frustration has boiled over twice into extended popular uprisings. The first was in 1987-1991 (the First Intifada), a largely peaceful uprising that prompted negotiations for a two-state solution. When the peace process failed, economic stagnation and political frustration resulted in a second popular uprising in 2000-2007 (The Second Intifada),⁷ which was accompanied by an escalation in violence as militant Palestinian elements used conventional and suicide

⁶Current population estimates are drawn from the World Bank report on Area C (2013).

⁷See the joint Israeli-Palestinian economic think tank AIX Group's report entitled 'Twenty Years after Oslo and the Paris Protocol' for a detailed and thorough discussion. <http://www.aixgroup.org/research>

attacks against Israeli military and civilian targets.

Labor force and employment

Censuses of Palestinian residents of the West Bank were conducted in 1997 and 2007, pre- and post-dating the Second Intifada. Table 1.1 lists aggregate descriptive statistics of the Palestinian West Bank. Depressed economic conditions in the late 1990s are evident in the low employment rate (80.9%) recorded in 1997. Formally, I define labor force as

Table 1.1: Palestinian West Bank: summary of census data 1997, 2007

	1997	2007
Population	1.5 million	2 million
Labor force	408,618	472,916
Employed	330,667	402,366
<i>As a percentage of labor force:</i>		
Employed	80.9%	85.1%
Male	86.0%	85.1%
Female	14.0%	14.9%
<i>As a percentage of those employed:</i>		
Commuting to Palestinian locations	74.0% ⁸	45.7%
Commuting to Israel	21.6%	13.0%
Commuting to settlements	2.7%	2.5%
<i>Major job sectors:</i>		
Construction	NA	23.2%
Retail	NA	10.9%
Education	NA	10.8%
Public sector	NA	10.1%
Farming	NA	6.8%
Transportation	NA	4.5%
Materials manufacture	NA	4.1%
Minerals	NA	3.2%
Health	NA	3.1%
Vehicle sale & repair	NA	3.1%
Hospitality	NA	2.3%
Wholesale	NA	2.2%
Furniture	NA	1.8%

$$labor_force = \#employed + \#unemployed_but_available_to_work$$

I therefore exclude those who report they are not available to work because they are instead studying, doing house chores, too old, disabled, or on pension. The correlation between

location labor forces and populations is over 99% in both 1997 and 2007, so labor force participation rates do not vary over locations in a meaningful way. I then define the employment rate as

$$employment_rate = \frac{\#employed}{labor_force}$$

Unemployment benefits

The West Bank has one of the world's highest rates of per-capita receipt of international aid, as a result of which living standards rarely reach such low levels as would be found in urban slums of sub-Saharan Africa or India. In 1997, 80% of West Bank Palestinians households owned houses while almost all of the remainder rented. 86% of households owned TVs and 81% owned fridges. Likewise, 94% were connected to a public electric grid, and were hooked up to either public (79%) or private (15%) water networks. 23% owned private vehicles. Impoverished families receive subsidies from the Palestinian Authority, itself significantly funded by international aid⁹; and various international organizations offer aid to the unemployed and impoverished more directly, such as food stamps available through the UN World Food Program.

Commuting

Table 1.1 indicates high rates of commuting inside the West Bank and into Israel in both 1997 and 2007. Even in 2007, with obstacles deployed throughout the West Bank, 61% of employed Palestinian laborers still reported working outside of their location of residence. Less than a quarter of these commuted across the border into Israel; the majority commuted

⁹See the AIX Group report 'Twenty Years after Oslo and the Paris Protocol', Section 2.3.

to other Palestinian locations inside the West Bank.

Internal trade

Alongside labor flows, there may have been internal flows of traded goods that were likewise disrupted by road obstacles, and one might be concerned that trade disruption, rather than commuting disruption, drives the paper's main results. In Section 3.4 I address this concern empirically, showing that point-estimates for the commuting story are unaffected by controlling for trade-related variables. Table 1.1 points to some historical reasons why that result is not surprising. As of 2007, despite low participation in agriculture (6.8%), most West Bank jobs were in the service sector, with only 10% in manufacturing. These statistics reflect a kind of 'false' structural transition (Gollin et al (2013)). Prior to 1967, the West Bank was chiefly an agrarian economy. When Israel captured the West Bank, however, low-skill urban job opportunities suddenly became accessible to West Bank Palestinians, who could commute daily to Israel and work below minimum wage. Palestinians drained out of agriculture and the West Bank became a resource export economy where unskilled labor, rather than cocoa or oil, was the primary export (Hilal (1977), Graham-Brown (1979)). Wages earned by Palestinian laborers in Israel were then spent on imports (often Israeli) and services in the West Bank, thus giving rise to a predominantly service-oriented economy without the industrial base characteristic of classical structural transformations. Israeli demand for unskilled Palestinian labor also disincentivizes human capital accumulation, thus retarding the growth of high-skill service or manufacturing (Galor and Mountford, 2008). Furthermore, reports by AIX Group document various measures taken by the Israeli authorities to thwart the growth of Palestinian tradeables. All of this contributes to the low share of manufactures (10%) in overall employment. The World Bank report on Area C (2013) likewise highlights

the lack of a substantial tradeables sector.

Low-skill jobs

Table 1.1 suggests West Bank Palestinians mostly occupy low-skill jobs. Indeed, the average educational attainment of 1997 laborers was just 9.24 years of school, with only a 15% high-school graduation rate. By definition, low-skill jobs require less training and laborers are more substitutable, which supports the narrative below.

Rising employment rates with proximity to jobs

I generate location-level statistics from the population (1997 and 2007) and firm (2004) censuses, and map them using UN-OCHA polygon and centroid point data of census locations. I am able to identify 485 Palestinian locations (accounting for 99.5% of the West Bank 1997 labor force) for which census data are available.¹⁰ Using the 2004 firm census I count the number of employees per location k (*not* the number of residents of k who happen to have a job, but rather the number of laborers, hailing from any part of the West Bank, whose workplace is k). Calculating the minimum road distance between each pair of locations with ArcMap Network Analyst, I generate a location-level index of proximity to jobs for 2004 (ignoring obstacles):

$$job_proximity_j = \sum_{k=1}^J \frac{employees_k}{\tau_{jk}}$$

Where τ_{jk} is the minimum road-distance in kilometers between locations j and k (without

¹⁰See online appendix for details.

road obstacles).

I define location j 's employment rate as

$$employment_rate_j = \frac{\#employed_j}{labor_force_j}$$

And I define for each location j an ‘internal’ employment rate, measuring the fraction of laborers dwelling at j who are employed in a Palestinian location:

$$internal_employment_rate_j = \frac{\#employed_at_palestinian_location_j}{labor_force_j}$$

I regress ‘internal’ 1997 (pre-uprising) employment rates on job proximity, supposing for a moment that the spatial distribution of jobs in 2004 was identical to that of 1997:

$$internal_employment_rate_{j,1997} = \beta_0 + \beta_1 job_proximity_j + \epsilon_{j,1997}$$

Statistics such as $internal_employment_rate_{j,1997}$ are measured with less accuracy when location j 's labor force is small. I therefore weight observations by the size of their 1997 labor force.¹¹ Using robust standard errors, I find $\beta_1 > 0$ and statistically significant, and $R^2 = 34\%$, meaning proximity to jobs explains 1/3 of the variation in 1997 ‘internal’ employment. locations tended to enjoy higher employment rates when closer to jobs.

Immobility of firms and laborers

The job proximity measure is compiled using filled job counts from 2004. By regressing 1997 employment rates on this measure, I am implicitly assuming that the spatial distribution of jobs was unchanged between 1997 and 2004. In Section 3.4 I use firm census data to present

¹¹In Stata, this is done by running a regression using the a-weight option.

evidence that firms neither opened or nor closed in a manner correlated with obstacle deployment. Without panel data on employee counts, however, I cannot say if they redistributed during the uprising. To address this issue I generate a panel dataset of nighttime lights over Palestinian locations, show that lights closely track production activity, and take lights as a proxy for employee counts pre- and post-uprising years. I then show that proximity of locations to lights changes minimally over the course of the uprising, indicating that jobs did not spatially redistribute.

Nighttime lights satellite imagery are well established as a correlate of economic activity.¹² Many recent studies¹³ rely on lights as a proxy for GDP and GDP growth in the absence of more traditional data sources. Consistent with findings in other regions of the world, Levin & Duke (2012) shows that lights emitted over Israel and the West Bank correspond to economic realities on the ground, and reflect economic inequalities between Israeli and Palestinian societies. In co-authored work with van der Weide et al (2014), we use nighttime lights to track relocation of economic activity in the West Bank after the removal of road obstacles in 2009. For a thorough description of nighttime lights, see Abrahams et al (2014).

Do lights over Palestinian locations measure production or consumption? If the workday is over when the satellite passes by (typically 730-830pm local time), businesses will be closed and their lights may be turned off. In that case, lights are more likely to correlate with income of local residents, as they consume electricity at their homes. If, however, laborers are still at work, then lights should correlate with production activity. To isolate these two potential sources I calculate light emissions recorded over polygon extents of Palestinian

¹²See Doll et al 2006, Chen and Nordhaus 2011, Henderson et al 2012,

¹³Storeygard (2014), Baum-Snow & Turner (2014), Alesina et al (2014).

locations and regress lights on residential labor force counts and employee counts:

$$radiance_{j,t} = \beta_0 + \beta_1 employees_{j,t} + \beta_2 labor_force_{j,t} + \epsilon_{j,t}$$

Where t corresponds to the period 2004-2007.¹⁴ Table 1.2's results indicate that employee

Table 1.2: Do lights measure production or consumption?

	(1)	(2)
<i>employees</i>	0.29 (0.025)	0.207 (0.033)
<i>labor_force</i>		0.05 (0.019)
<i>constant</i>	36.445 (8.668)	15.855 (10.194)
Observations	485	485
R^2	0.8265	0.8351

counts per location strongly correlate with nighttime lights. Interpreting the regression coefficients in Column 2, each additional employee at location j increases j 's total radiance by 0.21 units, while each additional laborer dwelling at location j only increases radiance by .05, a magnitude over four times smaller. Comparing Columns 1 and 2, we see furthermore that employee counts on their own explain 82.65% of the variation in lights, and introducing labor force to the regression raises R^2 by less than 1%.¹⁵ Lights net of labor forces are simply

$$radiance_net_of_labor_{j,t} = radiance_{j,t} - \hat{\beta}_2 labor_force_{j,t} - \widehat{constant}$$

In the absence of pre-uprising employee counts, I take 1997 census labor force counts and 1996 radiance-calibrated nighttime lights¹⁶ and generate pre-uprising radiance net of labor.

¹⁴Precisely, labor force counts are calculated from the 2007 population census, employee counts are drawn from the 2004 firm census, and radiance-calibrated lights are from 2006. Ideally I would have population, jobs, and lights data all from the same year. The anachronisms of the regression are unproblematic as long as locations' populations in 2004 were proportional to their populations in 2007, and as long as lights in 2006 were proportional to lights in 2004.

¹⁵I cannot credibly add 2007 population to the RHS of the regression in Column 2 since 2007 labor force and 2007 population correlate over 99%, meaning these variables are effectively collinear.

¹⁶Radiance-calibrated lights data are available only for years 1996, 2000, 2006, and 2010, whereas stable

These residuals correlate 98% with 1996 radiance. I apply them as a proxy for pre-uprising employee counts:

$$job_proximity_j = \sum_{k=1}^J \frac{radiance_net_of_labor_1996_k}{tau_{jk}}$$

Regressing $internal_employment_rate_{j,1997}$ on this radiance-weighted job-proximity index, I find $\beta_1 > 0$ and statistically significant, with $R^2 = 13.5\%$. This confirms that the positive relationship between proximity to jobs and 1997 employment rates was not an artefact of using 2004 employee counts; the result persists even when using pre-uprising, radiance-imputed employee counts.

While using radiance-calibrated imagery avoids the problem of topcoding, there is a further source of error prevalent in all nighttime lights data known as ‘overflow’ (Croft (1979), Small, Pozzi, Elvidge (2005), Pinkovski (2011)). Essentially, the imagery suffer from significant blurring, with light emitted at one location spilling onto nearby locations. This is particularly problematic in the West Bank since locations lie very close together: the average minimum as-the-crow-flies distance between Palestinian locations is just 1.81km, and between Palestinian locations and Israeli settlements is just 2.12km.

This paper, along with van der Weide et al (2014), is the first to apply the novel overflow correction method of Abrahams et al (2014). In that paper, the authors study the DMSP satellites’ optics and derive a nonparametric representation of the elliptical distortion, then provide a deconvolution algorithm by which nighttime lights imagery can generally be de-blurred. I apply that method here to all radiance-calibrated imagery of the West Bank.

lights data are available for every year 1992-2012. The major advantage to using radiance-calibrated data is that they are not top-censored (see Elvidge et al (1998) or Doll (2008)). Stable lights imagery over West Bank locations in the vicinity of Jerusalem suffer top-censoring, which is problematic since Ramallah and associated locations lie within this area and are known to be important job destinations for Palestinians throughout this time period.




Figure 1.2: Blurred image, F16-2006




Figure 1.3: De-blurred image, F16-2006

Figures 3.2 and 3.3 provide a comparison of the 2006 image before and after de-blurring.

I calculate job proximity before and after the uprising using radiance net of labor from 1996 and 2006. Though ten years apart, the correlation between the two measures is 97.6%. This means that, *ignoring obstacles*, the proximity of locations to jobs changed minimally during the uprising. Meanwhile, the 2007 census indicates just 7,774 respondents (1.6% of the 2007 labor force) moved during 2000-2007 for job-related reasons. The correlation between 1997 and 2007 location labor forces is over 99%; between populations it is 98%. The evidence therefore indicates that neither laborers nor firms substantially redistributed over space, despite the fact that changes in commuting costs due to Israeli road obstacles likely generated inequality in utility levels. This apparent immobility of laborers and firms despite changes to real income is not unique to the West Bank: Topalova (2010), Bryan et al (2013) and Munshi & Rosenzweig (2013) find similarly puzzling behavior in India and Bangladesh, where migration rates are low despite real rural-urban gaps in living standards. A new report on the Egyptian labor force (World Bank, 2014) likewise documents low migration rates (and high commuting) despite apparent rural-urban inequality. Munshi & Rosenzweig (2013) argues that the value of local networks makes would-be migrants reluctant to leave. Cammett (2014) shows that local networks in Lebanon tend to replace most of the functions that, in developed countries, would be handled by the state. The same is likely true in the West Bank, perhaps especially given the unresolved political situation and ongoing military occupation.

Furthermore, relocating had to be weighed against expectations of how long road obstacles would remain deployed. As we shall now see, there was always good reason to believe ob-




Figure 1.4: Author's digitization of UN-OCHA map (2007)

stacles would be removed as soon as the uprising ended.

Road obstacles

As violence escalated during the Second Intifada, the Israeli army responded to militant activity with both offensive retaliations (Paserman 2008) and defensive efforts aimed at intercepting militants before they could reach Israeli civilian destinations. This latter policy was known as 'Operation Defensive Shield', and involved the deployment of numerous physical obstacles. A 500km wall was built to separate Israel from the West Bank, but most Israeli settlers dwelt beyond the wall and could not rely upon it for protection. Instead, the army deployed hundreds of obstacles *inside* the West Bank, along the internal road network, in order to intercept militant traffic before it could approach Israeli settlements. Quoting

B'Tselem:

Israel's primary justification for the movement restrictions is that they are necessary to protect Israelis within its jurisdiction and Israelis living in the West Bank or traveling on West Bank roads.

...the settlement enterprise, including the roads built for it, was one of the primary factors in shaping the restrictions regime that Israel has forced on the Palestinians since the beginning of the Second Intifada.

(B'Tselem, Ground to a Halt, 2007)

Consistent with their defensive purpose, obstacles did not forbid commuting between locations, but rather introduced significant delays as vehicles were checked for weapons:

I saw the soldiers were carefully checking a Palestinian taxi. They dismantled the seats, the door panel and many other parts of the taxi. They spent a lot of time inspecting it...

(B'Tselem, Ground to a Halt, 2007)

As of September 2003, UN-OCHA¹⁷ began publishing poster-sized maps depicting the precise locations of obstacles along the West Bank's internal roads. Through to the end of 2007 a total of 11 posters were published.¹⁸ I georeference all 11 poster-map pdfs and digitize all obstacles (6,180 obstacles plus the separation barrier), superimposing the new data over the poster imagery. Figure 1.4, for example, displays the digitized December 2007 map, where

¹⁷The United Nations Office of the Coordination of Humanitarian Affairs

¹⁸The entire time series of posters is available in pdf format from UN-OCHAs website: <http://www.ochaopt.org/>

more than 500 obstacles were deployed along with 500km of wall.

Not all obstacles were the same. While checkpoints were manned by soldiers and could not be passed without inspection, the UN also identifies ‘partial checkpoints’, which had all the appearance of checkpoints yet rarely interfered with passing traffic. Unmanned obstacles included ‘roadblocks’, usually giant boulders set in the middle of roads by army bulldozers; or ‘earthmounds’, described as mounds of dirt dumped in the middle of roads by bulldozers. UN-OCHA reports that earthmounds were a particularly ineffective type of obstacle, since Palestinian traffic usually circumvented them with ease, or drove over them repeatedly until they flattened out. Although results are similar with their inclusion, I exclude earthmounds, partial checkpoints, and incomplete sections of the separation wall from the regression analysis of Section 3.4 in order to reduce attenuation bias.

Not all obstacles stayed in the same place. Indeed, at any given time some 40% of checkpoints were temporary, ‘flying’ checkpoints (World Bank (2007)), set up suddenly along road segments with the intention of surprising militants. While UN-OCHA faithfully recorded obstacle locations, their maps amount to a series of snapshots that capture both permanent and temporary obstacles. As such, when I calculate obstacle exposure indices I introduce a lot of attenuation bias by supposing that commuters faced the same obstacle for many months, when in fact it may as often as not have been deployed only for a few days. As will become clear in Section 3.4, my instrumental variables approach greatly mitigates attenuation bias by isolating the subset of obstacles deployed in the immediate vicinity of settlements. Settlements required perpetual defending, so obstacles in the vicinity of settlements were most likely to have been enforced with permanent vigilance.

1.3 Model

I develop a commuting model incorporating the stylized facts reviewed in Section 1.2: many locations, unemployment at all locations, more unemployment farther away from jobs, unemployment benefits, commuting, and substitutability of laborers. The purpose of the model is to predict how obstacles affect employment rates of Palestinian locations in the short run, when neither firms nor laborers can relocate, and to relate short-run employment rate changes to short-run welfare changes. The model draws on the modeling framework of Ahlfeldt et al (2014), but assumes laborers are homogeneously productive, and therefore perfectly substitutable. I introduce unemployment benefits coupled with randomness in job search in order to generate unemployment in equilibrium. While some models of spatial mismatch speculate that differential search costs or effort are to blame for employment rates declining with residential distance from jobs (Wasmer & Zenou (2002), Smith & Zenou (2003)), my model remains agnostic. In particular, I assume that laborers dwelling far from firms search for jobs with just as much success as laborers dwelling close. But when it comes time to accept or decline the job offer, laborers dwelling farther away are more likely to decline, since their commuting costs effectively raise their reservation threshold. Accordingly, firms set wage offers to maximize expected profit, uncertain about the reservation utilities of their prospective laborers. In equilibrium, laborers choosing to dwell near to firms enjoy higher probability of employment, while laborers dwelling far from firms are unemployed more often but pay lower land prices. When commuting costs subsequently rise, laborers dwelling far from firms suffer a decline in earnings and employment as commuting becomes costlier. The reduced supply of labor to firms induces a rise in wages, from which laborers dwelling nearby are able to benefit disproportionately, enjoying increases to earnings and employment. A full derivation is available in the online appendix. Here I only present main results and intuition.

I posit an economy of N homogeneous laborers with identical preferences for a composite

consumption good x

$$U = x \tag{1.1}$$

And an endogenous number of profit-maximizing firms M with identical production functions

$$y = \theta l \tag{1.2}$$

where the productivity of the quantity of laborers l employed by a firm is multiplied by exogenous technology θ .¹⁹

Utility maximization and profit maximization occur in two stages. In the first stage, laborers and firms decide where to locate. There exist J census locations, each endowed with a different supply of land L_j available for commercial and/or residential use, where α_j denotes the endogenous fraction of location j 's land used for commercial activity. Notably, land enters neither into laborers' utility functions nor firms' production functions, so laborers and firms do not value land intrinsically.²⁰ Rather, as shown below, firms (laborers) value locations only by their proximity to laborers (firms), and pay absentee landlords for the opportunity to locate there. Laborers dwelling in j pay $p_{j,r}$, while firms pay $p_{j,f}$.²¹ The prices of land faced by firms and laborers in each location j can alternatively be thought of as entry fees: a laborer or firm must purchase exactly 1 unit of land in j in order to locate

¹⁹The constant returns to scale of labor assumed here is in agreement with Henderson, Storeygard, Deichmann (2014), where the urban services sector has an identical production function.

²⁰In general, consumers may enjoy renting more spacious properties, or productivity of a firm's employees may be declining in the quantity of land rented. Without data on land parcels or rental rates, however, any modeling decision in this context is speculative. Moreover, this model demonstrates that these other ways of valuing land are not necessary for generating predictions about the effects of commuting costs on welfare.

²¹Ahlfeldt et al (2014) insists on a single, prevailing price of land p_j for each location, in order to match the Berlin dataset. To achieve spatial equilibrium, the Berlin model introduces unobserved location-specific amenities B_j . The price p_j is set so as to make firms indifferent over location choices, while the amenities B_j , valued only by residents, are used to make laborers indifferent over location choices. As discussed in that paper's technical appendix, in general there may be a wedge between rental rates of residents and firms. For example, the government may tax or subsidize commercial land usage in ways different from residential usage. In the depressed Old City of Hebron, for example, the Palestinian Authority is known to pay shopkeepers to stay open, while in Hebron's busy downtown areas the author has interviewed business owners and municipal workers claiming high rental rates for office space.

there.²²

In the second stage, firms post wage offers.²³ Laborers are each endowed with 1 unit of labor and are homogeneously productive, producing θ units of the composite consumption good. Laborers undertake job searches and each laborer discovers exactly one firm,²⁴ drawn at random *and with replacement* from among the economy's M firms.²⁵ Observing wage offers, laborers decide whether to accept or decline jobs. For a laborer dwelling at location j , if he accepts a job offer at location k , he produces θ units for the firm and is compensated with wage w_k , from which a commuting cost d_{jk} is subtracted. If he declines the offer, he receives a 'reservation wage' (unemployment benefits) worth $\bar{w}$, drawn from $U[0, 1]$. The laborer's decision rule is simple:

$$\text{accept } w_k \Leftrightarrow \bar{w} < w_k - d_{jk} \quad (1.3)$$

Firms located in k anticipate 1.3 and set wage offer w_k to maximize expected profit:

$$\max_{w_k} E(\pi) = (\theta - w_k) \frac{N}{M} (w_k - \sum_{j=1}^J d_{jk} n_j) - p_{k,f} \quad (1.4)$$

The difference $(\theta - w_k)$ quantifies the surplus enjoyed by the firm, i.e. the difference between the value of what the laborer produced and his compensation. A lower wage offer w_k allows the firm to extract greater surplus, but also reduces the probability $(w_j - \sum_{j=1}^J d_{jk} n_j)$ that a (random) laborer will accept the offer. The firm's maximization problem is to balance this

²²This feature of the model is for the sake of parsimony, so as not to distract from the model's purpose. The data contain no information on sizes of properties or heterogeneity in preferences for real estate, so any attempt to model those features would be speculative.

²³This is a departure from Ahlfeldt et al (2014), where laborers discover firm-laborer match productivity shocks and then decide which firm to work with. In that case, there are $N \cdot J$ idiosyncratic shocks (one for each firm-laborer match). I avoid this by introducing imperfect information.

²⁴Trivially, the model can be altered to allow laborers to discover 1 firm with probability < 1 .

²⁵The job search therefore is not a function of commuting costs, i.e. within this model, laborers do not search nearby locations with more success than distant ones. This is a modeling decision for the sake of parsimony: the model should predict that expected wages are higher among laborers near firms, and as I show, this prediction arises even when success in finding jobs is not a function of commuting costs to where jobs are offered.

tradeoff optimally. The FOC of 1.4 gives

$$w_k^* = \frac{\theta}{2} + \frac{\sum_{j=1}^J d_{jk} n_j}{2} \quad (1.5)$$

This shows that profit-maximizing wage offers are an increasing function of average commuting cost to labor. Holding firm and labor location choices fixed, the implication is that wages will rise with commuting costs. This is the distinguishing feature of the model, and the reason why commuting costs have dual welfare consequences. The firm's expected profits are therefore

$$E(\pi(w_j^*)) = \frac{N}{M} \left(\frac{\theta}{2} - \frac{\sum_{k=1}^J d_{jk} n_k}{2} \right)^2 - p_{j,f} \quad (1.6)$$

Laborers anticipate that firms located in k will make a wage offer of w_k^* . Therefore in the first stage, when each laborer's reservation wage is not yet revealed, the laborer calculates his expected utility conditional on dwelling in each location, and chooses where to live. Expected utility of a laborer dwelling in j turns out to be

$$E[U_j] = \frac{1}{2} + \frac{1}{2} \sum_{k=1}^J m_k (w_k^* - d_{jk})^2 - p_{j,r} \quad (1.7)$$

This says simply that a laborer dwelling in j can expect to earn $\frac{1}{2} + \frac{1}{2} \sum_{k=1}^J m_k (w_k^* - d_{jk})^2$, of which a quantity p_j is paid to an absentee landlord, while the remainder is spent on consumption good x .

Equilibrium

Since laborers and firms are free to locate anywhere, expected utility and profits must be

equalized across all locations in equilibrium:²⁶

$$\frac{1}{2} + \frac{1}{2} \sum_{k=1}^J m_k (w_k^* - d_{jk})^2 - p_{j,r} = \bar{U}, \text{ for each } j = 1, \dots, J \quad (1.8)$$

$$\frac{N}{M} \left(\frac{\theta}{2} - \frac{\sum_{k=1}^J d_{jk} n_k}{2} \right)^2 - p_{j,f} = \bar{\pi} \text{ for each } j = 1, \dots, J \quad (1.9)$$

Expressions 1.8 and 1.9 determine equilibrium land prices $p_{j,r}$ and $p_{j,f}$ for laborers and firms in all j . Firms enjoy free entry and exit. As long as expected profits are positive, firms will continue to enter the economy. With each additional entry, laborers are spread more thinly among firms ($\frac{N}{M}$ declines) and expected profits fall. In equilibrium, $\bar{\pi} = 0$, and M is determinate.

In equilibrium, demand for land must equal supply. Since each firm and each laborer uses 1 unit of land, the demand for commercial land in j is just M_j , and the demand for residential land in j is just N_j . We have

$$n_j = (1 - \alpha_j) \frac{L_j}{N}, \text{ for each } j = 1, \dots, J \quad (1.10)$$

$$m_j = \alpha_j \frac{L_j}{M}, \text{ for each } j = 1, \dots, J \quad (1.11)$$

The system of equations in 1.8, 1.9, 1.10 and 1.11, along with the condition $\bar{\pi} = 0$, characterizes equilibrium. See the online appendix for a simple sufficient condition under which a

²⁶In the data, some locations to which commuting labor is flowing are Israeli, and Palestinians are forbidden to locate there. The model easily accommodates this: the share of the Israeli labor force and Israeli firms located at the Israeli location are taken as exogenous to the model. In equilibrium, the price of land in Israeli locations need not equalize utility of its residents to utility of residents in Palestinian locations; nor need profits of Israeli firms equalize with profits of Palestinian firms, since there is no freedom to relocate to Israeli locations. In other words, equilibrium conditions for the West Bank economy do not pin down any characteristics of Israeli locations.

long-run equilibrium always exists.

Employment

In order for a laborer dwelling at location j to be employed at a job in location k , two things must happen: the laborer must be randomly matched with a firm at k ; and the firm must offer a wage high enough to induce the laborer's acceptance. The probability that this coincidence occurs is

$$Prob(\text{laborer dwelling in } j \text{ is employed in } k) = m_k \left(\left(\frac{\theta}{2} + \frac{\sum_{i=1}^J d_{ik} n_i}{2} \right) - d_{jk} \right) \quad (1.12)$$

Expression 1.12 is the bilateral commuting flow $outflow_{jk}$ of labor from j to k . It follows that the probability that a laborer dwelling in j will be employed at all is just

$$Prob(\text{laborer dwelling in } j \text{ is employed}) = \sum_{k=1}^J m_k \left(\left(\frac{\theta}{2} + \frac{\sum_{i=1}^J d_{ik} n_i}{2} \right) - d_{jk} \right) \quad (1.13)$$

Expression 1.13 is the employment rate of location j , and is observable in the census data.

Dual effects of obstacles on employment

How do commuting costs affect employment rates in the short run? I define the short run as a period of time when neither firms nor labor are mobile (m_k and n_k are fixed for all k).²⁷

Figure 1.5 helps with intuition. Suppose laborers are commuting from Origin A and Origin B to work at Destination. When a checkpoint is deployed along the road between A and Destination, laborers from A face higher commuting costs to reach work, causing some to give up their jobs (see 1.12). Firms at Destination respond to diminished labor supply by raising wages (see 1.5). When wages rise, some unemployed laborers from B are now willing

²⁷In the long run, changes in commuting costs d_{jk} will cause laborers (firms) to relocate to maximize expected utility (profits). The increased (decreased) demand for land in each location will cause land prices p_j to adjust.




Figure 1.5: Transference of employment: A's loss is B's gain

to accept jobs at Destination (see 1.12). The effect of the checkpoint's deployment, therefore, should be to decrease employment of A 's laborers but *increase* employment of B 's laborers. In short, A 's loss is B 's gain.

To show this formally, suppose laborers commuting from j to work at k face a commuting cost of d_{jk} . Differentiating 1.13 with respect to d_{jk} , I obtain

$$\frac{\partial \text{Prob}(\text{employed})}{\partial d_{jk}} = m_k \left(\frac{n_j}{2} - 1 \right) < 0 \quad (1.14)$$

This says that an increase to commuting costs faced by laborers residing in j can only harm their chances of employment. The cross partials are

$$\frac{\partial \text{Prob}(\text{employed})}{\partial d_{jk} \partial m_k} = \left(\frac{n_j}{2} - 1 \right) < 0 \quad (1.15)$$

$$\frac{\partial \text{Prob}(\text{employed})}{\partial d_{jk} \partial n_j} = \frac{m_k}{2} > 0 \quad (1.16)$$

According to 1.15, an increase to m_k exacerbates the commuting-cost-based employment losses in j . The larger is m_k , the more important it is as a destination for laborers from j , in which case an increase to d_{jk} is more harmful to their employment chances. Holding m_k

fixed, however, 1.16 says that an increase to the mass of laborers dwelling at j mitigates the employment losses caused by an increase to d_{jk} . The larger is n_j , the more important it is as a source of labor to firms in k , so that a rise in commuting cost d_{jk} will induce a larger increase in wage w_k^* , thus partially offsetting losses.

What happens to the employment rate in j when commuting costs rise between i and k ? Partially differentiating, I obtain

$$\frac{\partial \text{Prob}(\text{employed})}{\partial d_{ik}} = m_k \frac{n_i}{2} > 0 \quad (1.17)$$

This says that an increase to commuting costs faced by laborers residing in i on their commute to work at k will *increase* the employment rate at location j . Together, 1.14 and 1.17 formalize the simple intuition of Figure 1.5, where town A 's losses were converted into town B 's gains. Calculating cross partials, I obtain

$$\frac{\partial \text{Prob}(\text{employed})}{\partial d_{ik} \partial m_k} = \frac{n_j}{2} > 0 \quad (1.18)$$

$$\frac{\partial \text{Prob}(\text{employed})}{\partial d_{ik} \partial n_i} = \frac{m_k}{2} > 0 \quad (1.19)$$

Evidently the positive effect on employment in 1.17 is amplified by k 's importance as a destination for j . In particular, 1.18 says that for a larger mass of firms at k , laborers at j enjoy an even greater increase to their employment chances when laborers from i lose access. The positive effect is likewise amplified if i is a large source of labor to k : when commuting cost d_{ik} rises, wage offer w_k^* will rise significantly, benefiting laborers from j even more (1.19).

Checkpoints and roadblocks along the road between j and k delay laborers originating from j and laborers originating from k . The marginal effect on employment rates of obstacle

deployment along bilateral linkages is therefore the sum of 1.14 and 1.17:

$$\Delta employment_rate_j = m_k(\frac{n_j}{2} - 1) + m_j \frac{n_k}{2} \quad (1.20)$$

The first term is negative, and identifies the loss to j 's employment rates due to remotification from jobs at k ; while the second term is positive, identifying the gain to j 's employment rates due to remotification from labor at k . If the econometrician only measures each location's total obstacle 'exposure', she will end up estimating 1.20, without being able to distinguish the dual effects 1.14 and 1.17. The key contribution of my regression analysis below is to identify separately the dual effects.

Dual effects of obstacles on welfare

In equilibrium, utility of laborers is equalized across all locations (see 1.8). When commuting costs change, the economy is thrown out of equilibrium. In the short-run, m_j and n_j do not change, i.e. neither firms nor laborers relocate, so there is no increased (decreased) demand for L_j for any j , so absentee landlords do not adjust prices $p_{j,f}$ or $p_{j,r}$. Therefore we can calculate the effect of commuting costs on expected utility of residents of location j (1.7) by its partial derivative w.r.t. d_{jk} :

$$\frac{\partial E[U_j]}{\partial d_{jk}} = m_k(w_k^* - d_{jk})m_k(\frac{n_j}{2} - 1) = outflow_{jk} \frac{\partial Prob(employed)}{\partial d_{jk}} < 0 \quad (1.21)$$

Short-run welfare of location j therefore declines as commuting costs to work destination k rise. On the other hand, the partial derivative of expected utility w.r.t. d_{ik} yields

$$\frac{\partial E[U_j]}{\partial d_{ik}} = m_k(w_k^* - d_{jk})m_k \frac{n_i}{2} = outflow_{jk} \frac{\partial Prob(employed)}{\partial d_{ik}} > 0 \quad (1.22)$$

Increased commuting costs between i and k raise j 's welfare. The misfortune of laborers from i is therefore a boon for laborers at j .

The marginal effect on j 's welfare of obstacle deployment along the bilateral linkage between j and k is therefore the sum of 1.21 and 1.22:

$$\Delta welfare_j = outflow_{jk} \Delta employment_rate_j \quad (1.23)$$

1.23 shows that when commuting costs d_{jk} rise, short-run welfare changes are simply short-run employment rate changes multiplied by initial outflow of labor from j to k . Identifying the dual effects of obstacles on employment rates is therefore *sufficient* for identifying their effects on welfare.

1.4 Identification Strategy

In this section I use the commuting model to develop two generic treatment variables, ‘obstruction’ and ‘protection’, to capture the dual roles of obstacles as costly obstructions to job access on the one hand, and protective blockades against the inflow of competing labor on the other hand. I then develop instruments for these variables, arguing that Israeli settlement proximity to commuter routes largely predicted obstacle deployment during the uprising.

To measure each location j 's obstruction, I count obstacles lying along j 's primary routes²⁸ to other locations. If I knew the workplace of each laborer, I could compute pre-uprising bilateral flows between each pair of locations j and k . Bilateral flows would allow me to weight obstacles by the pre-uprising importance of location k as a destination for laborers

²⁸Without obstacles, each location has a shortest path by which to reach location k . I count the obstacles along this path to k .

from j : if $x\%$ of j 's laborers worked at k , I could weight obstacles along the road from j to k by $x\%$. As with Ahlfeldt et al (2014), however, I lack flow data. Turning to the commuting model for guidance, expression 1.15 states that the negative effect on j of commuting costs along the path from j to k will be magnified by k 's pre-uprising share of firms m_k . I therefore weight k 's importance to j by its year-2004 count of firms in k , having argued already that firms did not relocate during the uprising (see also Section 3.4 for evidence that firms did not open or close in a manner endogenous to obstacles). As a robustness check, however, I use year-2000 pre-uprising nighttime lights as an alternative weighting. Expression 1.14 meanwhile states if pre-uprising d_{jk} was high, then few laborers from j would have commuted to k . I therefore down-weight destinations k by an increasing function of their road-distance.

A large portion of 1997 laborers commuted to Israeli locations. To address this issue, I identify from the UN-OCHA maps 10 border crossings (Green-line Checkpoints) along the 1967 armistice line and treat them as if they were locations of destination.²⁹ I define the fraction of a Palestinian location j 's labor force employed in Israel as

$$work_isr_{j,1997} = \frac{\#employed_in_israel_{j,1997}}{labor_force_{j,1997}}$$

I calculate the Palestine-bound and Israel-bound parts of the index separately and then sum together, weighting by $(1 - work_isr_{j,1997})$ and $work_isr_{j,1997}$, respectively. Finally, the 2004-

²⁹I am not interested in how Palestinians commuted to work *after* crossing the border, since no obstacles were deployed inside Israel during the uprising.

firm-share-weighted index of obstruction is

$$ob_firms_j = (1 - work_isr_{j,1997}) * \sum_{k=1}^J \frac{\#obstacles_{jk} * \frac{firms_2004_k}{\sum_{k=1}^J firms_2004_k}}{\tau_{jk}} \\ + work_isr_{j,1997} * \frac{\#obstacles_{j,border}}{\tau_{j,border}}$$

and the 2000-radiance-share-weighted index of obstruction is

$$ob_bright_j = (1 - work_isr_{j,1997}) * \sum_{k=1}^J \frac{\#obstacles_{jk} * \frac{radiance_{k,2000}}{\sum_{k=1}^J radiance_{k,2000}}}{\tau_{jk}} \\ + work_isr_{j,1997} * \frac{\#obstacles_{j,border}}{\tau_{j,border}}$$

The left-hand side of Figure 1.6 depicts a spatial histogram of *ob_firms*, where each location's quantity of obstruction determines its color coding. Notably, in the northern half of the West Bank, along the border south and west and northwest of Ramallah, locations are heavily obstructed from the major job-offering areas of Ramallah and Nablus.

Complementing the obstruction indices, I design a protection index for each location, based on predictions 1.17 and 1.19. In 1.17, the model predicts that obstacles faced by laborers from *i* on their way to *k* will raise wages in *k*, thereby indirectly benefiting laborers from *j*. 1.19 shows that this beneficial effect is magnified by *i*'s share of labor force. A simple index of *j*'s protection is therefore to let *k* = *j* and measure obstruction of each location *i* on its primary route to *j*, weighting by pre-uprising labor force share of *i*:

$$protection_j = \sum_{k=1}^J \frac{\#obstacles_{kj}}{\tau_{kj}} * \frac{labor_force_{k,1997}}{\sum_{k=1}^J labor_force_{k,1997}}$$

The right-hand side of Figure 1.6 depicts a spatial histogram of *protection*, which in many




Figure 1.6: Spatial histograms of *ob_firms* and *protection*

ways complements the histogram for *ob_firms*. Laborers dwelling in the zone between Ramallah and Nablus enjoy the highest levels of protection as obstacles block out inflowing labor from more peripheral areas south, west, and north of this zone.

Instruments

In Section 3.4 I regress locations' post-uprising (2007) employment rates on obstruction-protection pairs *ob_firms* and *protection*; and *ob_bright* and *protection*; in order to estimate the dual effects of obstacles on location employment rates. These indices, however, may be correlated with unobserved variables influencing economic outcomes. For example, if locations with poor economic prospects had low opportunity costs to violence (see Miaari, Zussman, Zussman (2014)), then during the Second Intifada they may have been involved in more confrontations with the Israeli army, which may have responded by deploying more obstacles on roads leading to and from those locations. More obstructed locations, therefore, may be systematically different from less obstructed locations in a way that affects economic




Figure 1.7: Road segments where roadblocks and checkpoints are likely

outcomes like post-uprising employment rates.

Seeking exogenous variation in obstruction, I develop an instrument based on the aforementioned narrative that obstacles were deployed to protect Israeli settlements. Figure 1.8 displays UN-OCHA polygon data of settlements from 2005. Most settlements are effectively commuter suburbs of Jerusalem, the largest among them being Ma’ale Adumim, located at the West Bank’s midsection, just east of Jerusalem. A few settlements are ideologically oriented, such as Kiryat Arba’a and nearby settlements inside the Old City of Hebron, near the Tomb of Abraham. Finally, there are several large agricultural settlements in the Jordan Valley.

If indeed the Israeli army was deploying obstacles in order to protect settlements, then obstacle deployment should be predictable by settlement location. In particular, the army would have insisted on enforcing defensive buffer zones, intercepting all Palestinian traffic passing close to settlements. Figure 1.7 demonstrates this concept for a fictional pair of Palestinian locations, ‘origin’ and ‘destination’. Since two Israeli settlements lie near the connecting road, Palestinian commuter traffic from ‘origin’ cannot help but pass through the settlements’ buffer zones on their way to ‘destination’. The intersecting road segments,




Figure 1.8: Defensive buffer zones around Israeli settlements

colored dark blue, are where traffic faced a high likelihood of being detained and searched.

How wide were these buffer zones likely to have been? Arbel et al (2010) documents housing price fluctuations in the West Bank Israeli settlement of Gilo after repeated incidents of gunshot attacks by Palestinian militants at a range of 650m. I set buffer width to 500m (see Figure 1.8), then identify all segments of road intersecting buffer zones around settlements.³⁰ For each location j , I count the number of meters of road that must be traversed within 500m

³⁰Results are similar if I use narrower (300m) or wider buffers (1km).




Figure 1.9: Spatial histograms of ob_firms_seg and $protection_seg$

of Israeli settlements on the way to or from each other location k , forming indices

$$\begin{aligned}
 ob_firms_seg_j &= (1 - work_isr_{j,1997}) * \sum_{k=1}^J \frac{seg_length_{jk} * \frac{firms_2004_k}{\sum_{k=1}^J firms_2004_k}}{\tau_{jk}} \\
 &\quad + work_isr_{j,1997} * \frac{seg_length_{j,border}}{\tau_{j,border}} \\
 ob_bright_seg_j &= (1 - work_isr_{j,1997}) * \sum_{k=1}^J \frac{seg_length_{jk} * \frac{radiance_{k,2000}}{\sum_{k=1}^J radiance_{k,2000}}}{\tau_{jk}} \\
 &\quad + work_isr_{j,1997} * \frac{seg_length_{j,border}}{\tau_{j,border}} \\
 protection_seg_j &= \sum_{k=1}^J \frac{seg_length_{kj}}{\tau_{ij}} * \frac{labor_force_{k,1997}}{\sum_{k=1}^J labor_force_{k,1997}}
 \end{aligned}$$

These latter three indices are therefore defined exactly the same way as the former set, ex-

cept that in the former set I count along the path to each destination the total number of obstacles, whereas in the latter I count along the path to each destination the total length of road segments inside settlement defensive buffer zones. Whereas the obstruction indices quantify obstruction of outflowing labor from j , the protection indices quantify protection from inflowing labor to j . Comparing *ob_firms* to *protection*, we see that whereas non- j locations are weighted by their firm shares in *ob_firms*, non- j locations are weighted by their labor force shares in *protection*. Indeed, if we consider a location j from which zero laborers commute to Israeli destinations, then *ob_firms* and *protection* are distinguishable only by differences in firm shares and labor force shares of non- j Palestinian locations. Similarly, *ob_bright* and *protection* are distinguishable only by differences in radiance shares and labor force shares. Since each obstruction index never appears in any of my regressions except accompanied by a protection index, its regression coefficient is therefore always statistically distinguished by the differences in radiance and labor shares, thus obviating any calculation of radiance-net-of-labor. When the flow of laborers from j to Israeli destinations is nonzero, then *protection* is statistically distinguishable from *ob_firms* and *ob_bright* in part by these share differences, and in part by the fact that Israeli destinations demand but do not supply Palestinian labor, so that obstacles along primary routes to Israeli destinations affect *ob_firms* and *ob_bright* but not *protection*. Similar reasoning applies to the road-segment-based indices.

Figure 1.9 depicts the spatial histograms of indices *ob_firms_seg* and *protection_seg*, which largely mirror those of *ob_firms* and *protection* in Figure 1.6. In the next section, I show that the segment-based indices are valid instruments for the obstacle-based indices, finding them to be both strongly correlated and exogenous. I then use them to identify the dual effects of obstacles on post-uprising employment rates via 2SLS regressions.

1.5 Regression Analysis

In this section I test the model’s predictions 1.14, 1.17, and 1.20, showing that obstacles had dual effects on post-uprising employment rates of Palestinian locations. I use 2SLS regressions and provide strong evidence that the instrumental variables approach, relying on settlement proximity to commuter routes to predict obstacle exposure, is robust to the inclusion of many relevant covariates. I show supporting evidence that locations’ net-immigration rates of labor and household incidence of renting react in a manner consistent with what the model and overall narrative of the paper predict, although the results are statistically weak owing to the low total quantity of internal migration and change in renting behavior. Finally, I show that net-opening rates of firms across locations do not respond in a statistically stable way to obstacle deployment, further confirming that the spatial economy did not undergo any long-run reorganization in response to obstacles.

Main specification

I run 2SLS regressions with the following basic specification:

$$employment_rate_{j,2007} = \beta_0 + \beta_1 obstruction_j + \beta_2 protection_j + \beta_3 work_isr_{j,1997} + \epsilon_{j,2007}$$

The dependent variable is the vector of location employment rates in 2007, at the uprising’s end. The main 2SLS regression uses *ob_firms* for *obstruction*, instrumented for with *ob_firms_seg*. Meanwhile, *protection_seg* instruments for *protection*. As a robustness check on firm locations, a second 2SLS regression uses *ob_bright* for *obstruction*, proxying for each location’s pre-uprising job count by pre-uprising nighttime light emissions. *ob_bright_seg* instruments for *ob_bright*.

By Expression 1.14, I expect that any increase to *obstruction_j* will cause a marginal decline in *j*’s employment rate: $\beta_1 < 0$. Meanwhile, Expression 1.17 predicts that any increase to

$protection_j$ will cause a marginal increase to j 's employment rate: $\beta_2 > 0$. All regressions include $work_isr_{j,1997}$, which accounts for the effect of increased border costs on location employment rates. Locations with a higher fraction of their pre-uprising labor force employed in cross-border jobs faced higher exposure to increased border costs (Miaari, Zussman, Zussman (2014)). If I run regressions without controlling for $work_isr_{j,1997}$, then when a laborer dwelling at j loses his cross-border job, the loss will be attributed entirely to increased commuting costs resulting from obstacles along the primary route from j to the border, whereas in fact the loss of his job is due partly to obstacles but also partly to increased costs of crossing the border.

Table 1.3: First-stage regression results

Endogenous regressor	(1) ob_firms	(2) ob_firms	(3) protection	(4) protection
ob_firms_seg	0.434*** (0.0775)	0.378*** (0.0701)	-0.0531 (0.0536)	-0.0719 (0.0586)
protection_seg	-0.181** (0.0766)	-0.198** (0.0823)	0.594*** (0.0637)	0.485*** (0.0622)
ob_firms			0.293** (0.115)	0.268*** (0.0823)
protection	0.355*** (0.102)	0.338*** (0.114)		
work_isr	0.0217*** (0.00355)	0.0234*** (0.00428)	-0.0122** (0.00480)	-0.00976*** (0.00308)
Covariates	No	Yes	No	Yes
Observations	480	463	480	463
R-squared	0.702	0.859	0.663	0.847
Adjusted R-squared	0.700	0.825	0.660	0.811

All standard errors robust. Weighted regressions use 2007 labor force.

First-stage regression results

Table 1.3 lists 2SLS first-stage regression results where indices $ob_firms_seg_j$ and $protection_seg_j$ instrument for indices ob_firms_j and $protection_j$, respectively.³¹ In Column 1, ob_firms_j is regressed on $ob_firms_seg_j$, controlling for $protection_j$ and $protection_seg_j$, along with $work_isr_{1997}$. The coefficient on $ob_firms_seg_j$ is positive and highly significant. In Column 3, $protection_j$ is regressed on index $protection_seg_j$, controlling for ob_firms_j , $ob_firms_seg_j$

³¹First-stage results for radiance-based indices are available upon request.

and *work_isr*₁₉₉₇. The coefficient on *protection_seg_j* is positive and highly significant. Indices *ob_firms_seg_j* and *protection_seg_j* are therefore strong predictors of *ob_firms_j* and *protection_j*, respectively. The inclusion of many relevant covariates (described below) in Columns 2 and 4 do not destabilize the instruments' relationships with the endogenous regressors.

Table 1.4 lists the paper's main regression results. Column 1 presents the paper's main result, instrumenting for *ob_firms* with *ob_firms_seg*, and for *protection* with *protection_seg*. Column 2 repeats Column 1's regressions with numerous covariates. For reference, Column 3 displays the results of running Column 1's corresponding OLS regression, i.e. *without* instrumenting for obstruction and protection. Column 4 displays the reduced-form regression, where post-uprising employment rates are projected directly onto the instruments. For ease of interpretation, I scalar-multiply each index so that its mean value among all locations is 1.0.

Table 1.4: Main results: effects of obstacles on 2007 employment rates				
Regression type	(1) 2SLS	(2) 2SLS	(3) OLS	(4) OLS
ob_firms	-4.900*** (1.840)	-5.176*** (1.927)	-1.041 (0.642)	-1.960*** (0.634)
protection	3.744*** (1.313)	4.933*** (1.728)	0.941* (0.564)	2.099*** (0.631)
work_isr	-0.0681 (0.0702)	0.0647 (0.0761)	-0.216*** (0.0284)	-0.187*** (0.0280)
Weighted	Yes	Yes	Yes	Yes
Covariates	No	Yes	No	No
Observations	480	463	480	480
R-squared	0.273	0.526	0.351	0.369
Adjusted R-squared	0.269	0.415	0.346	0.365

All standard errors robust. All observations weighted by 2007 labor force.

In Column 1 we see that a 1-unit increase to *obstruction* resulted on average in a 4.90 percentage-point loss to employment rates as laborers struggled to reach their jobs. On

Table 1.5: Main results over different samples

Regression type	(1) 2SLS	(2) 2SLS	(3) 2SLS	(4) 2SLS
ob_firms	-4.900*** (1.840)	-5.175*** (1.910)	-5.663** (2.311)	-5.173** (2.322)
protection	3.744*** (1.313)	4.325*** (1.448)	4.537** (1.899)	3.823** (1.824)
work_isr	-0.0681 (0.0702)	-0.0445 (0.0766)	0.0145 (0.0975)	0.0118 (0.0979)
Weighted	Yes	Yes	Yes	No
Covariates	No	No	No	No
Min. laborers	0	100	200	200
Max. laborers	∞	10,000	5,000	5,000
Observations	480	368	298	298
R-squared	0.273	0.215	0.089	0.023
Adjusted R-squared	0.269	0.209	0.080	0.013

All standard errors robust. Weighted regressions use 2007 labor force.

Table 1.6: Main results, weighting obstruction by radiance

Regression type	(1) 2SLS	(2) 2SLS	(3) OLS	(4) OLS
ob_bright	-2.718*** (0.990)	-4.724** (1.863)	-0.939* (0.553)	-1.446*** (0.472)
protection	1.116*** (0.330)	2.943*** (1.092)	0.506** (0.232)	0.882*** (0.249)
work_isr	-0.152*** (0.0392)	0.0305 (0.0674)	-0.218*** (0.0281)	-0.207*** (0.0248)
Covariates	No	Yes	No	No
Weighted	Yes	Yes	Yes	Yes
Observations	480	463	480	480
R-squared	0.332	0.462	0.353	0.367
Adjusted R-squared	0.328	0.335	0.349	0.363

All standard errors robust. Weighted regressions use 2007 labor force.

the other hand, we see that a 1-unit increase to *protection* resulted on average in a 3.74 percentage-point gain to employment rates. The key point to recognize here is that for every Palestinian location j , the deployment of obstacles along the path from j to k increased both $obstruction_j$ and $protection_j$. The total effect of obstacles on 2007 employment rate per location j projects as

$$\Delta employment_rate_j = \hat{\beta}_1 obstruction_j + \hat{\beta}_2 protection_j$$

Integrating across locations' 2007 labor forces, the total effect of obstacles was to reduce Palestinian post-uprising employment by .44 percentage points. In other words, out of 472,916 laborers in 2007, the total number of employed laborers (402,366) was 2,103 smaller than it would have been had obstacles never been deployed. The obstructive effect of obstacles integrated over labor forces is -4.28 percentage points, meaning obstacle deployment caused 20,262 laborers to lose their jobs. But the protective effect of obstacles integrated over labor forces is 3.84 percentage points, meaning some 18,158 laborers *gained* jobs as an indirect consequence of obstacle deployment.

Expression 1.23 showed $\Delta welfare_j \propto \Delta employment_rate_j$. Location j is therefore a net-beneficiary of obstacle deployment (i.e. expected utility increases) as long as $\Delta employment_rate_j$ is nonnegative. The fitted values therefore indicate that 180 of 485 locations were net-beneficiaries. Figure 1.10 plots a spatial histogram of total effects of obstacles by location. The spatial correlation with job proximity (see Section 1.2) is 20.1%: locations enjoying the highest net-benefits of obstacle deployment tend to be drawn from those 'core' areas most proximate to jobs in the pre-uprising era. This is particularly noticeable in the north-middle area between Nablus and Ramallah. Those dwelling in 'core' areas (i.e. where most jobs were offered) would have tended to commute short distances and cross few obstacles on their way to work. Those dwelling in 'peripheral' areas, far from jobs, would have tended to commute




Figure 1.10: %-point change in employment due to obstacles (Table 1.4, Column 1)

farther and been more vulnerable to obstacles. The deployment of obstacles stymied the inflow of competing labor from peripheral to core areas, raising core wages and employment rates in a way that disproportionately benefited core laborers. From this perspective, core residents had reason to resent their peripherally-dwelling countrymen, and to rejoice when obstacles made their journeys to work prohibitively costly.³² The B'Tselem report quoted earlier mentions that "in almost every instance, the persons most harmed by the [besieging of cities] are the residents of the villages situated outside the [besieged cities], who depend on the services provided there." The spatial distribution of losses and benefits due to obstacles matches this core-periphery narrative well.

But are the instruments truly exogenous? If the instruments generate quasirandom variation in obstruction (protection), then I should be able to include numerous covariates on the right-hand side of Column 1's regression without substantially altering the point-estimates for *ob_firms_j* and *protection_j*. In Column 2 I re-run Column 1's regression but include 86 covariates that seem relevant *a priori* to 2007 employment rates. I now describe each of the included covariates, explaining why they might reasonably be expected to affect 2007 employment rates, and why they might correlate with the obstruction (protection) indices.

Settlements as employers

In addition to affecting the placement of obstacles, settlements themselves were a limited source of employment for Palestinian laborers. Indeed, 2.7% of the 1997 Palestinian labor force worked on settlements, as did 2.5% of the 2007 labor force. The covariate *work_settle₁₉₉₇* controls for the fraction of each census location's 1997 labor force employed on settlements. locations with larger pre-uprising employment in settlements may have been

³²Residents of Ramallah actually have a derogatory term with which to refer to Palestinians from other parts of the West Bank who commute to Ramallah for work. The word is 'Thailandiya', a bitter allusion to the immigration of foreign laborers, many of them from Thailand, who moved to Israel and filled low-skill jobs that had since 1967 mostly been monopolized by Palestinian laborers commuting from the West Bank. See Friedberg, Sauer (2003).

more or less affected by elevated security concerns during the uprising. The proximity of settlements positively affects *ob_firms_seg* and *protection_seg* since their defensive buffer zones overlap with roads leading out of Palestinian locations. Their proximity also offers an employment opportunity for Palestinian laborers, and an occasion for increased violent confrontation (Arbel et al (2010)), therefore either increasing or decreasing *employment_rate*₂₀₀₇. I account for each location j 's proximity to its nearest settlement with a nonparametric (local-linear) function of binary variables indicating if a location's nearest settlement is 0-1km, 1-2km, ..., or 9-10km away.

Obstacle proximity

If deployment of obstacles spatially correlated with deployment of Israeli army troops, then whenever obstacles were proximate to a Palestinian location, troops were likely also nearby. The presence of troops could have occasioned violent confrontation, leading to disruptions in the labor market beyond the direct channel of commuting costs. To account for this possibility I control for proximity of obstacles by counting the number of each type of obstacle lying in each of 10 distance bands (0-1km, 1-2km, ..., 9-10km) around each Palestinian location.

Violent confrontation

The Israeli army may have deployed more obstacles in the vicinity of locations heavily involved in violent confrontation during the uprising. Following Miaari, Zussman & Zussman (2014) and Cali & Miaari (2014), I use B'Tselem location-level data on Palestinian fatalities during the 2000-2007 period, generating a location-level count of all Palestinian fatalities. I generate an indicator *any_killed*_{2000–2007} for whether a location suffered any fatalities during 2000-2007 and use it as a proxy for the location's tendency to confront (be confronted by) the Israeli army. If this tendency was related to (lack of) economic opportunities, then this covariate should correlate with 2007 employment and treatment variables, therefore affecting point-estimates of interest. By similar reasoning I include each location's 1997 employment

rate as a covariate. If settlements tended to lie nearest to Palestinian locations with low pre-uprising employment, then those locations would have suffered higher quantities of obstruction (protection). If these locations' low 1997 employment rates were the result of poor location fundamentals, then these fundamentals would tend to negatively influence 2007 employment rates. Alternatively, if low pre-uprising employment rates meant that residents had low opportunity costs of confrontation during the uprising, then subsequent incursions by the army might likewise have negatively affected 2007 employment rates.

Education & experience

I control furthermore for average pre-uprising educational attainment (in years) avg_educ_{1997} and age avg_age_{1997} of each location's labor force. Reports indicate that younger men were more likely to be profiled as militants and tended to be detained longer at checkpoints, so obstacles should have had larger effects on locations with low average age of labor force. Laborers with higher educational attainment, meanwhile, would have been less substitutable, so employers would have tended to retain them despite late arrivals to work.

Public sector

The commuting model accounts for the behavior of profit-maximizing, private-sector firms, but in 2007 fully 10.1% of employed Palestinians were government workers. When its employees ran into obstacles on their way to work, the government may have responded differently, perhaps more leniently, laying off laborers more reluctantly. I control for the share of each location's pre-uprising labor force employed in the public sector.

Net-immigration rates

For each person who relocates from location j in order to keep his job at location k , location j 's employment rate declines and location k 's increases. In this way, obstacles can produce the reaction in employment data predicted by the model, but by the long-term

channel (relocation) instead of the short-term channel (wage adjustments). To confirm that the short-term channel largely accounts for observed employment changes, I control for net-immigration rates of locations. In particular, I calculate the number of persons immigrating to each location for job-related reasons during 2000-2007, and subtract the number of persons emigrating for job-related reasons during 2000-2007, dividing the difference by the 2007 labor force. The resulting covariate is the net-immigration rate, which I include on the RHS as a covariate. In a later table I show that although net-immigration rates do not interfere with my main results, they do indeed respond to obstacles as one might expect, and in that way lend further credibility to my overall narrative of dual effects.

Firm openings and closings

Related to the issue of decreased demand and laying off laborers, some firms may have altogether closed down. Alternatively, firms may have continued to open in locations less affected by changing trade conditions, while ceasing to open in more affected locations. Since the firm census records closed firms and firms preparing to open, I calculate the fraction of all firms in each location that are closed or preparing to open and include these as two covariates.

Industrial sectors, trade, and border costs

Using the firm census I categorize firms by 15 broad industrial categories, including agriculture, manufacturing, retail, mining, etc. I calculate the fraction of each location's firms belonging to each of these categories. By including the resulting 15 covariates, I am controlling for variation in locations' sensitivity to changes in trade conditions. For example, if the West Bank's retail sector sold imported goods to Palestinians, then increased border costs during the uprising may have raised the sale prices of merchandise, reducing demand and potentially leading to lower employment. locations with a larger percentage of firms involved in retail should therefore have experienced more employment losses. Meanwhile, census locations exporting agricultural produce, manufactured goods, or mined goods (Jerusalem stone,

for example. See World Bank (2013)) would have been sensitive to changes in demand among other Palestinian locations or among consumers in the outside world. Increased border costs would have raised the price of exports in foreign markets, decreasing demand; while obstacles would have raised the price of exports in the marketplaces of other Palestinian locations, decreasing Palestinian demand for these tradeables. Facing decreased demand, employers would have curbed supply, potentially laying off laborers. To confirm that commuting, and not trade, drives my results, I therefore include these 15 sector covariates to see if point-estimates change.

Private cars

I control for the fraction of each location’s households owning a private car in the pre-uprising era. If private cars were better able to circumvent obstacles, or if they were treated differently by soldiers when crossing a checkpoint, then treatment effects would covary.

Outside options

The model emphasized that laborers give up employment when wages-net-of-commuting-costs are exceeded by ‘reservation’ wages. One possible outside option available to Palestinian laborers was to return to agricultural activity. The 1997 census recorded the number of dunums³³ of agricultural land owned by each respondent. I calculate average number of dunums owned by laborers in each location j and include this as a covariate named *agri_per_laborer*₁₉₉₇. Laborers with more agricultural holdings would have had more to live on if unemployed, so they should have exhibited a greater tendency to give up employment when facing obstructions. If any of these covariates correlated strongly with both instruments and endogenous regressors, then their inclusion would cause point-estimates of the

³³A dunum is a traditional Palestinian farmland area unit.

instruments' coefficients to shift significantly.

Column 2 of Table 1.4 repeats Column 1's regression, but includes all covariates discussed above. R^2 increases over 25 percentage points, from 27% to 52%, yet neither the point-estimates of interest, nor their statistical significance, are affected. This result strongly affirms the validity of the instrumental variables approach. The large increase in R^2 suggests that indeed the added covariates are relevant and explain a lot of the variation in employment rates across locations, yet they leave the commuting narrative unaffected. Furthermore, the R^2 is very high in Column 2. If the introduction of more unobservables were to continue to move point-estimates at a similar rate to how they move between Columns 1 to 2, then R^2 would reach 100% long before either point-estimate of interest could reach zero.

Columns 3 and 4 allow us to peek under the hood of the 2SLS regression to see what the instruments are doing. Column 3 displays the OLS regression result with endogenous regressors, while Column 4 displays the OLS reduced-form regression, i.e. using the instruments directly as regressors. The point-estimates are signed as expected in both columns, but do not achieve statistical significance in Column 3. The likeliest explanation for this is attenuation bias. While UN-OCHA maps accurately report the locations of obstacles at the time of recording, the World Bank and B'Tselem consistently report that obstacles often changed places from day to day. In particular, the army often set up makeshift, 'flying' checkpoints in an effort to surprise militants with spot inspections (World Bank, 2007). So whereas the instrumental variables draw on areas that required perpetual defending (settlements and their immediate surroundings), obstacle counts draw partially from areas that were only temporarily guarded, but happened coincidentally to be guarded at the time when UN-OCHA drew up its maps. This idea also extends to the issue of how intensely checkpoints were guarded: around settlements, they were probably guarded with great vigilance, since Israeli civilian

lives were at stake; while elsewhere there may have been more variability in enforcement. The obstacle counts are therefore noisier and so estimates of β_1 and β_2 are attenuated in Column 3. In Column 4, by contrast, the focus on friction zones around settlements filters out the noise introduced by enforcement variability and flying checkpoints.

In Table 1.5, the main regression is re-run over different restricted samples. Column 1 reprints the results of Table 1.4, Column 1, for ease of comparison. Column 2 restricts focus to all Palestinian locations where at least 100 laborers dwelt in both 1997 and 2007, and where no more than 10,000 laborers dwelt in either year. The upper bound of 10,000 eliminates outlier cities like Nablus and Hebron, while the lower bound discards over 100 small villages that contribute a lot of noise to the main regression. Point-estimates do not substantially change between Columns 1 and 2, owing to the fact that, since Column 1's regression is weighted by 2007 labor force, the small-village observations were already being heavily discounted. Given this weighting scheme, however, it is comforting to see that the exclusion of large outlier towns makes so little difference to point-estimates, since one might be concerned that their outcomes were weighted too heavily in Column 1. Likewise in Column 3, when locations with labor forces smaller than 200 or greater than 5,000 are excluded, point-estimates are remarkably stable, suggesting very small and very large towns are not driving the paper's main results. In Column 4, I re-run Column 3's regression without weighting observations, to demonstrate that likewise the weighting scheme is unnecessary once the tails of the distribution of location sizes are chopped off.

Table 1.6 repeats the regressions of Table 1.4, only this time the destination locations are weighted by radiance, not by firm counts, instrumenting for *ob.bright* with *ob.bright_seg*. The similarity in signs, statistical significance, and even magnitudes between Column 1 of Table 1.6 and Column 1 of Table 1.4 is not far short of a miracle, and ought to be interpreted as an extraordinary affirmation of the efficacy of nighttime lights in tracking production ac-

tivity. Note that, since both *protection* and *ob_bright* are included in the regression, the point-estimates on these variables are identified off the *differences* in labor shares and radiance shares (for a full discussion, see Section 1.4). In other words, the point-estimate on *ob_bright* is identified off of the *residual* light, *after* a location’s residential labor force share is accounted for. So these regression results are not exploiting the fact that the total light of a location will be larger when its residential labor force is larger. Instead, the regression specification implies that total light *net of* residential labor force identifies the point-estimate on *ob_bright*. In view of these facts, it is incredible that the point-estimates in Column 1 of Table 1.4, where *ob_firms* weights destination locations by the count of firms as recorded by the Palestinian Central Bureau of Statistics, are so similar to the point-estimates in Column 1 of Table 1.6, where *ob_bright* weights destination locations by the count of light as recorded by US Air Force meteorological satellites. As such, Table 1.6’s results act as an extraordinary validation of Table 1.4’s results, drawing on an utterly different data source to lend credibility to the firm census; and as an extraordinary validation of lights data as a meaningful source of economic information, in that they are able to track Palestinian production activity net of residential size with such finesse as to rival the census bureau’s own firm counts.

Spatial autocorrelation

Census locations in this dataset are separate, named Palestinian neighborhoods and villages of the West Bank. In some cases, however, the locations are so close to each other that the assumption that each is an independent draw is not credible. For example, Beit Lahem and Beit Sahour are separate locations on paper, but in reality there is no recognizable boundary between them, and only someone familiar with these towns would know when one has passed from one to the other. Similarly, the city of Ramallah is comprised of several adjacent neighborhoods, including Al-Bireh, Ramallah, Qaddura Refugee Camp, among others. So as to be sure that the paper’s main results are not driven by an artificially large number of highly correlated observations, I cluster observations according to a PCBS formula devised

in consultation with the World Bank.³⁴ The formula groups 464 census locations into 278 clusters. Table 1.7 repeats Columns (1) and (2) of Table 1.4 alongside identical, clustered regressions in Columns (3) and (4). The point estimates increase in magnitude and remain statistically significant.

Table 1.7: Results with PCBS/World Bank spatial clusters

Regression type	(1) 2SLS	(2) 2SLS	(3) 2SLS	(4) 2SLS
obstruction	-4.900*** (1.840)	-5.001*** (1.864)	-8.040** (3.353)	-7.812*** (3.021)
protection	3.744*** (1.313)	5.017*** (1.697)	5.071*** (1.889)	6.152*** (2.310)
work_isr	-0.0681 (0.0702)	0.110 (0.0799)	0.0280 (0.119)	0.206 (0.130)
Covariates	No	Yes	No	Yes
Clusters	No	No	271	271
Observations	480	463	443	428
Adjusted R-squared	0.269	0.428	0.234	0.405

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

There is another, more subtle spatial autocorrelation concern: if for each location j the indices of obstruction and protection were built using pre-uprising characteristics of j itself, then there would be concern that j 's 2007 employment rate was effectively being regressed on a function of lagged characteristics of j , whose correlation with 2007 characteristics might subtly introduce endogeneity problems. For example, if j 's pre-uprising radiance (proxying for job counts) were highly correlated with its post-uprising radiance, and if that positively predicted j 's 2007 employment rate, then the presence of j 's pre-uprising radiance in j 's protection index would bias upward my estimate of the (positive) protective effect of obstacles. Although the pre-uprising characteristics of j are *not* used to form j 's treatment indices, this endogeneity concern persists through spatial autocorrelation: if k 's pre-uprising characteristics are used in j 's obstruction (protection) index, and if k is *near* to j , then I am effectively regressing j 's 2007 employment rate on correlates of j 's 1997 characteristics,

³⁴These clusters were identified by the PCBS to assist the World Bank with its poverty assessment of the West Bank. I thank Brian Blankespoor for generously sharing these data with me.

Table 1.8: Effect of obstacles on locations of firms and laborers

<i>Dependent variables</i>	(1) net immigration	(2) net immigration	(3) renting incidence	(4) renting incidence	(5) net firm openings	(6) net firm openings
ob_firms	-1.943* (1.008)	-1.167* (0.596)	-5.392** (2.484)	-2.550** (1.024)	-1.115 (1.672)	-0.395 (1.375)
protection	3.906*** (1.239)	1.522** (0.693)	6.681** (2.833)	2.031** (0.927)	1.310 (1.602)	2.647* (1.400)
work_isr	0.0502 (0.0369)	0.00946 (0.0246)	-0.0899 (0.0892)	0.0552* (0.0291)	0.0739 (0.0947)	0.0422 (0.0559)
Observations	480	463	478	463	464	463
R-squared	0.094	0.713	0.281	0.901	0.013	0.492
Adjusted R-squared	0.088	0.646	0.277	0.878	0.006	0.376

Robust standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

which are themselves correlates of j 's 2007 characteristics. To address this concern, I re-run all of the regressions in Tables 1.3 and 1.4 insisting that partner locations k are at least 20km away from each j . I find my results are virtually unchanged.³⁵

Spatial redistribution of firms and laborers

The commuting model predicts welfare losses and gains in the short run, when firms and laborers have not yet relocated or redistributed over space. Relocation and spatial redistribution of the economy would have tended to mitigate the welfare consequences of obstacles. Table 1.8 shows that there is some statistically weak evidence that laborers relocated, and very weak evidence that firms redistributed in response to obstacles. Taken together with evidence presented in Section 1.4, Table 1.8's results justify this paper's focus on the short run, and are consistent with the broader literature's finding that there can be persistent factor immobility in developing country settings, even in the presence of significant spatial welfare inequality.

In Columns (1) and (2), the dependent variable is the net-immigration rate, i.e. the difference between the number of persons immigrating to and emigrating from each census

³⁵Results available upon request.

location during 2000-2007, divided by the location's 2007 labor force and multiplied by 100. The commuting model predicts that when obstacles remotify a location from job opportunities, expected utility derived from dwelling in that location declines, leading to net-emigration (negative net-immigration) in the long run. To the degree to which obstacles protect local laborers from inflow of competing labor, expected utility should rise, leading to net-immigration in the long run. Agreeing perfectly with the model's predictions, the point-estimate of the marginal effect of *obstruction* on net-immigration is negative in Column (1), while the marginal effect of *protection* is positively signed. The point-estimate for *obstruction* is not significant at the 5% level, however. Furthermore the point-estimates are not stable when covariates are introduced in Column (2). The point-estimate on *protection* is more than halved, while the point-estimate on *obstruction* remains insignificant at the 5% level. In summary, the results in Columns (1) and (2) support the overall narrative, but their statistical weakness supports this paper's focus on short-run consequences.

In Columns (3) and (4), the dependent variable is the percentage of households who report in 2007 that their property of residence is rented. An anecdotally supported hypothesis about the effect of obstacles is that Palestinian laborers, unwilling to undertake high commuting costs daily, began to rent apartments in their workplace locations, returning to their hometowns only on weekends, bi-monthly, or at sparser intervals. The results in Columns (3) and (4) suggest that there is likely some truth to that hypothesis: rental incidence in 2007 is a declining function of *obstruction* but an increasing function of *protection*. The effect is significantly mitigated by the inclusion of covariates, but remains statistically significant at the 5% level, and given the very high R^2 achieved in Column (4), there is not much unobserved variation left over that could otherwise account for the results. It remains only to say that, like net-immigration of labor, the total amount of change in rental incidence between 1997 and 2007 is quite small: if for each location I take the absolute value of the difference between 1997 and 2007, the sum across all locations is 8,527, or just 2.3% of households in

2007.

In Columns (5) and (6), the dependent variable is the net firm openings rate, i.e. the difference between the number of firms recorded as preparing to open and the number of firms closed in each census location, divided by the total number of firms (closed, open, or opening) censused in that location in 2004. None of the point-estimates in Columns (5) or (6) is significant at the 5% level, undermining any alternative hypotheses in which firms tend to open or close at different rates in response to obstacles. Coupled with the radiance-based evidence presented earlier, I conclude that firms did not redistribute over space in response to obstacles.

1.6 Conclusion

Economists have always found that commuting costs and laborers' welfare are inversely related in the short run, but this paper finds robust empirical evidence of a simple labor market channel by which commuting costs act in much the same way as trade barriers, protecting some laborers while harming others. My empirical findings suggest that in cities of the developing world, where laborers are very interchangeable and unemployment rates are high, commuting cost increases will cause one set of laborers to lose out on job access, but the resulting vacancies will benefit another set of laborers more advantageously situated. As we move to advise mayors and municipalities around the developing world about infrastructure investments, the absence of a clear, Pareto welfare gain would seem to undermine somewhat the argument for costly investment in maintaining or expanding commuter infrastructure. In the case of the West Bank, however, this paper's results put some urgency in the case for investing in ways to buoy the economic prospects of peripheral Palestinian locations. In particular, the presence of Israeli settlements creates the necessity for defensive zones and spot inspections of vehicles and passengers. The sluggishness of commuter traffic along

these routes tends to disconnect peripheral Palestinian areas from core Palestinian areas, generating economic inequality and encouraging periphery-to-core migration. The ensuing depopulation of rural areas leaves them vulnerable to annexation, making the possibility of a two-state solution ever more tenuous.

1.7 Bibliography

ABRAHAMSON, A., N. LOZANO-GRACIA, C. J. ORAM. 2014. Correcting Overflow in Nighttime Lights Data. Unpublished manuscript.

AGNEW, J. G., T.W., Gonzalez, J. and Min, B. 2008. Baghdad Nights: Evaluating the US Military Surge Using Nighttime Light Signatures. *Environment and Planning A*, 40, 2285-95.

AHLFELDT, G., S. REDDING, D. STURM and N. WOLF. 2014. The Economics of Density: Evidence from the Berlin Wall, NBER Working Paper, 20354.

AIX GROUP. 2013. Twenty Years after Oslo and the Paris Protocol: Changing Course, Averting Crisis From Dependency to Economic Sovereignty.

ALESINA, A., S. MICHALOPOULOS and E. PAPAIOANNOU. 2014. Ethnic Inequality. Working Paper 155, Centre for Competitive Advantage in the Global Economy, Department of Economics, University of Warwick.

ANGRIST, J. D. 1996. Short-Run Demand for Palestinian Labor. *Journal of Labor Economics*, 14(3), 425-453.

ARBEL, BEN-SHAHAR, GABRIEL, TOBOL. 2010. The local cost of terror: Effects of the Second Palestinian Intifada on Jerusalem housing prices. *Regional Science and Urban Economics*, 40, 415-426.

BAUM-SNOW, N. 2007. Did Highways Cause Suburbanization? *Quarterly Journal of Economics*, 122(2), 775-805.

BAUM-SNOW, BRANDT, HENDERSON, TURNER, ZHANG. 2014. Roads, Railroads and

Decentralization of Chinese Cities. Working paper.

BAUM-SNOW, TURNER. 2014. Transportation and the Decentralization of Chinese Cities. Working paper.

BOWES, E. R., IHLANFEDLT, K. R. 2001. Identifying the impacts of rail transit stations on residential property values. *Journal of Urban Economics*. 50 (1), 1-25.

BRYAN, G., CHOWDHURY, S., MOBARAK, A.M. 2014. Under-investment in a protable technology: The case of seasonal migration in Bangladesh. *Econometrica*, forthcoming.

B'TSELEM. 2007. 'Ground to a Halt'. www.btselem.org.

CALI, M., S. MIAARI 2012. The Labour Market Impact of Mobility Restrictions: Evidence from the West Bank. HiCN Working Paper No. 130.

CAMMETT, M. 2014. Compassionate Communalism: Welfare and Sectarianism in Lebanon. Cornell University Press.

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.

CROFT, T.A. 1979. The Brightness of Lights on Earth at Night, Digitally recorded by DMSP Satellite. U.S. Geological Survey.

DEBREZION, G. PELS, E. RIETVELD, P. 2007. The impact of railway stations on residential and commercial property value: a meta-analysis. *Real Estate Finance and Economics*. 35 (2), 161-180.

DOLL, C. N.H., J. P. MULLER, and J. G. MORLEY. 2006. "Mapping regional economic activity from night-time light satellite imagery." *Ecological Economics* 57: 75-92.

ELSTER, Y., A. ZUSSMAN, N. ZUSSMAN. 2014. Rockets: The Housing Market Effects of a Credible Terrorist Threat. Working paper.

ELVIDGE, C. D., K. E. BAUGH, J. B. DIETZ, T. BLAND, P. C. SUTTON, H. W. KROEHL. 1998. Radiance calibration of DMSP-OLS Low-Light Imaging Data of Human Settlements. *Remote Sensing of Environment*, 68, 77-88.

FRIEDBERG, R. and R. SAUER. 2003. The Effect of Foreign Guest Workers on the Labor Market Outcomes of Palestinians from the West Bank and Gaza Strip. Hebrew University

Falk Institute Discussion Paper 03.08.

GALOR, O., A. MOUNTFORD. 2008. Trading Population for Productivity: Theory and Evidence. *Review of Economic Studies*, 75(4), 1143-1179.

GOLLIN, D. R. JEDWAB, D. VOLLRATH. 2013. Urbanization with and without Industrialization. Working paper.

GONZALEZ-NAVARRO, M., C. QUINTANA-DOMEQUE. 2014. Paving Streets for the Poor: Experimental Analysis of Infrastructure Effects. Working paper.

GRAHAM-BROWN, S. 1979. The West Bank and Gaza: The Structural Impact of Israeli Colonization. *MERIP Reports*. 74, 9-22.

HARRIS, J. R., M. P. TODARO. 1970. Migration, Unemployment and Development: A Two-Sector Analysis. 1970. *American Economic Review* 60 (1): 126142.

HENDERSON, V., A. STOREYGARD, U. DEICHMANN. 2014. Has climate change promoted urbanization in Sub-Saharan Africa? mimeo, Tufts University.

HENDERSON, J. V., A. STOREYGARD, D. N. WEIL. 2012. "Measuring Economic Growth from Outer Space." *American Economic Review*, 102(2): 994-1028.

HILAL, J. 1977. Class Transformation in the West Bank and Gaza. *Journal of Palestine Studies*. 6(2), 167-175.

IHLANFELDT, K. R., D. L. SJOQUIST. 1998. The Spatial Mismatch Hypothesis: A Review of Recent Studies and Their Implications for Welfare Reform. *Housing Policy Debate*. 9, 849-492. IRWIN, D. A. 2005. The Welfare Cost of Autarky: Evidence from the Jeffersonian Trade Embargo, 180709. *Review of International Economics*, 13(4), 631-645.

KAIN, J. F. 1968. Housing segregation, negro employment, and metropolitan decentralization. *Quarterly Journal of Economics*. 82, 175-197.

LEVIN N., Y. DUKE. 2012. High spatial resolution night-time light images for demographic and socio-economic studies. *Journal of Remote Sensing of Environment*. 119, 1-10.

MCMILLEN, D. P., J. MCDONALD. 2004. Reaction of housing prices to a new rapid transit line: Chicago's Midway line, 1983-1999. *Real Estate Economics*. 32 (3), 463-486.

- PASERMAN, D. 2008. The Cycle of Violence? An Empirical Analysis of Fatalities in the Palestinian-Israeli Conflict. *American Economic Review*, 98(4), 1591-1604.
- PINKOVSKIY, M. 2013. Economic Discontinuities at Borders: Evidence from Satellite Data on Lights at Night. Working paper.
- MIAARI, S. ZUSSMAN A., ZUSSMAN N. 2014. Employment Restrictions and political violence in the Israel-Palestinian conflict. *Journal of Economic Behavior & Organization*, 101, 24-44.
- MUNSHI, K., ROSENZWEIG, M. 2013. Networks and misallocation: Insurance, migration, and the rural-urban wage gap. Unpublished paper, Yale University.
- SMALL, C., F. POZZI, C. D. ELVIDGE. 2005. Spatial analysis of global urban extent from DMSP-OLS night lights. *Remote Sensing of Environment*. 96, 277-291.
- SMITH, T., Y. ZENOU. 2005. Spatial mismatch, search effort, and urban spatial structure. *Journal of Urban Economics*. 54, 129-156.
- TOPALOVA, P. 2010. Factor immobility and regional impacts of trade liberalization: Evidence on poverty from India. *American Economic Journal: Applied Economics*. 2, 141.
- VAN DER WEIDE, R. B. BLANKESPOOR, B. RIJKERS, A. ABRAHAMS. 2014. How Valuable is Market Access? Evidence from the West Bank. unpublished manuscript.
- WASMER, E. Y. ZENOU. 2002. Does city structure affect job search and welfare? *Journal of Urban Economics*. 51, 515-541.
- WORLD BANK REPORT. 2007. Movement and Access Restrictions in the West Bank: Uncertainty and Inefficiency in the Palestinian Economy.
- WORLD BANK REPORT. 2013. West Bank and Gaza: Area C and the Future of the Palestinian Economy.
- WORLD BANK REPORT. 2014. More Jobs, Better Jobs: A Priority for Egypt.

Chapter 2

Correcting Overglow in Nighttime Lights Data

2.1 Introduction

Economists have always gone out of their way in search of good data that will allow them to test their hypotheses. Today they are going further, looking to outer space, and increasingly relying on satellite nighttime lights data as a proxy for economic activity at national, provincial, and municipal levels.

Nighttime lights data have long been recognized in both the remote sensing (Doll et al (2006)) and economics literatures (Chen and Nordhaus (2011), Henderson et al (2012)) as a strong correlate of economic activity and growth. Recently, a flurry of papers has taken advantage of these data in contexts where more traditional measures of economic production and consumption are poorly measured or altogether missing. For example, while annual, national GDP estimates are typically available for most countries and most years, Henderson et al (2012) show how optimal weighting of lights and national accounts data improve the reliability of GDP estimates, particularly in countries where data collection and processing

is weak. Informal economic activity is also notoriously underestimated by official accounting methods, and recent papers explore ways of using lights to improve such estimates (Ghosh et al (2009), Haraati and Hardy (2014)).

Official subnational estimates of GDP tend to be scarce, particularly in developing countries. Recent papers have therefore relied exclusively on nighttime lights to proxy for GDP levels and growth in subnational regions. Alesina et al (2012) rely on nighttime lights to estimate per-capita income of subnational regions occupied by various ethnic groups in over 150 countries, finding that income inequality along ethnic lines tends to harm national economic performance. Pinkovskiy (2011) measures nighttime light emissions in buffer zones around national borders, arguing that institutions (and not geography) best explain the discontinuous cross-border differences in economic performance.

The granularity of lights data¹ has tempted researchers to pursue questions at even lower levels of spatial disaggregation. For example, there are several recent studies for which cities (towns, villages) are the unit of analysis. Storeygard (2014) exploits nighttime lights to track differential economic growth rates of almost 300 inland sub-Saharan African cities near and far from major ports when world oil prices exogenously increased in the 2000s. Levin and Duke (2012), Abrahams (2014) and van der Weide et al (2014) study the Israel-Palestine Conflict, using town-level nighttime lights measures in the West Bank as a proxy for business activity. Pushing the data even further, several recent studies have attempted *intra*-city analysis. Agnew et al. (2008) use lights data as a proxy for energy consumption in neighborhoods of Baghdad to explore whether the strategic surge of US troops in early 2007 led to subsequent improvements in quality of life through lower levels of violence. Henderson et al (2014) and Baum-Snow and Turner (2014) rely on nighttime lights to track the decentralization of Chinese cities in the 1990s in response to infrastructural expansion by comparing

¹DMSP-OLS nighttime lights data, for example, are available in grid cells 1/2 nautical-mile times the *cosine* of the latitude.

light growth in the cores and peripheries of cities.

However, as with data collected on Earth, data from outer space are not free of problems. Most studies to date have relied on NOAA’s DMSP-OLS nighttime lights annual imagery with global coverage since 1992. These imagery have long been known to suffer from two key sources of error: saturation and overglow (Croft (1979)).² On the one hand, saturation error, leading to topcoding of lights data in bright urban areas, is easily recognized and well understood (Elvidge et al. (1997)). Economics research to date has avoided heavily topcoded areas, typically concentrating on the developing world where topcoding is less prevalent. In recent work, radiance-calibrated low-gain imagery and electrification data have been used convincingly to recover true light emissions over topcoded areas in parts of eastern China (Letu et al (2012)). Researchers have also had some success in ameliorating topcoding using the negative correlation of vegetation intensity (NDVI) and brightness (Zhang et al (2013)). Most recently, (non-topcoded) radiance-calibrated imagery have become available for multiple years.³

On the other hand, the problem of overglow (or blooming, as it was earlier known), which refers to the apparent blurring of nighttime lights imagery, has never been well understood. Researchers have been aware of the problem since as early as Croft (1979). The overglow phenomenon was first acknowledged by the remote sensing literature as a nuisance in estimating urban extents (Imhoff et al (1997); Henderson et al (2003); Small, Pozzi, Elvidge (2005)), since spillage of light beyond its sources tended to bias upward estimates of urban extent. Consistent with their results, we find in this paper that the extents of urban localities

²More recently, geolocal shifting has been identified as another source of error (Tuttle et al (2012)).

³Radiance-calibrated annual composite imagery of the globe are now available for free download from the DMSP website for years 1992, 1996, 2000, 2006, and 2010

are typically overestimated by a factor of 9.5 when using nighttime lights.

Within economics, however, overflow has the potential to further confound economic analysis, since lights data are also used to quantify GDP levels and growth rates at different locations. For a start, overflow can corrupt absolute measures of economic activity. Nighttime lights researchers typically obtain a boundary definition (polygons) for each area of interest (neighborhood, city, country, ethnic homeland, etc.) and then calculate the total light observed within the boundaries of each area. But light emitted near the boundaries often spills beyond borders and across areas, erroneously falling in or out of calculation. Coastal towns spill light into the ocean. Border towns bleed light onto neighboring state or beyond national boundaries. For the selected sample of cities studied in this paper, we find that if boundaries are defined using accurate urban extent data, an average of 72% of emitted light spills beyond those boundaries.

Overflow also confuses intensive and extensive growth of cities. Recent work by Henderson et al (2014) and Baum-Snow and Turner (2014) use nighttime lights to compare economic growth rates in the cores and peripheries of Chinese cities. The latter paper finds that light growth in peripheral areas is relatively high, taking this as evidence of economic decentralization due to the arguably exogenous introduction of intracity transport infrastructure. The nature of overflow, however, is that light is redistributed *in proportion* to what is emitted. This guarantees that even if city centers were the only locations of economic growth, peripheral areas would nevertheless have appeared to have grown relatively fast, because newly emitted light in the city center spilled over into the periphery.

Overflow error is particularly harmful to border discontinuity designs. Pinkovskiy (2011) uses nighttime lights to predict differences in GDP levels and growth rates across national borders. He argues that within sufficiently small buffer zones around shared land borders,

discontinuous differences in economic activity cannot be explained by continuously changing geography, but must be due to discontinuously changing national institutions. The analysis is challenged, however, by cross-border overflow from ‘richer’ to ‘poorer’ countries, which tends to bias downward the gaps in levels and growth rates.

In this paper we offer a novel and comprehensive explanation for why overflow occurs. We argue that overflow is fundamentally the result of a mismatch between the DMSP satellites’ onboard optics and onboard data management. In particular, the sensor has a circular field of view and therefore an elliptical ground scanning area. As we show, this ultimately implies that the satellite’s resolution on any given night is not small and square-shaped as the pixelization would suggest, but rather large and elliptical. This is why, though the data are stored in square pixels, the imagery appear replete with ellipses of light.

With this explanation in hand, we derive the geometry by which overflow occurs and for any light source we can generate the corresponding blurred image. The key parameter to accurately predicting the blur is the radius of the ground scanning area at nadir. For dim light sources, we explain why that radius is rather small (we estimate 2.0km), leading to an elliptical distortion that illuminates pixels no more than 4.0km from the source. For brighter sources, however, we explain why the ground scanning area is much larger (we estimate 41km), thus illuminating pixels up to 77km away from the source.

Since we know how overflow is generated, we can also reverse the process and de-blur the imagery. We provide a simple mathematical derivation for deconvoluting the imagery and implement our method in a series of Python and Matlab scripts.⁴ We apply the method to de-blur the imagery of a selected sample of cities in sub-Saharan African and South Asia. A comparison of blurred and de-blurred imagery suggests that indeed the resolution of the

⁴Available upon request.




Figure 2.1: DMSP-OLS sensor data collection

imagery is dramatically improved.

2.2 Physical origins of overflow

The objective of this section is to provide a novel and comprehensive explanation of how overflow error arises in DMSP-OLS data. Overflow error (or ‘blooming’) was first recognized by Croft (1979), but to date no research article has provided a step-by-step, comprehensive explanation of how this phenomenon occurs.⁵ We argue that overflow results almost entirely from the optics of the DMSP satellites, the behavior of the onboard light sensor, and the

⁵Small, Pozzi, Elvidge (2005) mention geolocation error, atmospheric scattering, and pixel resampling as likely sources of overflow error, but do not explicitly discuss their reasoning, focusing instead mainly on an empirical examination of overflow in the data. Doll (2008) lists these potential sources of error with some further discussion, but again refrains from thoroughly explaining how the error is generated.




Figure 2.2: Misattribution of light in consecutive scans

onboard data management.

DMSP nighttime lights

Satellites of the US Air Force's Defense Meteorological Satellite Program (DMSP) have been collecting nighttime lights data since the early 1970s, though only since 1992 have these data been organized into global, annual composites easily accessible by researchers. The purpose of these satellites is to monitor weather patterns by spotting cloud formations hovering above the earth's surface (Croft (1979), Doll (2008)). To accomplish this task, the satellites are equipped with several sensors, including an Operating Line Scanner (OLS hereafter) designed to detect light reflecting off the tops of moonlit clouds at night. However, on nights when cloud cover is thin or absent over a given locality, the scanner can see through to the Earth's surface and record nighttime light emissions.

Overflow error results fundamentally from the considerable difference between what the scanner is actually viewing, and what it claims to be viewing. The scanner claims to be

viewing square raw pixels of side-length .56km,⁶ but actually views areas of varying sizes and shapes, all of which are vastly larger than the raw pixels to which the data are being ascribed onboard. Figure 2.1 depicts the data collection process. While traveling from pole to pole at an altitude of 833km, the OLS gazes downward to nadir, then swings its gaze eastward and westward in a 3000km- wide sweep (Doll (2008)). Figure 2.1 depicts a DMSP satellite traveling north, making side-to- side sweeps of the ground. In Figure 2.1, the scanner's current (instantaneous) ground scanning area is depicted in this figure by a red ellipse. At nadir, where the distance from the satellite to the Earth is minimal, the ground scanning area is smallest and is set by the geometry of the satellite's detection system. Here we have depicted the ground scanning area as circular.⁷

As the scanner swings east or west, the distance grows from the satellite to the ground scanning area, so the circular ground scan area expands and morphs into an ellipse (longer in the East-West direction than the North-South direction). A quantity of photons rising from within the elliptical ground scanning area enters the scanner and gives rise to a pulse of electrons. The size of this pulse is then mapped to a digital number. The satellite's internal data management can spare only 6 bits of memory for each datum, so the size of the pulse is mapped to numbers between 0 and 63 representing light intensity.

That measure of light intensity, however, is not ascribed to the red ellipse, from which the light actually originated, but rather to the green, square pixel centered on the centroid of the ellipse (see Figure 2.1). In short, the area from which light is collected by the OLS is considerably larger than the area of the pixel to which light is ascribed. Misattribution of light emission is immediate: even if the geolocation of the pixel contains no light-emitting

⁶The DMSP website claims raw pixels are of length .55km, while Doll (2008) claims they are of .56km in length. This difference has no impact on our qualitative explanation of overglow. Our quantitative work is easily adjusted to accept either value.

⁷The true shape of the ground scanning area is important for quantitatively estimating and removing overglow error, but irrelevant for comprehending how overglow is generated.

sources, the pixel will nevertheless be recorded as lit if even one significant light source lies anywhere within the much larger ellipse.

When the scanner shifts its focus to the next ground scanning area (one shift every 16 microseconds), the new ellipse is centered 0.56km further east/west, significantly overlapping the previous ellipse. Figure 2.2 depicts the consequence of consecutive overlaps for pixels in the vicinity of one light-emitting source (symbolized by a yellow star in the middle of the diagram). In the example illustrated in Figure 2.2, a west-to-east scan would result in 11 consecutive ellipses (scanning areas) containing the light source. Therefore, while only one pixel emits the light observed, the scan would erroneously assign the same quantity of light to each one of the 11 consecutive pixels. Figure 2.3 shows the light pattern that emerges as the satellite rescans the same longitudes from different latitudes on its pole-to-pole journey. Each white triangle in Figure 2.3 represents a pixel registering positive light intensity due to the misattribution of emissions by the yellow star light source. We see that, ultimately, all pixels are lit within the ellipse centered at the light-emitting pixel. The value assigned to each pixel is again exactly the same.

Storing the light intensity values for millions of fine pixels is troublesome, especially for 1970s-era technology. Instead, the satellite's onboard data management system takes 5 x 5 sets of fine pixels and *averages* their total light, producing 2.8km x 2.8km *coarse* pixels (Doll (2008)). The coarse pixels, *not* the fine pixels, are the data finally transmitted back to the NGDC for compositing. The averaging of side-by-side, 5 x 5 grids produces some smoothing or smudging of data values. Figure 2.3 illustrates this aggregation of data for the running example of Figures 2.2 and 2.3.

While imagery ordered from NGDC for any single night are therefore available only at

Time sequence of scans over a light source off nadir




Figure 2.3: Misattribution of light in consecutive scans

Time sequence of scans over a light source off nadir
Re-binned into 5x5 pixels




Figure 2.4: Conversion of fine to coarse pixels

a coarse pixelization of 2.8km x 2.8km, the annual composite imagery have a much finer pixelization.⁸ This finer pixelization is achieved simply by the fact that coarse pixels from different nights do not exactly stack on each other, but rather partially overlap, so the sum across many nights throughout the year produces a finer pixelization.

Dimensions of overglow in annual composites

From the 3000km-wide swath viewed by the satellite on each orbital circuit, only the middle 1500km of data are retained by NGDC for annual image compositing (Doll (2008)). The outer quarters are discarded by NGDC since, at such extreme viewing angles, significant measurement error is anticipated. Thus for the data ultimately used to produce the annual composite nighttime lights imagery, each light source on the Earth's surface is viewed by a satellite always 0-750km off nadir.

Table 2.1: Ellipses at various displacements off nadir

Off-nadir distance (km)	Major radius (km)	Minor radius (km)
0.0	2.3	2.3
100.0	2.33	2.32
200.0	2.43	2.37
300.0	2.6	2.44
400.0	2.83	2.55
500.0	3.13	2.68
600.0	3.5	2.83
700.0	3.93	3.0
750.0	4.17	3.09

Figure 2.5 depicts the geometry of a ground scanning area x -km off nadir. In Doll (2008) we are told that DMSP satellites orbit at an altitude of 833km and possess ground scanning areas at nadir of radius 2.3km⁹. As is evident from Figure 2.5, the dimensions of the elliptical off-nadir ground scanning area are easily determined by the longitudinal displacement x of the light source from nadir. Table 1 lists the dimensions of elliptical ground scanning areas

⁸Annual nighttime lights composite imagery are composed of pixels of size 1/2 nautical mile times the sine of the latitude.

⁹Ultimately for our sample we find a radius of 2.0km seems to fit the data better, but this is an easily resolved issue.




Figure 2.5: Trigonometry of off-nadir ground scanning areas

0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	.09	0	0	0	0	0	0
0	0	0	2.0	14.1	33.9	46.5	33.9	14.1	2.0	0	0	0
0	0	3.4	25.9	68.4	98.0	100	98.0	68.4	25.9	3.4	0	0
0	.27	15.6	51.9	98.3	100	100	100	98.3	51.9	15.6	.27	0
0	1.08	21.1	63.8	100	100	100	100	100	63.8	21.1	1.08	0
0	.27	15.6	51.9	98.3	100	100	100	98.3	51.9	15.6	.27	0
0	0	3.4	25.9	68.4	98.0	100	98.0	68.4	25.9	3.4	0	0
0	0	0	2.0	14.1	33.9	46.5	33.9	14.1	2.0	0	0	0
0	0	0	0	0	0	.09	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0

Figure 2.6: Overflow frequency distribution (pixel size .867km)

for various displacements from nadir.

At which angles will a given light source be viewed in the course of a calendar year? On each night, the satellite will view the same location at a *different* longitudinal displacement. We assume for simplicity that the set of displacements at which a given light source is viewed is sampled from a uniform distribution. In particular, we suppose that a light source is viewed once at each of displacements 0km, 1km, ..., 750km. At each displacement, a unique ellipse is illuminated (see Figure 2.5). *Every* ellipse is centered at the light source, so the ellipses corresponding to smaller off-nadir displacements are concentrically contained by ellipses corresponding to larger displacements. The intersection of all the ellipses is therefore the on-nadir ground scanning area ellipse (circle).

The geometry of overflow therefore produces a simple and powerful prediction about the frequency with which pixels will be lit: pixels within the on-nadir ground scanning area of a light source (i.e. within 2.3km) will be lit precisely as often as the light source itself. Pixels further away from the light source will be lit *less* frequently. Figure 2.6 depicts the predicted distribution of frequencies for pixels of side-length .867km, for a light source that is *permanently* lit (i.e. is lit on every cloud-free night).

Figure 2.6 provides something of a retrospective justification for the ‘thresholding’ techniques employed by Imhoff et al (1997) and Henderson et al (2003). These researchers hoped to use nighttime lights to calculate urban extent around the world, and recognized overflow error as a serious nuisance. After inspecting the data, they recognized that often the pixels that should not have been lit were lit relatively *infrequently*. Accordingly, they proposed the idea of filtering out pixels that were lit sufficiently infrequently. They proceeded to set various thresholds, dismissing infrequently lit pixels as spurious. This literature faltered when no single threshold could be agreed upon as universally appropriate. Figure 2.6 clarifies why




Figure 2.7: Overflow distribution for 100 units of light (pixel size .867km)

these researchers’ intuition was basically correct, but also why their method failed. Indeed, the pixel containing the permanent light source is lit 100% of the time, so all pixels lit less than 100% of the time are non-light-emitting and should be ignored when measuring urban extent. On the other hand, this filter is not quite enough: pixels adjacent to the light-emitting pixel are *also* lit 100% of the time, yet they are not light-emitting. To further complicate things, light sources in developing economies often suffer power outages. Annual composite imagery will therefore record some light-emitting pixels as being lit less than 100% of the time. A fixed filtration threshold of 100% will therefore erroneously discard such light sources. Rather than using a fixed threshold, a more effective method is to search for *local maxima* in the frequency surface. In particular, the geometry of overflow suggests that a pixel containing a light source will always be lit *more* frequently than adjacent pixels. Identifying local maxima in the frequency surface as a ‘prior’ indicating the likeliest sources of light in the data is a useful way of reducing the dimensionality of the image rectification problem, and we find it speeds up computation considerably.¹⁰ Nevertheless, in the empirical

¹⁰We have tested code for this approach, and it is available upon request.

Schematic of a Photomultiplier




Figure 2.8: The internal function of a photomultiplier tube

section below we allow all lit pixels to be potentially light emitting.

The overflow function

For a given quantity of light emission, we can also depict the overflow pattern predicted by our geometric explanation. Figure 2.7 depicts the overflow pattern for an isolated light source emitting 100 units of light (we temporarily suspend the discussion of topcoding), with pixel side-lengths once again set to .867km. Note that the area covered by positive digital numbers (dns) is smaller than that of Figure 2.6, due to the fact that we bottom-coded all dns less than .05 as zero for presentational purposes. Since the nighttime lights data are themselves bottom-coded during NGDC data processing, the exact elliptical dimensions of Figures 2.6 and 2.7 will seldom appear in the data, since their outer pixels will tend to be dimly and infrequently lit and may be reclassified as dark by NGDC data cleaning algo-

Malfunction of a Photomultiplier




Figure 2.9: Malfunction of a photomultiplier tube

rithms.

Extended overflow

Table 1 indicates that the largest ellipse to be illuminated around any light source will have major radius 4.17km and minor radius 3.09km. Much larger elliptical distortions with radii of 15km or more, however, are evident over the greater Las Vegas area and around gas flares in southeastern Nigeria, among other places. What explains this ‘extended’ overflow? And why do we observe it only around some light sources and not others?

Extended overflow results from a kind of ‘malfunction’ of the photomultiplier tube. Figure 2.8 is helpful to follow. An incoming photon (depicted as a crooked arrow) strikes the photocathode, which is constructed from a rare material that enjoys the peculiar property that, when struck by a photon, it detaches a commensurate quantity of electrons. The de-

tached electrons accelerate toward a relatively positively charged dynode, which is likewise composed of a material such that when struck by an electron, a commensurate quantity of electrons are detached. These detached electrons then accelerate toward the next dynode, and the process snowballs until a sizeable pulse of electrons is accelerating rightward toward the anode. Upon reaching the anode, the size of the pulse is measured, and a digital number is subsequently recorded.

Typically this chain reaction of strikes and detachments is only set in motion by a photon striking the photocathode in the ‘acceptance area’, which corresponds to the smaller ground scanning areas thus far discussed. Occasionally, however (on the order of 10^{-4} of the time), photons striking the photocathode outside the ‘acceptance area’ will nevertheless initiate a chain reaction (see Figure 2.9). For relatively dim light sources, the total quantity of photons striking the photocathode outside the acceptance area is far too small to generate a large enough current of electrons to be recorded as a positive digital number. In a sense, this ‘malfunction’ is bottom-coded out of the data collection process for dim light sources. But for very bright sources such as are found in cities like Las Vegas, there are so many photons striking the photocathode that even the fraction 10^{-4} corresponds to an absolute quantity of photons large enough to register positive digital numbers. We estimate the ground scanning radius of the resulting ‘extended’ overflow to be approximately 41km, implying the largest ellipse’s major radius is 77km. So for very bright cities, pixels as far as 77km away can be erroneously illuminated.

2.3 Modeling and removing overflow

Section 2 established how light is erroneously redistributed by the satellite’s optics and onboard data management. In this section we summarize this redistribution as a linear

transformation. We introduce a simple method for de-blurring nighttime lights imagery.

Stylized 1-dimensional example

To establish intuition, let us first consider a stylized example of a linear city suffering from overglow (see Figures 10 and 11). We observe a row of 5 locations: A, B, C, D, and E. We suppose that a city occupies locations B, C, and D. The city center is at C; suburbs are located at B and D; while locations A and E form a rural periphery around the city. Without overglow error, Figure 2.10 is observed. The city center emits 10 units of light, while the suburbs each emit 5 units. The rural periphery is non-light-emitting.

Now suppose that light is redistributed erroneously at a rate of 40%. For simplicity, let us assume that only *adjacent* pixels are affected. Then of C's 10 units of emitted light, 4 are redistributed: 2 to B; 2 to D. Of B's 5 units of emitted light, 1 is redistributed to A, and 1 to C. And of D's 5 units of emitted light, 1 unit is redistributed to C, and 1 unit to E. The resulting image is depicted in Figure 2.11.

Figure 2.11 neatly depicts several key problems with nighttime lights data owing to overglow error. First of all, as has long been observed by the remote sensing literature (see, for example, Imhoff et al (1997) and Henderson et al (2003)), overglow error tends to make nighttime lights a dubious data source for measuring urban extent, since it systematically exaggerates city boundaries. An economics researcher observing Figure 2.11 would be in danger of concluding that the city sprawls across 5 locations, whereas in reality it only occupies 3. On the other hand, an economics researcher seeking to calculate the total amount of light emitted by the city, and possessing an accurate city boundary polygon that identifies locations B, C, and D as the city proper, would be in danger of under-counting the city's total light by *missing* the light stored at locations A and E. She would calculate the city's total light emission as $5+5+8=18$, a 10% underestimate. Finally, an urban economics researcher




Figure 2.10: Stylized linear city *without* overflow

studying intracity decentralization by comparing levels of economic activity in locations B and C would find erroneously that the ratio of light emission of a suburb to the city center is 5/8 instead of 5/10, perhaps concluding incorrectly that the city's economic activity is substantially decentralized.

Modeling overflow in a linear city

We can model the observed light in Figure 2.11 by a system of 5 equations:

- (1) $y_A = \rho \cdot \beta_B$
- (2) $y_B = \beta_B - 2\rho \cdot \beta_B + \rho \cdot \beta_C$
- (3) $y_C = \beta_C - 2\rho \cdot \beta_C + \rho \cdot \beta_B + \rho \cdot \beta_D$
- (4) $y_D = \beta_D - 2\rho \cdot \beta_D + \rho \cdot \beta_C$
- (5) $y_E = \rho \cdot \beta_D$

Where y_i is the total light (digital number) observed by the researcher for pixel i ; β_i is the light emitted at pixel i (unknown); and ρ is the (unknown) overflow parameter. Note also




Figure 2.11: Stylized linear city *with* overflow

that this system satisfies the ‘conservation of energy’ constraint:

$$y_A + y_B + y_C + y_D + y_E = \beta_B + \beta_C + \beta_D$$

This constraints states that the total light observed in the image equals the total light emitted at Earth’s surface; this emphasizes that overflow is an erroneous *redistribution* or blurring

of light. The correct quantity of light is recorded, but assigned to incorrect locations.

We write the above system in matrix form:

$$\begin{aligned}
\begin{pmatrix} y_A \\ y_B \\ y_C \\ y_D \\ y_E \end{pmatrix} &= \begin{pmatrix} \rho & 0 & 0 \\ 1-2\rho & \rho & 0 \\ \rho & 1-2\rho & \rho \\ 0 & \rho & 1-2\rho \\ 0 & 0 & \rho \end{pmatrix} \cdot \begin{pmatrix} \beta_B \\ \beta_C \\ \beta_D \end{pmatrix} \\
&= (1-2\rho) \begin{pmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix} \cdot \begin{pmatrix} \beta_B \\ \beta_C \\ \beta_D \end{pmatrix} + \rho \cdot \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} \beta_B \\ \beta_C \\ \beta_D \end{pmatrix} \\
&= (1-2\rho) \cdot W_0 \cdot \beta + \rho W_1 \beta \\
&= [(1-2\rho)W_0 + \rho W_1] \cdot \beta \\
&= X(\rho) \cdot \beta
\end{aligned}$$

The left-hand-side vector y of digital numbers is observed, while the 3×1 parameter vector β and scalar ρ are unknown. The 5×3 neighbor matrices W_0 and W_1 perform the task of selecting potentially-light-emitting pixels to be part of each pixel's light decomposition. The first and last rows of W_0 , for example, are zero vectors, because the corresponding pixels A and E are non-light-emitting. The second row of W_0 , however, corresponds to pixel B, and so it has a 1 in the first column, because pixel B is potentially light-emitting. Since pixel B is not pixel C or D, there are zeros entered at positions (2, 2) and (2, 3). Meanwhile, the W_1 matrix selects potentially-light-emitting pixels *adjacent* to each pixel. In the first row, for example, it is registered that pixel A has one potentially-light-emitting neighbor (B). In the third row, corresponding to pixel C, two potentially-light-emitting neighbors are recorded




Figure 2.12: Elliptical bands

(B and D).

General solution

To generalize, consider a system of N pixels, with $K < N$ of them potentially emitting light. Then the $N \times 1$ vector y of digital numbers is observed, and the $K \times 1$ vector β of true light emissions is unobserved. In general, light may be redistributed not only to adjacent pixels, but also to pixels further away. We take a sequence of ellipses $0 < e_1 < \dots < e_L$ and define the l th elliptical band of pixel i to be the region of space lying between the l th and $(l - 1)$ th concentric ellipses centered at the centroid of i (see Figure 2.12). Then the (i, j) th entry of the l th $N \times K$ neighbor matrix W_l indicates whether LE pixel j lies in the l th elliptical band

of pixel i . A fixed percentage $M_l \rho_l$ of light is lost from j to its l th distance band. If i lies in that band, it receives a fraction ρ_l of j 's light (alternatively stated, pixel i receives $\frac{1}{M_l}$ th of all the light transferred to band l from pixel j). Then observed light may be decomposed as

$$(1) \quad y = [(1 - \sum_{l=1}^L M_l \rho_l) W_0 + \sum_{l=1}^L \rho_l W_l] \beta = X(\rho; M) \cdot \beta$$

So that the $N \times K$ matrix $X(\rho; M)$ is determinate conditional on vectors $\rho = (\rho_1 \ \rho_2 \ \dots \ \rho_L)$ and $M = (M_1 \ M_2 \ \dots \ M_L)$. In section 2 we derived exactly what the elliptical distortion should be for a pixel emitting any quantity of light (see Figures 2.6 and 2.7 for the case where 100 units of light are emitted). Figure 2.7 is a nonparametric representation of the overflow function, where values are binned by pixels. We can now re-bin Figure 7 by elliptical bands, assigning the various partial overlaps of pixels to the correct bands, and calculate the M and ρ vectors. In short, the geometric derivation of overflow in Section 2 allows us to make the X matrix determinate. In the image rectification literature, the X matrix is known as a ‘blurring’ function. The geometry of the satellite’s movement and optics allow us to define the blurring function.

We see from (1) that the observed image y is simply the true image β transformed by a blurring operator X . The true image β minimizes the sum of squared residuals

$$(2) \quad \beta = \operatorname{argmin}_{\hat{\beta}} = SSR(\hat{\beta}) = u'(\hat{\beta})u(\hat{\beta}) = [y - X\hat{\beta}]'[y - X\hat{\beta}]$$

We can estimate β with standard nonlinear optimization packages.

Correcting extended overflow

Extended overflow occurs in precisely the same manner as its shorter-ranged counterpart, just that the ground scanning area at nadir is much larger. We estimate the radius of the

at-nadir radius to be 41km, implying that the largest elliptical distortions have a major radius of about 77km (east-west).¹¹ Because extended overflow only affects light systems containing very bright pixels, the conservation-of-energy constraint is violated: the total light recorded around dim light sources is less-than-proportional to what is recorded around bright sources, because the extended overflow of dim sources is very small and is therefore bottom-coded out of the data. As a consequence of this fact, we do not handle extended overflow by reassigning light as in the short-ranged case; instead, we subtract this light out of the image altogether, in order to bring bright sources back into proportion with dim sources.

Computational challenges arise since a much larger quantity of pixels are affected and therefore must be included in the minimization problem. Since the problem is one of exponential complexity, there is a more-than-proportional increase in computations. For this paper, we study systems of light that are relatively small and topologically separable from other light systems. For light systems that can be contained within a circle of radius of 41km, all pixels are equally affected by extended overflow, which means we can remove extended overflow simply by applying a threshold filter. In particular, we take the observed image and subtract a fixed quantity of radiance from all pixels. If the resulting pixel values are below zero, we reset them to zero. For relatively large systems of light, this thresholding technique does not quite capture the nuance that some pixels are receiving more extended overflow than others. We leave it for future research to develop a computationally efficient method for handling these larger systems.

2.4 Empirical Results

We remove overflow from radiance-calibrated nighttime lights images F15-2000 and F16-2010 over 9 sub-Saharan African cities; Kabul, Afghanistan; and Kandy, Sri Lanka. We

¹¹Calculations available upon request.

begin by drawing a polygon around each city such that all pixels lit by city light sources in both F15-2000 and F16-2010 lie inside the polygon’s boundary. Figures 2.13 and 2.14 depict Addis Ababa in F15-2000 and F16-2010, respectively. The turquoise boundary around the main city defines the polygon.

Next, we filter out extended overflow using a simple threshold, as justified in Section 3. By visual inspection, setting a threshold of 8 radiance units seems to reduce extended overflow sufficiently well throughout our data, in both sub-Saharan Africa and South Asia.¹² In Figures 2.13 and 2.14, lit pixels of radiance 8 or less are colored grey.

We then remove short-range overflow by implementing the de-blurring method of Section 3 in a series of Python and Matlab scripts. The Python scripts download and organize the geographic content of the imagery, while the Matlab script generates the SSR function and runs a standard Matlab nonlinear optimization method to minimize it. Additional Python scripts then reconstruct point feature class data of the Matlab results.

Table 2.2: Urban extent and light over MODIS polygons in 2000; raw vs. de-blurred images

City, country	Ratio of lit areas raw vs. overflow-corrected	Fraction of city’s raw F15-2000 light in MODIS extent
Kabul, Afghanistan	1.87	0.47
Nakuru, Kenya	12.86	0.24
Kisumu, Kenya	13.59	0.13
Nairobi, Kenya	6.91	0.29
Bida, Nigeria	8.53	0.22
Bauchi, Nigeria	7.97	0.2
Illorin, Nigeria	14.59	0.26
Maiduguri, Nigeria	6.25	0.32
Kandy, Sri Lanka	10.06	0.02
Dire Dawa, Ethiopia	13.73	0.34
Addis Ababa, Ethiopia	7.63	0.62
AVERAGE	9.45	.28

Figures 2.16 and 2.17 depict the overflow-corrected F15-2000 and F16-2010 images for Ad-

¹²The only city where this does not seem to be the case is Kabul, Afghanistan. We do not apply any filtering threshold there.




Figure 2.13: Addis Ababa, Ethiopia; F15-2000 radiance-calibrated data




Figure 2.14: Addis Ababa, Ethiopia; F16-2010 radiance-calibrated data




Figure 2.15: Addis Ababa, Ethiopia; F15-2000 rad-cal and MODIS urban extent

Table 2.3: Intensive vs. Extensive Growth, 2000-2010; raw vs. de-blurred images

City	Intensive growth (raw)	Intensive growth (overflow-corrected)	Extensive growth (raw)	Extensive growth (overflow-corrected)
Kabul, Afghanistan	3.95	5.96	7.94	8.04
Nakuru, Kenya	0.78	0.76	1.13	0.83
Kisumu, Kenya	0.73	0.74	1.15	0.87
Nairobi, Kenya	0.89	0.96	1.03	1.13
Bida, Nigeria	0.59	0.48	0.84	0.69
Bauchi, Nigeria	0.79	1.15	1.33	1.51
Illorin, Nigeria	1.41	1.87	1.44	2.64
Maiduguri, Nigeria	0.95	0.91	1.13	0.74
Kandy, Sri Lanka	0.89	1.03	1.01	1.06
Dire Dawa, Ethiopia	0.67	0.91	1.11	11.36
Addis Ababa, Ethiopia	0.8	1.04	1.88	4.22
AVERAGE	1.13	1.44	1.82	3.01




Figure 2.16: Addis Ababa, Ethiopia; F15-2000 radiance-calibrated; *de-blurred*

dis Ababa, which should be contrasted with the raw data in Figures 2.13 and 2.14. The overflow-corrected imagery are much sharper, no longer exhibiting elliptical smoothing of the raw imagery. Lit pixels are also more diffuse in Figure 2.17 than in Figure 2.16, suggesting that Addis Ababa decentralized over 2000-2010. As further visual evidence of this, note that 5 pixels in Figure 2.17 achieve the highest category of brightness (376-500 radiance units) while in Figure 2.16, none of the pixels belongs to this category.

Economic statistics: raw vs. overflow-corrected imagery

With overflow-corrected imagery in hand we can finally answer the question of whether overflow error really matters for empirical economic research. Are economic statistics calculated from raw nighttime lights imagery meaningfully different from those calculated from overflow-corrected imagery? We assemble a set of summary statistics of all our cities in




Figure 2.17: Addis Ababa, Ethiopia; F16-2010 radiance-calibrated; *de-blurred*

Tables 2 and 3.

To begin with, we compare estimates of urban extent. In Table 2, the first column of statistics presents for every city the ratio of lit pixels in raw versus overglow-corrected imagery. Consistent with what we observed for Addis Ababa, the total lit area of each city is vastly overstated in the uncorrected, raw image. On average, raw imagery overstate city size by a factor of 9.45, a number which corresponds well to the findings of Imhoff et al (1997) and Henderson et al (2003), where raw-lights-based estimates of urban extent were compared to Landsat-based estimates.

Economists using nighttime lights typically obtain polygon data of their spatial units of interest (county borders, urban extents of cities, etc), then calculate total light over each polygon. But even if polygons accurately delineate urban extents of cities, economists will

under-count total light of the city by failing to add light spilt across the city’s boundaries. To quantify this problem, we obtain an accurate, non-lights-based measure of urban extent: the MODIS-based IGBP land classification raster created by Friedl et al (2010).¹³ These data define urban extent (impervious surfaces) at a resolution of 500 meters, and are accurate for the year 2001. In Figure 2.15, the MODIS-derived urban extent (purple) of Addis Ababa is superimposed on the F15-2000 raw image. Column 2 from Table 2 lists for each city the fraction of total F15-2000 light that lies *inside* the MODIS polygon, i.e. the fraction of the city’s total light that would actually end up being counted by economists. On average, only 28% of the city’s total light lies inside the city polygon. So if the city’s GDP is calculated as a linear transformation of light (see Henderson et al (2012)), then economists will underestimate the city’s GDP by 72% on average.

Urban economists want to distinguish intensive and extensive growth of cities: is the city becoming more economically active over the same areas it has historically occupied, or is the city decentralizing? Perhaps most notable in this area is the recent work of Henderson et al (2014) and Baum-Snow and Turner (2014) on the decentralization of Chinese cities in response to road and railway expansions. In Table 3, we present estimates of intensive and extensive growth of our sample cities, contrasting results from raw and overflow-corrected images. We define the city’s historic, ‘core’ area as the region delineated by the MODIS-based polygons. Increases to total light over these polygons is counted as *intensive* growth, while increases to total light outside these polygons is counted as *extensive* growth. In the first two columns of Table 3, we compare intensive growth in raw versus overflow-corrected images. In the raw image of Addis Ababa, we find light inside the MODIS polygon is only 80% in 2010 of what it was in 2000, i.e. there is a 20% absolute decline in light inside the MODIS polygon. In the overflow-corrected image, however, light in 2010 is 104% of its 2000 level, suggesting the city’s core grew slowly, but did not actually *decline*. In the third and

¹³These data are freely available from the LPDAAC website as the MCD12Q1 data product https://lpdaac.usgs.gov/products/modis-products_table/mcd12q1

fourth columns we calculate extensive growth as the total light outside the MODIS polygon in 2010 divided by the total light outside the MODIS polygon in 2000. In the raw image, light in 2010 is 188% of its 2000 level. In the overflow-corrected image, it is 422%. Qualitatively, a researcher equipped only with the raw imagery would conclude that economic activity had fled the city center and had grown at an annual average rate of 6.5% from 2000-2010. With the overflow-corrected imagery, a researcher would say instead that economic growth had stagnated in the city center, but in the city’s periphery activity had increased at an extraordinary rate of 15.5% annually. Indeed, Table 3 indicates that on average, raw imagery understate extensive growth by 119 percentage-points relative to overflow-corrected imagery (1.82 versus 3.01), though there are several cities in the sample (Nakuru, Kisumu, and Maiduguri) where raw imagery indicate a decline in central areas and growth peripherally, while overflow-corrected imagery indicate decline in both central and peripheral parts of the city.

2.5 Conclusion

Nighttime lights data have become a popular proxy for economic activity at various levels of spatial resolution, and have filled a major information gap, particularly in the developing world where accurate and spatially disaggregated measures of economic activity are less frequently available. This paper analyzes an important type of measurement error inherent in these data, called ‘overflow’. We carefully review the origins of overflow error in DMSP-OLS nighttime lights data, identifying it as a problem engendered by the optics and onboard data processing of the satellites themselves. We derive the geometry of overflow and introduce a simple method by which the imagery can be rectified. We implement the method in Python and Matlab and apply the code to a set of nighttime lights images over a sample of sub-Saharan African and South Asian cities. Overflow-corrected imagery exhibit dramatically improved resolution, confirming that overflow is an important source of error and that our

correction method is a useful tool for researchers to enhance imagery before proceeding with economic analysis. Future work can apply the methodology established in this paper to improve analysis of urbanization, urban sprawl, and to enhance understanding of the role of infrastructure in urban decentralization.

2.6 References

- Abrahams, A. 2014. Hard Traveling: Commuting and Welfare in the Second Intifada. Working paper.
- Agnew, J., Gillespie, T.W., Gonzalez, J. and Min, B. 2008. Baghdad Nights: Evaluating the US Military 'Surge' Using Nighttime Light Signatures. *Environment and Planning A*, 40, 2285-95.
- Alesina, A., S. Michalopoulos and E. Papaioannou. 2013. Ethnic Inequality. Working Paper 155, Centre for Competitive Advantage in the Global Economy, Department of Economics, University of Warwick
- Angel, Civco, Sheppard. 2005. The Dynamics of Global Urban Expansion. World Bank.
- Baum-Snow, Brandt, Henderson, Turner, Zhang. 2014. Roads, Railroads and Decentralization of Chinese Cities. Working paper.
- Baum-Snow, Turner. 2014. Transportation and the Decentralization of Chinese Cities. Working paper.
- 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.
- Croft, T.A., 1979. The Brightness of Lights on Earth at Night, Digitally recorded by DMSP Satellite. U.S. Geological Survey.
- Doll, Christopher N.H., Jan Peter Muller, and Jeremy G. Morley. 2006. Mapping regional economic activity from night-time light satellite imagery. *Ecological Economics* 57: 75-92
- Doll, Christopher N.H. 2008. CIESIN Thematic Guide to Night-time Light Remote Sensing

and its Applications. Columbia University.

Elvidge, C. D. , P. Cinzano, D. R. Pettit, J. Arvesen, P. Sutton, C. Small, R. Nemani, T. Longcore, C. Rich, J. Safran, J. Weeks and S. Ebener. 2007. The Nightsat Mission Concept. *International Journal of Remote Sensing*. Vol 28 (12), pp 2646-2670.

Elvidge, C. D., Baugh, K. E., et al. 1998. Radiance calibration of DMSP-OLS low-light imaging data of human settlements. *Remote Sensing of Environment*, 68, 7788.

Friedl, M. A., Sulla-Menashe, D., Tan, B., Schneider, A., Ramankutty, N., Sibley, A., and Huang, X. 2010. MODIS Collection 5 global land cover: Algorithm refinements and characterization of new datasets. *Remote Sensing of Environment*, 114, 168182.

Ghosh, T., S. Anderson, R. L. Powell 1, P. C. Sutton, and C. D. Elvidge. 2009. Estimation of Mexico's Informal Economy and Remittances Using Nighttime Imagery. *Journal of Remote Sensing*, 418-444.

Hardy, Morgan. R, Haraati. 2014. Illuminating the Shadow Economy: Using lights from space to estimate a new measure of informal economic activity. Working paper.

Henderson, M., E. T. Yeh, P. Gong, C. D. Elvidge, K. Baugh. 2003. Validation of urban boundaries derived from global night-time satellite imagery, *International Journal of Remote Sensing*, 24:595-609.

Henderson, J. Vernon, Adam Storeygard, and David N. Weil. 2012. Measuring Economic Growth from Outer Space. *American Economic Review*, 102(2): 994-1028.

Imhoff, M.L., W. T. Lawrence, D. C. Stutzer, C. D. Elvidge. 1997. A Technique for Using Composite DMSP-OLS 'City Lights' Satellite Data to Accurately Map Urban Areas, *Remote Sensing of Environment*, 61:361-370

Letu, Hara, Tana, Nishio. 2012. A Saturated Light Correction Method for DMSP/OLS Nighttime Satellite Imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 50(2).

Levin, Duke. 2012. High spatial resolution night-time light images for demographic and socio-economic studies. *Journal of Remote Sensing of Environment*, 119, 1-10.

Lieske, R. W. DMSP primary sensor data acquisition. 1981. *Proceedings of the International*

Telemetering Conference.

Pinkovskiy, Maxim. 2013. Economic Discontinuities at Borders: Evidence from Satellite Data on Lights at Night. Working paper.

Small C., Pozzi F., Elvidge C. D. 2005. Spatial analysis of global urban extent from DMSP-OLS night lights. *Remote Sensing of Environment*. 96, 277-291.

Storeygard, Adam. 2014. Farther on down the road: transport costs, trade and urban growth in sub-Saharan Africa. Revise & resubmit, *Review of Economic Studies*.

Tuttle, Benjamin T., Anderson, S. J., Sutton, P., Elvidge, C. D., Baugh, K. 2012. It Used To Be Dark Here: Geolocation Calibration of the Defense Meteorological Satellite Program Operational Linescan System. Working paper.

van der Weide, Roy. B. Rijkers, B. Blankespoor, A. Abrahams. 2014. The Merits of Market Acces. Working paper.

Zhang, Q., Schaaf, C., Seto, K. 2013. The Vegetation Adjusted NTL Urban Index: A new approach to reduce saturation and increase variation in nighttime luminosity. *Remote Sensing of Environment*, 129, 32-41.

Chapter 3

How Valuable is Market Access? Evidence from the West Bank

3.1 Introduction

How valuable is market access? Answering this question is difficult because market access typically only evolves slowly and non-randomly over time which makes it challenging to isolate its impact on economic performance. Exploiting variation in the provision of transportation infrastructure (Baum-Snow et al., 2013, Donaldson, 2010, Jedwab and Moradi, 2014) placement of borders (Redding, 2008), or transportation costs (Storeygard et al. forthcoming) existing studies typically find sizeable gains to reducing barriers to mobility. Consistent with the results of these studies, transport infrastructure projects account for a sizable share of international development assistance. For example, in 2011, investment in transportation projects alone accounted for 20% of the World Bank's project portfolio.

This paper contributes to the literature on the returns to reducing trade frictions by examining the impact of the easing of mobility restrictions imposed by Israel in the West Bank on local GDP growth, proxied for by nighttime lights data. Mobility restrictions took the form

of manned and unmanned physical road obstacles, including roadblocks, checkpoints, earth mounds, trenches, and a separation barrier wall, and were part of the broader closure regime initially instituted by the Israeli army in response to the Second Palestinian Uprising. We exploit changes in the composition and configuration of these obstacles between 2004 and 2012 as a quasi-natural experiment to assess the marginal returns to market access.

This context enjoys several qualities that aid identification. Firstly, the deployment of road obstacles was arguably exogenous to local Palestinian economic conditions. As referenced in Abrahams (2015), obstacles' purported objective was (is) to safeguard the security of Israeli settlers. We provide evidence in our regression analysis that indeed our obstacle-derived market access variable is assigned in a quasirandom manner. Fatalities data and international export data allow us to test if omitted variables related to conflict or international trade are driving results. Our regression results remain robust to the inclusion of these covariates, suggesting that our market accessibility index is quasirandomly determined.

Deployment of obstacles varied over short time horizons and across locations, which facilitates identification of the impact of variation in market access on development. Mobility restrictions dramatically intensified in response to the Second Intifada in the early 2000s, after which they were gradually relaxed, especially during 2009-2010. Our sample period, which runs from 2004 up until 2012, thus coincides with an overall easing of mobility restrictions.

Based on detailed data gathered from official sources and interviews with UN-OCHA officials, we are able to estimate the time cost of traversing each type of obstacle. Since UN-OCHA also provided average travel speeds for all road segments in the West Bank, we are able to calculate the minimum travel time between any pair of locations in our dataset via the road network, and to recalculate based on any configuration of obstacles observed through-

out the 2004-2012 period. Using these calculations we construct a comprehensive index of market accessibility (see e.g. Deichmann, 1997), that takes into account both the precise geolocations of the obstacles, traversal times, and road quality. This is important since one strategically placed checkpoint can reduce accessibility more severely than a multitude of roadblocks when alternative connections are available. In addition, restrictions may be mutually reinforcing; and market access in a given locality may in part be impacted by obstacles in relatively remote areas, the impact of which conventional proxies for mobility restrictions, such as simple counts of the number of obstacles, would fail to capture. The accessibility index is appealing as it provides an intuitive, uni-dimensional, encompassing metric for the evolution of market access. To our knowledge, this is the first attempt to construct finely spatially disaggregated market access measures varying over relatively short time intervals.

Using nighttime lights satellite data we are able to calculate an annual panel of light emissions for all Palestinian locations. Lights are now widely accepted as a credible proxy for GDP and GDP growth (Doll et al (2006), Chen & Nordhaus (2011), Henderson et al (2012)), and recent research has relied on lights in the absence of spatially disaggregated GDP measures (Alesina et al (2012), Storeygard (2015)). This paper, along with Abrahams (2015), is the first to use nighttime lights to study the Israel-Palestine Conflict, and indeed the first to use satellite data of any kind to study the same. The paper is also the first to employ the interannual lights calibration technique of Wu et al (2012) and the optical de-blurring algorithm of Abrahams et al (2015) to clean the data before use.

Previewing our main findings, the paper documents a strongly positive relationship between market accessibility and local economic growth which is both economically and statistically significant. Taken at face value, our preferred estimates imply that a 10% relative increase in accessibility, achieved by the removal of obstacles along relevant road segments, will on

average cause a 1% increase to economic growth (as proxied for by nighttime lights growth).

In generalizing our main findings, it is important to recognize they identify the marginal impact of changes in mobility restrictions, which arguably underestimate the total cost associated with mobility restrictions. Identification is based on comparing the relative performance of different localities and firms within the West Bank and does not account for adaptation and mitigation measures already undertaken. For example, one may speculate that the most footloose firms would relocate in response to severe restrictions. On the other hand, our estimates are based on the removal of restrictions in an environment where previously there were none. This may result in higher estimates of the marginal returns to market access than would accrue in the short-run if one were to expand market access in an unconstrained environment, as reaping the returns to better access may require complementary investments.

In documenting how relatively short-run fluctuations in market access impact economic performance, the paper builds on and contributes to different strands of literature. To start with our results complement existing studies of the returns to reducing trade frictions (Anderson and van Wincoop, 2004, Arkolakis, 2012). The results of the paper also resonate with the growing New Economic Geography literature on the agglomeration externalities (see e.g. Fujita et al, 1999) and demonstrate how imposing restrictions on trade impede the potential gains from agglomeration from being capitalized on. Moreover, it adds to the considerable body of evidence on the (inadvertent economic costs of) the Israeli-Palestine conflict (see e.g. Zussman and Zussman, 2006, Jaeger and Paeserman, 2008). To the extent that prosperity propels pacification, the institution of mobility restrictions may thus inadvertently induce a tradeoff between short-term and long-term security objectives.

Four papers are especially closely related to our work. Storeygard (2015) constructs a panel

of night time lights data for countries in Sub-Saharan Africa and exploits variation in oil prices to identify the impact of transportation costs on economic development. Abrahams (2015), in a closely related study to ours, focuses on how Israeli obstacles disrupted Palestinian commuting, causing employment rates to decline in labor-supply locations but to increase in labor-demanding locations. Cali and Miaari (2013) document a negative association between local employment rates and mobility restrictions in the short-run. Etkes & Zimring (2013) uses the blockade of Gaza as a natural experiment to quantify the gains from trade by using the West Bank as the counterfactual.

The remainder of this paper is organized as follows. The next section discusses nighttime lights data. Section 3.3 discusses the construction of the accessibility index, which is our key independent variable. Section 3.4 presents regression results.

3.2 Nighttime Lights Data

Our main outcome variable of interest is the total quantity of nighttime light (NTL) emissions recorded over each Palestinian West Bank location for each year in 2004-2012. NTL data for the period 1992-2012 were recorded by the US Air Force's Defense Meteorological Satellites Program (DMSP), whose fleet of polar-orbiting satellites have scanned the Earth's surface since the early 1970s, recording nightly imagery of the globe. The National Oceanic and Atmospheric Administration (NOAA) has processed these imagery and composited them into global images where each pixel records the average quantity of light observed at that location across all cloud-free nights of the year.¹

The work of Doll et al 2006, Chen and Nordhaus 2011, and Henderson et al 2012 draws upon country-level GDP data to establish that DMSP-NTL data proxy credibly for GDP and GDP growth. Stated simply, brighter countries tend to have higher GDP, and changes

¹The data are available for free download from <http://ngdc.noaa.gov/eog/download.html>




Figure 3.1: Light over Palestinian West Bank locations, 1992-2012

in brightness tend to correlate positively with changes in GDP. Since GDP data are often absent at provincial and municipal levels of developing or conflicted regions, economists have seized upon NTL data as a convenient proxy.² In the context of the West Bank, Levin & Duke (2012) and Abrahams (2015) find strong evidence that nighttime lights reflect economic realities on the ground.

In Figure 3.1 we plot a time series of total NTL over Palestinian locations for 1992-2012. The raw time series are depicted in green, and are calculated by summing the NOAA-recorded digital numbers (dns) for each pixel overlapping or partially overlapping UN polygon extent data of all Palestinian locations in the West Bank.³ Since the DMSP satellites were designed in the 1970s with the intention of monitoring nightly weather patterns by spotting moonlit clouds, the satellites can only distinguish 6 bits of brightness (values 0-63). The raw digital numbers are *not* proportional to the number of photons observed over each location, so raw dns require ‘radiance calibration’ (this issue has been overlooked in all

²See Storeygard (2015), Alesina et al (2012), Squires (2015), Baum-Snow & Turner (2014), Turner & Gonzalez-Navarro (2015).

³Polygon data were obtained from the UN-OCHA.




Figure 3.2: West Bank F162006, blurred image

previous economic papers using NTL data). Ours is the first study to apply the radiance intertemporal calibration of Wu et al (2012), which transforms digital numbers of all images 1992-2012 via a power function into common units of radiance, where the parameters of the power function are determined separately for each satellite-year by an ‘invariant region method’, and are shown to be superior to previous attempts at intertemporal calibration. When summed over all Palestinian location polygons, the resulting data generate the time series depicted in blue.

A well-known drawback of the DMSP-NTL data is that they suffer from blurring, commonly dubbed ‘overflow’ or ‘blooming’, where emitted light is redistributed erroneously to neighboring pixels. This spatially autocorrelated error is a nuisance for any research where areas of interest are adjacent and/or small.⁴ In the West Bank, where bright Israeli settlements lie proximate to dimly lit Palestinian villages, and where Palestinian towns lie close together, spillage of light potentially introduces a lot of error. Along with Abrahams (2015), this paper

⁴For example, Pinkovskiy’s (2011) comparison of lights near national borders is challenged by the spillage of light from the ‘richer’ to the ‘poorer’ country, which biases downward estimates of the income gap.




Figure 3.3: West Bank F162006, de-blurred image

is the first to employ the de-blurring algorithm of Abrahams et al (2015), which attempts to remove blurring by performing an inverse filter based on the shape of the point-spread function determined via an analysis of the DMSP satellites’ optical system and the geometry of the satellites’ cross-track data collection. In the case of the West Bank, inspection of the raw data confirms that the entire territory is illuminated, and inspection of frequency data (indicating the percentage of cloud-free nights on which each pixel was lit) confirms that most of the territory is lit 100% of the time. This type of illumination is discussed thoroughly in Abrahams et al (2015), where it is dubbed ‘extended overglow’. Extended overglow results from an artefact of the DMSP’s optical system, coupled with high overall brightness of nearby cities Tel Aviv and Jerusalem. To remove extended overglow, the inverse filtration algorithm is preceded by a subtraction of a fixed quantity of light (8 units of radiance) from all pixels.⁵ Figures 3.2 and 3.3 provide a before-and-after comparison for the F162006. In Figure 3.2 we see that the entire West Bank is illuminated due to short-ranged blurring from

⁵The subtraction of 8 radiance units from all pixels means that the red curve is shifted downward significantly from the blue curve. Subsequent inverse filtration means that the slopes of the blue and red curves do not quite match. They are very similar, however, owing to the fact that some Palestinian locations lose light due to blurring, while others gain, so that when summed the net effect across all locations is small.

local light sources and extended blurring from Jerusalem and Tel Aviv. In Figure 3.3 we see the de-blurred image, after inverse filtration has been applied.

The resulting intercalibrated and de-blurred NTL data summed across all Palestinian locations is graphed as the red time series in Figure 3.1. The shape of the red curve is powerful evidence that NTL data have significant economic content. The dramatic increase of NTL over Palestinian locations during 1993-1999 corresponds to a period of hope following the Oslo Accords (1993), when it seemed as if the peace process would successfully give rise to a two-state solution and an end to conflict. The decline in NTL starting in 1999 corresponds likewise to a period of disillusionment in the peace process and the outbreak of the Second Uprising (September, 2000). It is not until the final, calmer years of the uprising (2005-2007) that lights once again reach their 1999 level. That the overall light level actually declined should dispel skepticism that light merely tracks population changes. During the 1997-2007 period, Palestinian fertility caused a massive population increase in the West Bank, from 1.5 million to 2 million, so the decline in lights during 1999-2002 does not reflect population decline but rather the decline of economic fortunes precipitating and resulting from the uprising.

3.3 Accessibility data

Obstacles

Responding to Palestinian militant attacks originating from within the West Bank in the uprising's early years, the Israeli army moved to defend Israeli civilians by deploying physical obstacles along roads and borders. In an effort to prevent Palestinian suicide attackers from entering Israel, a 500km wall or 'separation barrier' was built along (and oftentimes beyond) the 1967 Armistice Line of the West Bank. Henceforth, all cross-border traffic was forced to pass through any one of a dozen 'green-line checkpoint' border crossings, vigilantly




Figure 3.4: West Bank roads and Israeli army obstacles

guarded by the Israeli army (see, for example, the report by B'Tselem: “Ground to a Halt” (2007)). Throughout the uprising, however, nearly a quarter million Israeli civilian settlers were dwelling deep inside the West Bank, beyond the protection of the wall. To defend those settlements, the army deployed hundreds of manned checkpoints, roadblocks, boulders, earthmounds, and other physical obstacles along the West Bank’s internal road network with the intention of monitoring or otherwise discouraging Palestinian traffic along roads passing in the vicinity of Israeli settlements. See Abrahams (2015) for further discussion and references.

The UN-OCHA Map Center did an excellent job of recording the progress of the wall’s construction, and keeping track of road obstacles’ geolocations, updating their maps many times throughout and after the uprising as obstacles were moved and removed in response




Figure 3.5: Nablus: before and after the lifting of the blockade

to shifting strategic and political tensions. The left panel of Figure 3.4 displays a map constructed from UN-OCHA GIS data depicting major arteries of the West Bank's internal road network. The map also identifies the Palestinian Authority's 11 West Bank governorate capitals. The right panel zooms in on a northwestern section of the West Bank, depicting locations of Israeli army obstacles lying along roads connecting the governorate capitals of Tulkarm, Nablus, and Ramallah. The fact that obstacles were deployed on some roads and not others provides cross-sectional (spatial) variation in the degree to which each Palestinian location was obstructed by obstacles.

Temporal variation in obstacle deployment

Many obstacles were relocated or altogether removed during and after the uprising. Figure 3.5 compares obstacle deployment in 2009 versus 2010 in the area around Nablus, Palestine's historical economic capital and the West Bank's second largest city after Hebron. Nablus was blockaded during and after the uprising, but the blockade was finally lifted during the

2009-2010 period. The deployment and removal of obstacles provides temporal variation in obstruction, allowing us to compare the economic outcomes of Palestinian locations before and obstacles were removed.

Travel Times

The time required to traverse a segment of road depends on several factors, including whether or not the road is paved, road width (number of lanes), and the presence and type of Israeli army obstacles deployed along that road segment. From the UN we received estimates of the average amount of time required for a typical civilian vehicle to traverse each segment of road in the map *without any obstacles deployed*, i.e. the pre-uprising travel time for each road segment. We then used ArcGIS 10.3's Network Analyst software package to solve for the optimal pre-uprising route and minimum travel time between all origin and destination pairs of interest. All 792 Palestinian West Bank locations recorded by UN-OCHA were included as origins. Travel times from each of these origins to 18 distinct destinations were calculated, including the 11 West Bank Palestinian governorate capitals, 6 border crossings into Israel, and the King Hussein Bridge to Jordan.

The immediate effect of obstacle deployment was to increase travel times between origin-destination pairs by increasing the time required to traverse road segments wherever obstacles were present. Working in coordination with the UN, van der Weide & Blankespoor (2014) employed survey workers to monitor and record crossing times at various checkpoints around the West Bank in 2009. Extrapolating from this sample, we obtain the time required to traverse different types of obstacles throughout the West Bank, 2004-2012. Using the UN obstacle maps available for 2004-2012, we recalculate optimal routes and minimal travel times between every origin-destination pair in the presence of obstacles. Figure 3.6 provides a box plot of average travel times between our set of origins and destinations, for each year 2004-2012. Notably, average travel times are high in the latter years of the uprising (2005-




Figure 3.6: Travel times to governorates & border crossings

2006) and after (2007-2008), eventually declining in 2009-2010 as the Israeli army lifted its blockade of various cities, including Nablus.

Consistent with NGO reports cited in Abrahams (2015), our calculations suggest obstacles greatly increased travel times between Palestinian origins and destinations. For example, traveling from Hebron in the southern West Bank, we calculate a pre-uprising journey to Nablus would have taken 2 hours and 17 minutes. In 2005-2008, however, the same journey would have taken around 7 hours (!). After the lifting of the blockade in 2009, travel times declined to ‘just’ 4 hours and 40 minutes, and to less than 4 hours in 2010.

Accessibility Indices

It is intuitive that from the perspective of a producer at some origin location, not all possible destinations are of equal market importance: their size and distance help determine their relevance. *Ceteris paribus*, destinations farther away are less relevant, since the price of




Figure 3.7: Ease of access to Palestinian governorate capitals

the exported good will be higher there (endogenizing iceberg transit costs), and therefore less competitive with local substitutes. And destinations with smaller populations are less relevant, since they contain fewer consumers. For these reasons, we factor in travel time and a measure of the destination's 'size' when calculating each origin's ease of access to export markets.

For the 11 governorate capitals, we weight by their population (i.e. number of consumers) in year t , inverse-weighting by an exponential function of the minimum travel time required to reach them. The sum of these gives the 'internal' access to market index for origin i (henceforth 'internal accessibility index'):

$$intMA_{it} = \sum_j P_{jt} \exp(-T_{ijt}) \quad (3.1)$$

Where P_{jt} is destination j 's population in year t and T_{ijt} is the minimum travel time between




Figure 3.8: Ease of access to Israeli border crossings

i and j in year t . Figure 3.7 box-plots the natural-log of Palestinian locations' internal accessibility indices for 2004-2012. Consistent with the travel times plotted in Figure 3.6, internal accessibility improves during 2008-2010 when many Israeli obstacles were removed.

To quantify Palestinian locations' access to international markets, we draw on PALTRADE data to weight each crossing by the number of truckloads of material exported through that crossing in year t . Inverse-weighting again by an exponential function of the minimum travel time required to reach these crossings, we obtain each origin i 's 'external accessibility index':

$$extMA_{it} = \sum_j P_{jt} \exp(-T_{ijt}) \quad (3.2)$$

Where P_{jt} is destination j 's outgoing truckloads of produce in year t . Figure 3.7 box-plots internal accessibility indices of all Palestinian origins for 2004-2012, while Figure 3.8 box-plots their external accessibility indices. Note that PALTRADE data on exported truckloads at

border crossings were available only for 2006-2011, and only for 3 of a dozen border crossings, so this variable is likely noisier than $intMA_{it}$, increasing the probability of attenuation bias and the possibility of a non-result.

3.4 Regression Analysis

Taking Palestinian locations as our units of analysis, we attempt to identify the causal effect of market accessibility on annual per-capita changes in locations' observed light output. Our main specification is

$$\Delta \ln(LPC_{it}) = \beta_{int} \ln(intMA_{it}) + \beta_{ext} \ln(extMA_{it}) + \beta_X X_{it} + \beta_L \ln(LPC_{it-1}) + \beta_{\Delta L} \Delta \ln(LPC_{it-1}) + u_i + \tau_t + \epsilon_{it} \quad (3.3)$$

Where the outcome variable LPC_{it} is the total lights per-capita observed at location i in year t ; $intMA_{it}$ quantifies the ease with which producers at location i may export their goods to the consumer markets of Palestinian governorate capitals in year t ; $extMA_{it}$ quantifies the ease with which producers at location i may export their goods to international markets in year t ; X_{it} is a vector of time-varying control variables for location i ; LPC_{it-1} and ΔLPC_{it-1} are the lagged level and change in lights per-capita; u_i captures permanent features of location i that affect the annual growth rate of light output; and τ_t is a year t effect. The parameters of interest are β_{int} and β_{ext} .

Table 3.1 presents results of running specification (3). So as to ensure that results are not driven by very small villages or temporary bedouin encampments, the regressions exclude all locations whose population in 2007 was 5000 residents or fewer. This leaves just 100 locations. Alternatively, all 792 locations can be included but observations should be weighted by their 2007 population (see Table 3.2). All standard errors are robust to heteroskedasticity. Location effects and year effects are left unreported due to space constraints.

Table 3.1: Effects of accessibility on per-capita lights, 2004-2012; 100 largest towns

VARIABLES	(1) $\Delta \ln(\text{lights_pc})$	(2) $\Delta \ln(\text{lights_pc})$	(3) $\Delta \ln(\text{lights_pc})$	(4) $\Delta \ln(\text{lights_pc})$
ln(intMA)	0.0968*** (0.0221)	0.0864*** (0.0216)	0.0957*** (0.0331)	0.105*** (0.0330)
ln(extMA)			-0.000428 (0.00169)	0.00114 (0.00176)
lagged ln(lights_pc)	-0.362*** (0.0365)	-0.380*** (0.0364)	-0.420*** (0.0416)	-0.421*** (0.0415)
lagged $\Delta \ln(\text{lights_pc})$	-0.185*** (0.0365)	-0.185*** (0.0358)	-0.104** (0.0439)	-0.101** (0.0434)
Pal_fatal_dummy_10k		-0.0375*** (0.0144)		-0.0445*** (0.0145)
Isr_fatal_dummy_10k		0.00662 (0.0117)		0.0162 (0.0118)
checkpoint_dummy_5km		-0.0260*** (0.00710)		0.00940 (0.00855)
constant	-2.495*** (0.298)	-2.414*** (0.291)	-2.621*** (0.436)	-2.707*** (0.431)
Observations	900	900	599	599
R-squared	0.412	0.427	0.598	0.610

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 3.2: Effects of accessibility on per-capita lights, 2004-2012; all locations, weighted

VARIABLES	(1) $\Delta \ln(\text{lights_pc})$	(2) $\Delta \ln(\text{lights_pc})$	(3) $\Delta \ln(\text{lights_pc})$	(4) $\Delta \ln(\text{lights_pc})$
ln(intMA)	0.0832*** (0.0158)	0.0686*** (0.0148)	0.0709** (0.0297)	0.0731*** (0.0272)
ln(extMA)			-0.00232 (0.00152)	-0.00147 (0.00153)
lagged ln(lights_pc)	-0.325*** (0.0261)	-0.339*** (0.0265)	-0.414*** (0.0247)	-0.414*** (0.0249)
lagged $\Delta \ln(\text{lights_pc})$	-0.184*** (0.0280)	-0.187*** (0.0274)	-0.120*** (0.0246)	-0.121*** (0.0246)
Pal_fatal_dummy_10k		-0.0307*** (0.0109)		-0.0332*** (0.0115)
Isr_fatal_dummy_10k		-0.00120 (0.0129)		0.00733 (0.0101)
checkpoint_dummy_5km		-0.0221*** (0.00651)		9.09e-05 (0.00828)
constant	-1.752*** (0.202)	-1.595*** (0.190)	-1.705*** (0.351)	-1.707*** (0.323)
2007-population-weighted	Yes	Yes	Yes	Yes
Observations	4,194	4,194	2,792	2,792
R-squared	0.387	0.398	0.599	0.605

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The results in Column (1) suggest that a 10% relative increase in a location's ease of access to West Bank internal consumer markets (governorate capitals) would on average increase the location's growth rate of lights by 1% in relative terms. An accessibility increase could be driven by differential population growth, where more easily accessible markets grow faster; or by the removal of Israeli army obstacles, as occurred in 2009-2010.

The point-estimate for the coefficient on $\ln(intMA)$ in Column (1) may be biased if there are omitted variables driving both accessibility and changes in lights. For example, checkpoints proximate to a location may decrease its access to consumer markets while also contributing to the increase in incidence of violent confrontation, which in turn may harm economic growth, decreasing the rate of growth of light. Column (2) includes several conflict-related binary variables, including whether or not checkpoints lie within a 5km radius of a location in year t , whether or not Palestinian fatalities occurred in year t within a 10km radius, and similarly for Israeli fatalities. These variables do not shift the point-estimate on $\ln(intMA)$ by more than 1 standard deviation, suggesting that omitted conflict variables are not driving results.

Column (3) repeats Column (1)'s regression but includes our locality-level measure of access to international export markets as described in Section 3.3. This index relies on border crossing data regarding truckloads of produce flowing out of the West Bank. These data are missing for 2004-2005 and 2012, so the number of observations is reduced by 1/3 from Column (1) to (3). The point-estimate of the coefficient on $\ln(intMA)$ and its statistical significance are stable across Columns (3) and (4), despite the loss of 1/3 of observations and the inclusion of $\ln(extMA)$. The results of these latter columns suggest that the point-estimate of β_{int} is not biased by failure to account for changing access to international

markets.

Table 3.2 repeats the regressions of Table 3.1, but includes the full panel of 792 Palestinian locations, weighting observations by their 2007 population. Results across all columns of Table 3.2 are very similar, and likewise similar to results in Table 3.1.

3.5 Conclusion

This paper seeks to quantify the benefits of market access by exploiting temporal and spatial variation in the degree to which Israeli army obstacles increased travel times for the delivery of exported goods from Palestinian towns to Palestinian governorate capitals during and after the Second Palestinian Uprising. We merge GIS data on road quality and geolocations of Israeli obstacles and average traversal time for each type of obstacle, then calculate optimal routes and minimum travel times throughout the years 2004-2012. Along with population data for the governorate capitals, we construct an accessibility index for each Palestinian location, and run panel regressions to recover the marginal effect of accessibility on locations' economic growth. In the absence of traditional growth measures, we assemble a unique dataset of nighttime lights, performing novel interannual calibration and de-blurring techniques to clean the data. Regressions reveal a positive relationship between market accessibility and local economic growth rates as proxied for by nighttime lights: a relative increase of 10% in accessibility causes on average a 1% relative increase in light growth rates. Robustness checks indicate the results are driven neither by statistical outliers, omitted variables related to conflict, nor changes in access to international markets. Results suggest that Israeli army obstacles did indeed stymie Palestinian economic growth in the West Bank during and after the uprising. Going forward, policymakers should weigh the costs to the Palestinian economy that these obstacles impose against the safety they provide to Israeli citizens.

3.6 Bibliography

- Abadie, A., and J. Gardeazabal (2003). The Economic Costs of Conflict, *American Economic Review*, 93(1): 113-132.
- Abrahams, Alexei (2013) Mobility, Employment, and Income, Evidence from an Israeli counter-terrorism policy Mimeo.
- Anderson, J. E., and E. van Wincoop (2004): Trade Costs, *Journal of Economic Literature*, 42(3), 691-751.
- Arkolakis, C., A. Costinot, and A. Rodriguez-Clare (2012): New Trade Models, Same Old Gains?, *American Economic Review*, 102(1), 94-130.
- Baum-Snow, Nathaniel (2007) Did Highways Cause Suburbanization? *The Quarterly Journal of Economics*, Vol. 122, No. 2, pp. 775-805, 05.
- Donaldson, D. (2010). Railroads of the Raj: Estimating the Impact of Transportation Infrastructure, NBER Working Paper No. 16487.
- Deichmann (1997). Accessibility Indicators in GIS. United Nations, New York.
- Fujita, Krugman, Venables. *The Spatial Economy: Cities, Regions and International Trade* Cambridge: MIT Press, 1999.
- Henderson, Vernon, Adam Storeygard, and David N. Weil (2012) Measuring Economic Growth from Outer Space,” *American Economic Review*, Vol. 102, No.2, pp. 994-1028.
- Jedwab Remi and Alexander Moradi, 2014. ”Transportation Technology and Economic Change: The Impact of Colonial Railroads on City Growth in Africa,” Working Papers 2014-03, The George Washington University, Institute for International Economic Policy
- Miaari, S.H. and Sauer, R.M. (2011). The labor market costs of conflict: Closures, foreign workers, and Palestinian employment and earnings. *Review of Economics of the Household*, 9: 129-148.
- Redding, Sturm. The Costs of Remoteness: Evidence from German Division and Reunification. *American Economic Review*, 98(5), 1766-1797, 2008.
- Storeygard, Adam. Farther on down the road: transport costs, trade and urban growth in

sub-Saharan Africa. Revise & resubmit, Review of Economic Studies, 2013.

Jaeger, D. and D. Paserman (2008). The cycle of violence? An empirical analysis of fatalities in the Palestinian-Israeli conflict. American Economic Review 98 (4): 1591-1604.

Zimring, Assaf Gains from Trade Gains from Trade: Lessons from the Gaza Blockade 2007-2010 Mimeo.

Zussman,Zussman. Assassinations:Evaluating the Effectiveness of an Israeli Counterterrorism Policy Using Stock Market Data. Journal of Economic Perspectives, 20(2),193-206 Spring2006.