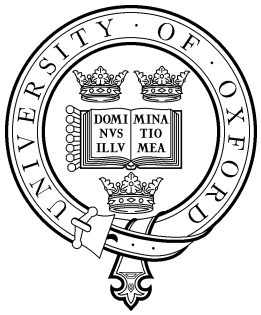
Essays on Natural Resource Economics in Developing Countries



Thomas M J McGregor
St Hugh's College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy in Economics

Trinity 2017

Abstract

This thesis comprises three papers within the field of natural resource economics in developing countries. In broad terms, this thesis investigates the challenges and opportunities surrounding the macroeconomic management of natural resource revenues in developing economies. It is a mix of theoretical and empirical work covering three distinct but related topics.

The first paper, titled "Commodity price shocks, growth and structural transformation in low-income countries", is an empirical investigation of the macroeconomic response of resource-rich, low-income countries to movements in commodity prices. A panel-VAR approach is used to estimate both the dynamic and structural macroeconomic response to global commodity price shocks. The results point to small, but positive, growth and investment responses to commodity price booms. I also find evidence of significant transformation of the sectoral value-added shares in these economies following price shocks, and present evidence of a Dutch disease mechanism at work.

The second paper, titled "Surfing a wave of economic growth", investigates whether the geographic determinants of growth extend to natural amenities. It combines data on spatial and temporal variation in the quality of over 5,000 surf breaks globally with data on local economic performance, proxied by night-time lights. The findings document a strong association between natural amenity quality and local economic development. Economic activity grows faster near good surf breaks; following the discovery of new breaks, or the technology needed to ride them; and during El Niño events that generate high-quality waves. The effects are concentrated in nearby towns and emerging economies, and population changes are consistent with tourism.

The third paper, titled "Pricing sovereign debt in resource-rich economies", investigates the link between commodity price movements and risk premiums in resource-dependent, developing economies. I develop a stochastic general equilibrium model of a resource-rich economy with sovereign debt and endogenous default risk. I show that the bond price is decreasing in the level of indebtedness and negatively correlated with the oil price. The model can explain a large proportion of the business cycle fluctuations in interest rate spreads in resource-dependent, developing economies. Higher risk-aversion, more impatience, larger oil shares and a stronger correlation between domestic output and oil price shocks all help to explain the key macroeconomic co-movements in these settings.

Declaration of authorship

Essays on Natural Resource Economics in Developing Countries

Thomas McGregor
St Hugh's College
University of Oxford

Thesis submitted for the degree of Doctor of Philosophy in Economics
Hilary Term 2017

Chapter 1: *Introduction*
This chapter is wholly my own work.

Chapter 2: *Commodity price shocks, growth and structural transformation in low-income countries*
This chapter is wholly my own work. It is based on a paper published recently in the Quarterly Review of Economics and Finance (McGregor (2017)).

Chapter 3: *Surfing a wave of economic growth*
This chapter is based on an earlier paper written jointly with Dr. Samuel Wills (University of Sydney). The paper is currently under review at the Economic Journal (McGregor & Wills (2017)). We both contributed to the production of the paper in equal shares. We developed the initial idea for the paper jointly. I gathered the majority of the data, including the initial data on surf locations from the wannasurf website, carried out the empirical estimations in sections 3.3, 3.4 and 3.5 and contributed to the drafting of the paper.

Chapter 4: *Pricing sovereign debt in resource-rich economies*
This chapter is wholly my own work. It is based on a recent OxCarre working paper (McGregor (forthcoming)).

Chapter 1: *Concluding remarks*
This chapter is wholly my own work.

No part of this thesis has been submitted for a degree at this, or any other, university.

Acknowledgements

First and foremost, I would like to thank my two outstanding supervisors, Professor Frederick van der Ploeg and Professor Christopher Adam for their continued guidance, advice and support throughout the years. Their knowledge of their field and passion for research has had a huge influence on my work and kept me motivated throughout the D.Phil. Without them this thesis would never have progressed further than a few muddled research ideas. Thanks also to my friend and co-author Sam Wills who has been a great colleague and mentor to me throughout the years.

I wish to thank Anthony Venables, Sophocles Mavroidis, Andrea Fererro, Stefan Dercon, Francis Teal, Simon Quinn, Doug Gollin, Tim Willems and Julia Bird, amongst others, for insightful discussions and advice at various stages of the D.Phil.

It would be impossible to thank everyone who has helped and supported me over the years, but a special thanks goes to all of my friends and fellow D.Phil colleagues who have made my time here at Oxford so enjoyable, to my fellow teammates and friends of the OUBC for two wonderful campaigns, and to the Grad IV.

My family has been an unwavering support throughout my life. I am eternally grateful to my parents, Magda and Roy, who together have encouraged, supported and inspired me every step of the way, to my siblings for keeping my feet on the ground, and to Amanda for putting up with me through all of the highs and lows along the way.

I would also like to take this oppertunity to thank my examiners, Professor Andrea Ferrero and Professor Jonathan Temple, for their helpful comments, suggestions and discussion of this thesis.

I acknowledge the generous financial support from the Department of Economics and the Oxford Centre for the Analysis of Resource Rich Economies (OxCarre).

Contents

1 Introduction **1**
- 1.1 Setting the scene . . . 1
- 1.2 Growing dependence and increased volatility . . . 2
- 1.3 Early work on resource dependence and growth . . . 5
- 1.4 Short run commodity price movements . . . 8
- 1.5 Zooming in on natural assets . . . 10
- 1.6 A specific volatility channel . . . 13

2 Commodity price shocks, growth and structural transformation in low-income countries **17**
- 2.1 Introduction . . . 17
- 2.2 Empirical methodology . . . 22
  - 2.2.1 Structural panel-VAR . . . 23
  - 2.2.2 Seemingly unrelated panel regressions . . . 29
  - 2.2.3 Weighted price index . . . 31
- 2.3 Data . . . 33
  - 2.3.1 Sources . . . 33
  - 2.3.2 Summary statistics . . . 35
- 2.4 Results . . . 38
  - 2.4.1 Cholesky . . . 39
  - 2.4.2 Block recursive . . . 41
  - 2.4.3 Robustness . . . 42
  - 2.4.4 Sectoral value added results . . . 44
- 2.5 Dutch disease mechanism . . . 48
- 2.6 Nigeria and the oil price collapse . . . 50
- 2.7 Conclusion . . . 51

3 Surfing a wave of economic growth **54**
- 3.1 Introduction . . . 54
- 3.2 Data and measurement . . . 60
  - 3.2.1 Surf breaks . . . 60
  - 3.2.2 Night-time lights . . . 63
  - 3.2.3 Population . . . 67
  - 3.2.4 Urban and rural classification . . . 68
  - 3.2.5 Political stability and business environment . . . 68
  - 3.2.6 El Niño and wave heights . . . 68
- 3.3 Experiment I: Cross section . . . 69

3.3.1 Identification . . . 69
3.3.2 Estimating equations . . . 70
3.3.3 Results . . . 72
3.4 Experiment II: Discoveries . . . 81
3.4.1 Identification . . . 81
3.4.2 Estimating equations . . . 83
3.4.3 Results - discoveries and disappearances . . . 84
3.4.4 Results - Cold water breaks . . . 85
3.5 Experiment III: Large waves . . . 85
3.5.1 Identification . . . 86
3.5.2 Estimating equations . . . 88
3.5.3 Results and discussion . . . 89
3.6 Conclusion . . . 91

4 Pricing sovereign debt in resource rich economies 94
4.1 Introduction . . . 94
4.2 Emerging market spreads . . . 98
4.3 The model . . . 101
4.4 Quantitative analysis . . . 105
4.4.1 Solving the Model . . . 105
4.4.2 Base calibration . . . 106
4.5 Pricing sovereign debt . . . 108
4.6 Model performance . . . 111
4.6.1 Independent shocks . . . 111
4.6.2 Covarying shocks . . . 114
4.6.3 Impatience and oil dependence . . . 118
4.6.4 Recursive preferences . . . 121
4.6.5 Final model . . . 124
4.7 Conclusion . . . 126

5 Concluding remarks 128

Appendix A General 132
A.1 Figures . . . 132
A.2 Tables . . . 137

Appendix B Commodity price shocks 138
B.1 Figures . . . 138
B.2 Tables . . . 141

Appendix C Local impacts of resource booms 145
C.1 What makes a good wave? . . . 145
C.2 Overglow . . . 147
C.2.1 Test 1: Wasteland boundaries . . . 147
C.2.2 Test 2: Isolated light sources . . . 148
C.3 Experiment I: Additional results . . . 153
C.3.1 Surfing and economic activity . . . 153
C.3.2 Ease of doing business and political stability categories . . . 154

C.4 Experiment I: Robustness 157
C.4.1 Pixel-level analysis 157
C.4.2 Omitted Geographic Characteristics 161
C.4.3 Selection 161
C.5 Experiment II: Additional Results 166

Appendix D Sovereign debt dynamics **169**
D.1 Technical appendix 169
D.1.1 Ramsey representation 169
D.1.2 Bivariate Markov shocks 178
D.2 Appendix 181
D.2.1 Solution algorithm 181
D.2.2 Oil export share calibration 181
D.2.3 Model sensitivity 184

Bibliography **186**

List of Figures

1.1 Global primary commodity price indices 4
1.2 World petroleum export shares, 1995-2015 4
1.3 Global oil and gas price 14

2.1 World commodity prices 37
2.2 Panel-VAR IRFs - Cholesky identification strategy 39
2.3 Panel-VAR IRFs - block recursive identification strategy 41

3.1 Global surf breaks 61
3.2 Surf breaks - top countries 62
3.3 Surf breaks by geographic type 63
3.4 Night-time lights and population data 64
3.5 PPP-adjusted GDP vs Night-time lights 65
3.6 Illumination growth in southern Morocco 66
3.7 NOAA Southern Oscillation Index. 69
3.8 North Shore Oahu, Hawaii 71
3.9 Growth regression, PDL model 5km 74
3.10 Wave heights 87

4.1 Emerging market economic performance 100
4.2 World oil price 100
4.3 Bond pricing schedule 111
4.4 Covarying oil price and growth 115

4.5 Sensitivity to correlation - $\rho_{zh} \in [.01, .10]$ 118
4.6 Sensitivity to impatience, $\rho \in [.01, .05]$, and oil dependence, $\bar{s} \in [0.02, 0.6]$. 121
4.7 Sensitivity to risk aversion, $\sigma \in [2, 100]$, and IES, $\xi \in [0.1, 0.5]$ 124

A.1 Global commodity price indices (non-fuel) 132
A.2 Oil reserves (bb) 133
A.3 Gas reserves (tcf) 133
A.4 Mineral reserves (Tril usd) 134
A.5 Natural Gas proved reserves 135
A.6 Crude oil proved reserves 136
A.7 Wealth and per capita wealth by type of capital and income group, 1995 and 2005 137

B.1 Panel VAR IRFs - Cholesky - IRFs with 95% confidence bands 138
B.2 Panel VAR IRFs - Cholesky - no OPEC members 139
B.3 Panel VAR IRFs - block recursive - no OPEC members 139
B.4 Panel VAR IRFs - Cholesky - low and middle income countries 140
B.5 Panel VAR IRFs - block recursive - low and middle income countries 140
B.6 Panel VAR IRFs - block recursive - fixed exchange rate 140
B.7 Panel VAR IRFs - block recursive - flexible exchange rates 141
B.8 Panel VAR IRFs - Cholesky - exposure levels 141
B.9 Panel VAR IRFs - block recursive - exposure levels 141

C.1 Bay of Biscay trench 146
C.2 Global wastelands 147
C.3 Gas flares, Algeria 149
C.4 Overglow comparison 150
C.5 Polynomial regressions 153
C.6 Endogenous town location - Cornwall, UK 154
C.7 Wave quality and experience 154
C.8 Pixel level, 1km and 5km 157

C.9 Pixel regressions 159
C.10 Wave location breakdown 161
C.11 Pixel regressions - discoveries 166
C.12 Pixel regressions - disappearances 167

D.1 Oil and gas shares over time 182
D.2 Sensitivity to correlation - $\rho_{zh} \in [.01, .10]$ ($\bar{s} = .35$) 183
D.3 Sensitivity to impatience, $\rho \in [.01, .05]$ ($\bar{s} = .35$) 184
D.4 Output truncation under default 184
D.5 Bond pricing schedule (no oil) 185

List of Tables

2.1 SUREG FE regression - all sectors 46
2.2 IV regression - real exchange rate channel 50
2.3 Nigeria projections vs outturn 51

3.1 Surf break quality 62
3.2 Total growth regression, 5km 73
3.3 Growth regressions in developing countries 75
3.4 Growth in nearby towns 78
3.5 Population regressions 79
3.6 Population regressions, closest & largest town 80
3.7 Discoveries and disappearances 83
3.8 Discoveries - diff-in-diff regressions 84
3.9 Growth regressions, El Niño years 90

4.1 Model Parameters - base 106
4.2 Oil and gas shares 108
4.3 Model parameters - stochastics 112
4.4 AR(1) model: autocorrelations and standard dev 113
4.5 AR(1) model: cross correlations with real interest rate spread 113
4.6 Cross AR(1) estimates - ρ^{YP} 116
4.7 Model parameters - best fit 125
4.8 AR(1) model: autocorrelations and standard dev 125
4.9 AR(1) model: cross correlations with real interest rate spread 125

A.1 Petroleum export share 137

B.1 Summary statistics 142
B.2 SUR model - all sectors - no controls 143
B.3 SUR model - all sectors - national controls, no OPEC 144

C.1 Wasteland regressions 148
C.2 Light growth regressions, 5-10km 151
C.3 Light growth regressions, 10-50km 152
C.4 Surf breaks by World Bank Doing Business Index 155
C.5 Surf breaks by World Bank Political Stability index 156
C.6 Pixel regressions, with controls 160
C.7 Breakdown of surf breaks by type 161
C.8 Growth regressions, without Australia or US 162
C.9 Growth regressions, with 2-star base 163
C.10 Pixel level growth regressions, with initial light controls 164
C.11 Discovery regressions, diff-in-diff 165
C.12 Cold surf breaks 168
C.13 Discovery of cold surf breaks 168

D.1 AR(1) model: autocorrelations and standard dev ($\bar{s} = .35$) 182
D.2 AR(1) model: cross correlations with real interest rate spread ($\bar{s} = .35$) . . 183
D.3 AR(1) model (no oil): autocorrelations and standard dev 185
D.4 AR(1) model (no oil): cross correlations with real interest rate spread . . . 185

Chapter 1

Introduction

1.1 Setting the scene

Natural resource wealth offers countries an opportunity to boost growth and foster sustainable economic development. For developing and emerging economies, the key is to transform their natural assets, including subsoil hydrocarbon and mineral deposits, into surface assets such as physical, human and institutional capital which in turn generate economic growth. The list of resource rich economies that have successfully managed this transformation includes countries of different income levels, political and institutional models and geographic settings, such as Canada, the United States and Norway as well as Botswana, Malaysia and Chile. For many others however, this process has proven difficult. Countries like Bolivia, Nigeria and Venezuela are cases in point.

The collapse of the oil price to around 30 US$ per barrel in Feb 2016 sparked renewed policy interest in the challenges of managing natural resource wealth. Many countries have become more dependent on commodity exports as subterranean reserves continue to be discovered. The importance of other types of natural assets in the development process is also beginning to be investigated by the economics profession. Add to this the volatile and unpredictable nature of commodity prices as well as the important role government policy plays in the resource sector in many countries, and the question of how best to manage natural resource wealth for development is more pertinent today than it has ever been.

Much of the evidence on the impact of natural resource wealth on economic growth and development is based on cross-country data. Drawing inference from cross-country analysis is difficult given the identification challenges that exist in these type of data. As a result, the debate around whether resource wealth constitutes a curse or a blessing is still ongoing. More recently, scholars have begun investigating a host of mechanisms linking resources and development, using novel datasets, creative identification strategies and state of the art macroeconomic modelling techniques.

This thesis comprises three distinct papers on the economics of natural resource dependency. In broad terms, it investigates the challenges and opportunities surrounding the macroeconomic management of natural resource revenues in developing economies. The remainder of this chapter introduces the areas of focus of this thesis. Chapters 2, 3 4 contribute to what has been a rich literature in the field of economics. Chapter 5 concludes and contributes some closing remarks.

1.2 Growing dependence and increased volatility

Natural assets are an important part of the world's capital stock and comprise both renewable and non-renewable assets. The World Bank (2011) estimates that in 2005, natural assets made up 5% of the global capital stock, and 30% in developing countries.[1] These estimates comprise a number of renewable natural assets, some of which are harder to value that others. Data on the stock of non-renewable natural resources, such as hydrocarbon and mineral reserves, are more readily available, and given the existence of a market price, are also easier to value and quantify.

Due to continued exploration, estimated reserves of hydrocarbon and mineral assets have grown steadily over the past 25 years. The stock of proven oil and gas reserves has increased by more than 2.5 times over this period (see figures A.2 and A.3). These discoveries

[1] See table A.7. This comprises agricultural land, protected areas, forests, minerals, and energy.

have been particularly important for developing countries. Proven oil reserves in Africa, for example, increased from around 57 billion barrels (bb) in 1980 to 126bb in 2016 (a rise of more than 2.2 times), whilst in Central & South America they increased from 25bb to an astounding 331bb (a rise of more than 13 times). Gas discoveries in these two regions have increased by a factor of 2.8 and 3.8 respectively.[2] The value of global subterranean mineral reserves has also risen substantially (see figure A.4).

The market value of these assets however, has been highly variable. Figure 1.1 plots real monthly non-fuel and energy price indices between Jan 2000 and Feb 2017. Between 2000 and the start of the financial crisis in late-2007, commodity prices doubled in real terms. After a sharp acceleration through mid-2008, prices collapsed as the crisis developed. After a period of recovery through mid-2011, which brought prices close to their mid-2008 peak, commodity prices began a long and steady decline as demand from emerging and developing countries slowed. This decline accelerated through mid-2014, bottoming out in 2016 close to the early-2000s level. The spot price of a barrel of Cushing crude oil fell from its peak of 134 US$ in Jul 2008 to a low of 30 US$ in Feb 2016. The collapse of world commodity prices has led some economists to call this the end of the so called commodities 'super-cycle'.[3]

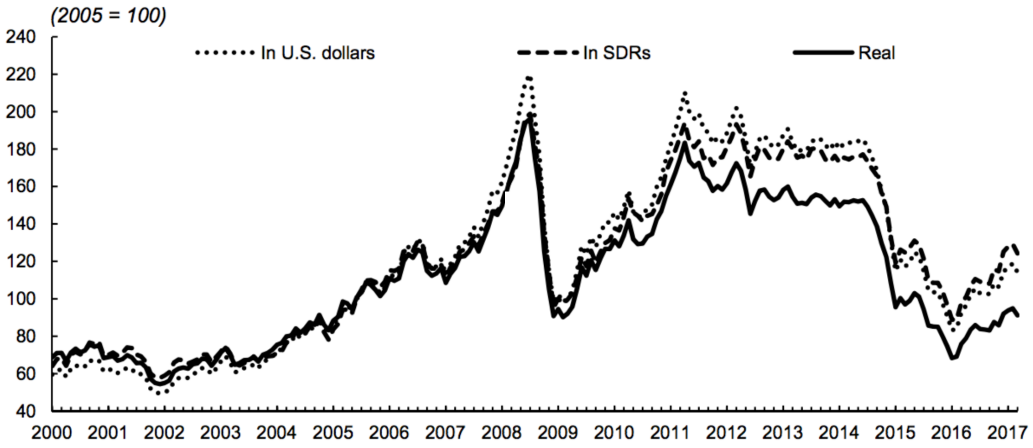
Many countries have also become more dependent on hydrocarbon exports. Figure 1.2 plots the kernel density estimates (KDEs) of the probability density functions (PDFs) of petroleum export shares around the world in 1995, 2005 and 2015. Whilst the number of countries with low dependence on petroleum exports (less than 10% of exports) has fallen since the mid-90s, the number of countries that are moderately dependent (greater than 10% of exports) has risen. According to data from the US Energy Information Administration (EIA), the proportion of countries in which oil makes up more than 10% of total exports rose from around 30% in 2000 to 40% in 2015[4].

[2]Oil reserves in North America increased by a factor of 3.13 between 1980 and 2016 whilst those in the Middle East increased by a factor or 2.2 over the same period. Gas reserves increased by a factor of 3.8 in the Middle East and 3.15 in the Asia & Oceania region.

[3]The term commodities 'supercyle' is often defined as a decades-long, above-trend movement in a wide range of base material prices (Erten & Ocampo (2013); Heap (2005); Rogers (2004)).

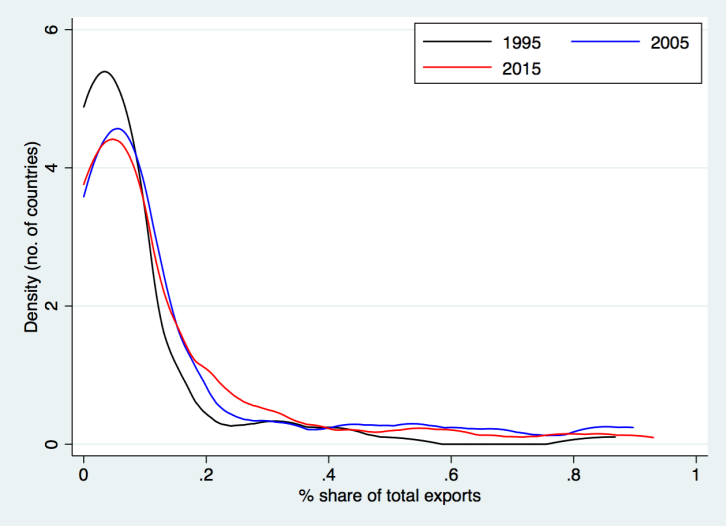
[4]Heavily dependent oil exporters (countries in which oil accounts for more than half of total exports) have fallen from 14.2% to 9.8% over the same period.

Figure 1.1: Global primary commodity price indices



Source: Combines indices of non-fuel primary commodity prices and energy prices.
Real indices deflated by U.S. CPI. IMF primary commodity price index charts, April 20, 2017

Figure 1.2: World petroleum export shares, 1995-2015



Source: US EIA International Energy Statistics, May 10, 2017. Own calculations

1.3 Early work on resource dependence and growth

Pioneering work by Gelb (1988) found that episodes of dramatic oil windfalls in six oil-dependent, low income countries, did not lead to sustained economic development.[5] Later work by Sachs & Warner (1995) formalised empirically a negative link between natural resources and economic growth in a cross-section of countries, showing that resource rich countries in the 1970s grew about one percentage point less per annum, compared with resource poor countries, for the next twenty years. This result holds, even after controlling for initial income per capita, investments during the period, openness, and rule of law, and is particularly strong for countries with weak institutions (Sachs & Warner (1997b); Mehlum et al. (2006b)). This negative correlation between natural resources and economic growth has become known as the 'resource curse' (originally coined by Auty (1993)).

A flurry of research extended and qualified these early findings across a number of important dimensions (see Frankel (2010) and van der Ploeg (2011) for reviews). The definition and measurement of 'natural resources' turns out to be quite important in arguing whether or not they can be treated as exogenous (Manzano & Rigobon (2001); van der Ploeg (2011); James (2015)). Allowing for country fixed effects in these early panel regressions, reduces the effect of resources on growth to zero, showing that resource dependence is likely to be correlated with unobservable country characteristics (Manzano & Rigobon (2001)). Using a measure of subterranean resource wealth results in a positive relationship between resource abundance and growth (Brunnschweiler & Bulte (2008)). Instrumenting resource abundance with the exogenous measure of economically recoverable reserves results in neither a positive or a negative relationship, unless one allows for an indirect effect via volatility (van der Ploeg & Poelhekke (2010)).

More recent work on natural resources and growth has attempted to shed light on the specific causal channels at play. The most well-known of these is the 'Dutch disease' in which resource booms lead to a rise in non-tradable demand thus increasing domestic prices and

[5]The striking contrast made in that book between the economic performance of Nigeria and Indonesia following these large oil windfalls is a useful reminder that the resource curse is not necessarily guaranteed.

wages and putting upward pressure on the real exchange rate (Corden & Neary (1982)).[6] The resulting loss in competitiveness can reduce the production of tradables, which are often assumed to be conducive to growth (Rodrik (2004, 2008)), thus depressing the long-run growth potential of the economy. Recent work has found evidence of a reallocation of resources away from the tradable sector following resource booms. Country level cross-section and panel evidence finds that following resource windfalls, countries typically save more, import more and reduce non-resource exports (Harding & Venables (2010); Ismail (2010)). Evidence from Brazil finds that municipalities that discover oil do see an expansion of the service sector and contraction of industrial production, although total non-oil GDP remains unaffected (Caselli & Michaels (2013)). However, as noted previously, the evidence that this reallocation leads to lower output and lower productivity growth in the long run is much weaker.

Another important channel is the quality of institutions and governance in resource abundant settings. Resource wealth has been shown to worsen the quality of governance and weaken institutions through a number of different channels.[7] The quality of institutions is an important factor in determining whether natural resource booms lead to increased productivity and thus growth (Murphy et al. (1989); Mehlum et al. (2006b,a)). Stronger institutions that encourage productive entrepreneurship mean that increased demand via resource revenues leads to increased entrepreneurial activity and thus higher aggregate productivity. Weak institutions, dysfunctional legal systems and low transparency can mean that rent-seeking activities yield a higher return, thus encouraging more rent-seeking and less productive activity following a resource boom. Resource booms have also been shown to hinder democracy (Ross (1999)) and increase the likelihood of violent conflict (Bannon & Collier (2003); Brunnschweiler & Bulte (2009); Lei & Michaels (2014)).

[6] The term Dutch disease was first used by The Economist in 1977 to describe the real exchange rate appreciation and decline in manufacturing in the Netherlands when natural gas was discovered in 1952.

[7] These include: lower government accountability (Isham et al. (2005)), increased incentives for corruption (Sala-i Martin & Subramanian (2013); Arezki & Brückner (2011); Vicente (2010); Bhattacharyya & Hodler (2010)), rent seeking (Torvik (2002); Mehlum et al. (2006a)), vote buying, bribery and an inefficient public sector (Robinson et al. (2006)) and reduced pressure to reform or modernise (Acemoglu & Robinson (2006))

The conventional wisdom for managing resource revenues is based on the permanent income hypothesis (PIH) of Friedman (1957). The typical policy advice involves smoothing consumption by borrowing before a windfall and building up foreign assets in a sovereign wealth fund, or repaying existing debt, during the windfall. For developing countries however, natural resources provide an opportunity to increase the rate of economic development through increased investment in productive physical, human and institutional assets. Broadly speaking, there are three competing principles that developing country governments need to consider when deciding how best to manage resource revenues. The first is that the return on domestic investment in these countries is likely to be substantially higher than offshore saving, due to high levels of capital scarcity and high risk premiums on international borrowing. This means that some proportion of resource income should be invested in domestic assets or used to pay down external debt, in practice via an *investment* fund (van der Ploeg & Venables (2011); van der Ploeg (2012)). The second is that, given the low levels of per capita income today compared to future generations, some proportion of resource revenues should be used to boost consumption immediately. This policy conclusion for developing country settings stands at odds with traditional advice around the setting up of *intergenerational* funds. And third, the more volatile are natural resource revenues the stronger the rationale for precautionary savings in a *volatility* fund (Carroll & Kimball (2008); Van den Bremer & van Der Ploeg (2013)).[8]

There is also a substantial literature on the effectiveness of foreign aid in developing economies which tackles similar questions to the resource literature, including intergenerational transfers, Dutch disease, absorptive constraints and institutional capacity.[9] Large aid transfers are similar in nature to large resource discoveries, making the study of the former useful in considering the latter. Early work by Obstfeld (1999) showed that in standard neoclassical or endogenous growth models, aid is likely to lead to increases in current consumption rather than investment due to the permanent income hypothesis. Adding

[8]See Collier et al. (2010); van der Ploeg & Venables (2012); Wills (2015) for a review of these competing principles for managing resource wealth in developing countries and Venables & Wills (2016); Venables (2016) for discussions of the challenges posed at different stages of the extraction process.

[9]The debate around whether aid works or not is still ongoing. See significant contributions by Boone (1996); Sachs & Warner (1999); Hansen & Tarp (2001); Easterly (2003); Burnside & Dollar (2004)

absorptive constraints on domestic investment increases the incentive to smooth aid transfers inter-temporally (Carter et al. (2015)). Concerns about potential Dutch disease effects of aid, and the corresponding effects on productivity and growth, have been heavily cited (Van Wijnbergen (1984); Adam & O'Connell (2004); Rajan & Subramanian (2011)) whilst others have highlighted the potential for positive spillovers on the supply side (Adam & Bevan (2006)) and growth effects of targeted public investment (McGregor (2011)).

1.4 Short run commodity price movements

A key challenge facing the resource literature has been the identification of exogenous measures of natural resource wealth. Early cross-country work typically used measures of resource 'dependence' - that is the share of resources or commodity exports in output - as the dependent variable. The problem here is that by dividing by output, this measure of resource wealth is no longer independent of the economic policies and institutions that drive both economic growth and the scale of resource production (Manzano & Rigobon (2001)). The initial solution was to use measures of resource 'abundance' - that is a measure of the stock of subterranean reserves (Brunnschweiler & Bulte (2008); van der Ploeg & Poelhekke (2010)).

Movements in commodity prices offer an alternative way to identify changes in countries' natural resource wealth. A number of papers have shown that resource booms linked with rising commodity prices can lead to significant increases in income levels and growth in developing countries over the short run (Deaton & Miller (1995); Raddatz (2007)), albeit with potentially negative effects in the medium to long run (Collier & Goderis (2012)). The crucial assumption here is that global commodity prices are exogenous to local economic conditions in low income countries. Global demand and financial conditions are likely to affect economic conditions in low income countries through a number of channels (including export demand, foreign investment, capital flows and exchange rate dynamics). In addition, commodity price movements themselves are often driven by aggregate demand in industrialised economies (Hamilton (2003); Barsky & Kilian (2004)).

The second chapter of this thesis, titled "Commodity price shocks, growth and structural transformation in low-income countries" aims to make progress on the issues identified in these two strands of the literature. It is based on a paper published recently in the Quarterly Review of Economics and Finance (McGregor (2017)).

I develop a new measure of resource wealth for low-income countries that seeks to address the endogeneity of conventional measures. The approach involves constructing estimates of country specific commodity reserves which are then used to weight movements in global commodity price indices. A panel vector auto-regression (panel-VAR) is then used to orthogonalise movements in these global commodity prices, specifically minerals and hydrocarbons, against global and local economic conditions. The short run macroeconomic and structural response for low income, resource rich economies is investigated. The panel-VAR approach allows for the interdependence between global economic conditions and low income country growth to be controlled for, as well as allowing for a greater degree of simultaneity amongst certain blocks of variables.

To shed light on potential Dutch disease mechanisms at play, I investigate the impact of commodity price shocks on the economic structure of these resource rich, low income countries. Specifically, I ask whether positive shocks lead to de-industrialisation. In an important contribution to the early literature, Corden & Neary (1982) formalised the de-industrialisation process that may result from resource booms. They distinguish between two channels: the 'resource movement' effect, which describes the process by which labour may be pulled out of the traditional manufacturing sector following a boom, and the 'spending effect', which describes the increase in demand for non-tradables following a boom. Using data from the UN on sectoral value added, I find that resource booms lead to significant de-industrialisation. This link appears to operate through the real exchange rate, with resource shocks leading to appreciations which in turn cause a restructuring of the domestic economy.

The results lend support to recent findings of a positive impact of resource booms on growth over the short run, as well as higher government spending and investment. The magnitude of the growth effects are however small, although slightly larger for less developed countries, economies with fixed exchange rate regimes and those that are more dependent on commodity exports. I also find evidence of de-industrialisation following resource booms, which I argue operates through the real exchange rate.

1.5 Zooming in on natural assets

Faced with the same endogeneity concerns that plagued the early growth literature, researchers have begun to shift their focus away from cross-section investigations of resources and growth towards a within-country or local approach (see reviews by Cust & Poelhekke (2015); van der Ploeg & Poelhekke (2017)). The benefit here is that many of the inherently endogenous country characteristics, such as aggregate growth, institutions, monetary and exchange rate policy, as well as many unobservable characteristics, can be controlled for once we focus within a particular country or local area. Moving away from cross-country studies offers new perspectives on the resource curse debate and can help overcome endogeneity concerns.

Within country studies tend to focus on the effects of 'non-renewable' natural resources, like oil, gas and minerals, both in the short run (Aragón & Rud (2013); Caselli & Michaels (2013); Allcott & Keniston (2014), and over the very long run (Dell (2010)). These studies typically find strong evidence for both positive and negative impacts, both at the local level and further afield. Mining and hydrocarbon extraction generates positive labor market and welfare effects on the surrounding economy, through backward linkages and positive demand spillovers in the short run. The continued development of downstream and upstream industries is key for resource booms having long lasting effects on income levels. Investing in local infrastructure, the provision of public goods and labour market flexibility are also important dimensions.

More recent work has begun looking at 'natural assets' and their impact on economic growth and development. Some places enjoy natural advantages over others, which in turn may yield a number of economic advantages. These advantages may be direct, such as rivers or fertile soils, or indirect, such as locations investing in environmental quality with a high amenity value for residents, business or tourists. These indirect effects are called 'natural amenities' in the literature, and there is no consensus on whether they are important for growth (see Waltert & Schläpfer (2010) for a review).

One specific channel through which natural amenities may affect growth is tourism. Tourism involves the export of otherwise non-tradable services which can be an important source of foreign exchange earnings for an economy as well as an engine for direct and indirect employment growth. The industry has been an important and fast-growing sector for many developing countries. The number of tourists visiting developing countries has grown at an annualised rate of 5.6% between 1990 and 2015, compared to 4.1% in advanced economies (UNWTO (2016)).[10] Despite this, the role of tourism in economic development remains a vastly understudied area.[11] Recent work by Faber & Gaubert (2015)) finds evidence of significant local economic gains via tourism activity in areas with offshore islands and white sandy beaches in Mexico, partly driven by significant positive spillovers on manufacturing.

This literatures on natural amenities tackles many of the same questions as the 'geography vs. institutions' debate in development economics. The source of economic growth is argued by some to be due to exogenous geographic characteristics, such as climate, natural resource endowment, disease burden, transport costs, and the ease of diffusion of knowledge and technology (Diamond (1999); Gallup & Sachs (2000); Sachs (2000, 2001); Hsiang & Meng (2015)). In contrast, those in the institutions camp argued that institutional technology, such as the rule of law, property rights and social architecture, is the major driver

[10]Tourism receipts, in real terms, have grown by 6.9% in developing countries over the same period, compared to 6.3% in advanced economies

[11]Some empirical evidence does exist which uses cross-section country data to study the determinants and consequences of tourism, including Eilat & Einav (2004); Sequeira & Maçãs Nunes (2008); Arezki et al. (2009)

of growth in developing counties (North (1989, 1990); Acemoglu et al. (2002); Rodrik et al. (2004)). This debate, which has raged for decades, has eased somewhat, with the conclusion being that data and econometric weaknesses make the question extremely difficult to answer, but highlighting the importance of human capital and good policies in delivering growth (Durlauf et al. (2005); Glaeser et al. (2004); Garcia-Ponce & Wantchekon (2014)).

The third chapter of this thesis, titled "Surfing a wave of economic growth" aims to address some of these challenges. The chapter presented in this thesis is based on an earlier paper written jointly with Prof. Samuel Wills[12] (McGregor & Wills (2017)) which is currently under review at the Economic Journal. This chapter uses three clean natural experiments to identify how a particular natural amenity, surf breaks, contributes to economic growth. Surf breaks are well-suited to this type of study because their quality is exogenously determined by a finely balanced combination of weather and sea-floor geography (bathymetry), both locally and at distance. Two locations may be metres apart but of vastly different appeal to surfers, which we can exploit for identification. We measure quality directly through an independently verified rating system from a unique global database of 5,000+ locations, combined with data on El Niño events and global wave heights, and episodes of discoveries and disappearances of surf breaks.

Collectively these three experiments show that natural amenities raise economic growth. Experiment II shows that the increment to growth is larger when the amenity is initially discovered, and Experiment I shows that the effect on growth can persist for decades. Experiment III shows that the largest growth effects are apparent when the quality of the amenity itself is temporally boosted; in the case of surfing, through bigger waves. There are many possible channels through which amenities may affect growth, including: agglomeration effects as individuals and industries cluster near high quality amenities, augmented productivity of physical capital as recreational and tourism industries are established, increased labour productivity as migrants, workers and entrepreneurs are attracted to the

[12]Samuel Wills, Assistant Professor and Lecturer of Economics, University of Sydney. samuel.wills@sydney.edu.au.

local area.

1.6 A specific volatility channel

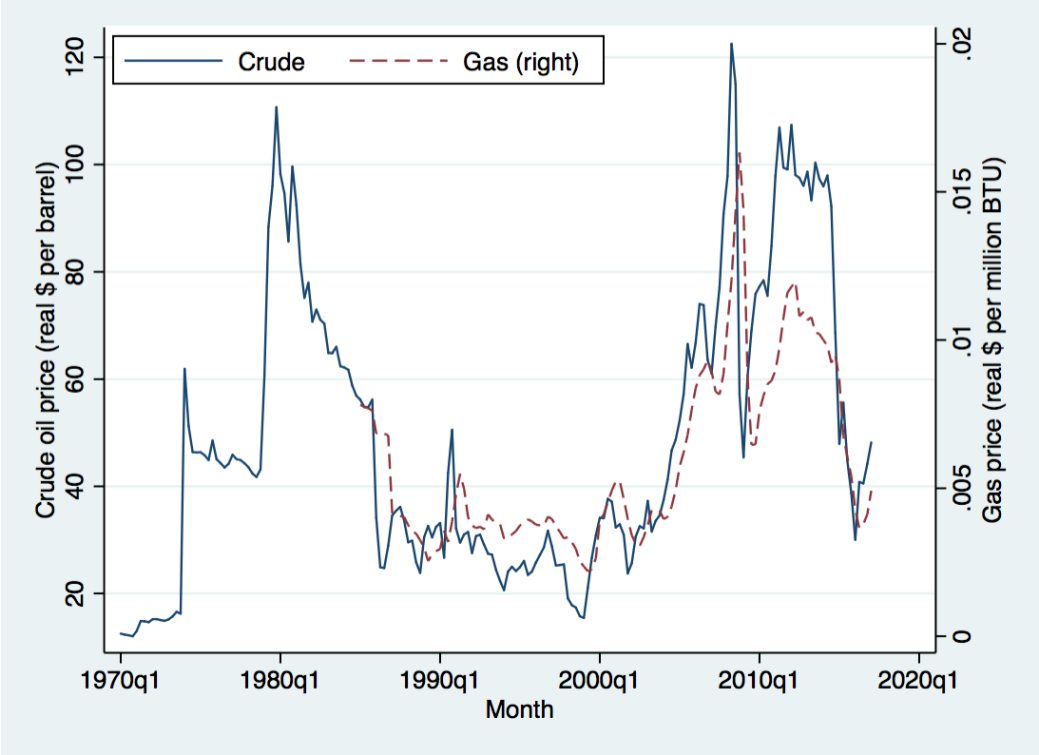
The negative relationship between macroeconomic volatility and economic growth is well documented (Ramey & Ramey (1995); Easterly et al. (1993)). The sources and types of macroeconomic volatility are numerous. One source which has been heavily debated in the literature is mismanaged macroeconomic policy. Some have argued that poor policy is to blame for the poor growth performance of many developing countries (Williamson (1990)) whilst others have highlighted the endogeneity of the quality of policy making and growth with respect to institutions (Acemoglu et al. (2003)).[13]

For resource dependent economies, an obvious source of macroeconomic uncertainty is commodity prices. The Latin American debt crisis is an example. Figure A.1 plots the monthly crude oil and natural gas price between 1970 and 2017. During the high commodity price era of the 1970s, many commodity exporters borrowed heavily on international capital markets, using resource revenues as collateral. These were often in the form of syndicated loans by groups of international banks destined for investment in large scale national development projects. Aggressive tightening of monetary policy by the Fed, along with falling commodity prices in the 1980s, caused many countries to face significant repayment difficulties and debt distress. A series of crises began with the sovereign default of Mexico in August of 1982 and spread quickly throughout the rest of Latin America.

A similar cycle appears to be playing out across resource rich emerging economies today. Steadily rising commodity prices since the turn of the millennium, often referred to as the commodities 'super-cycle', lead to rising debt levels in resource exporting economies. After the 2008/09 financial crisis, international investors turned to developing and emerging economies in search of higher yields. With commodity prices back on the rise between 2009

[13]The debate surrounding the importance of monetary policy shocks in causing recessions in developed countries is another example. See Christiano et al. (1999); Clarida et al. (2000); Orphanides (2002, 2003); Romer & Romer (2004); Coibion & Gorodnichenko (2011) for an introduction to what is a vast literature.

Figure 1.3: Global oil and gas price



Source: Real prices, IMF International Financial Statistics, April 20, 2017

and 2011, resource rich economies with their promising commodity fundamentals were an attractive prospect for investors. The recent collapse in commodity prices, commencing at the end of 2014, has sparked a new wave of debt distress across the resource-dependent developing world. Mozambique has been the first to fall victim with a de-facto default in March 2017.

The role of volatile commodity prices and debt dynamics in resource rich economies has been shown to be crucial in understanding the resource curse. Work by van der Ploeg & Poelhekke (2009, 2010) shows that once one accounts for the endogeneity of resource wealth, the negative effect of resource wealth on growth operates indirectly via volatility, debt cycles and credit constraints. In addition, a volatile real exchange rate caused by movements in commodity prices can lead to lower levels of investment, thus contracting the tradable sector (which is typically capital intensive) and lowering productivity growth

(Gylfason et al. (1999)).[14]

Whilst a large literature exists on 'sudden stops' and capital flow reversals in emerging economies[15] there is a surprising lack of work looking specifically at the role that commodity price movements have on debt dynamics themselves in resource rich economies[16]. A significant body of work has investigated the link between default risks and optimal borrowing behaviour in a dynamic macroeconomic framework, starting with the seminal work of Eaton & Gersovitz (1981) and Grossman & Van Huyck (1988). Almost completely separately, a body of research has also developed exploring the link between commodity prices and macroeconomic performance, in both oil-importing economies (Hamilton (2003); Barsky & Kilian (2004)) and oil-exporting economies (Deaton & Miller (1995); Raddatz (2007); McGregor (2017)).

Two important questions remain. How do commodity price movements affect the pricing of sovereign debt in resource rich economies? And, what does this mean for the optimal management of volatile oil revenues in developing countries? The fourth chapter of this thesis, titled "Pricing sovereign debt in resource rich economies" aims to shed light on these two questions. It is based on a recent OxCarre working paper (McGregor (forthcoming)).

I develop a stochastic general equilibrium model of a resource rich economy with sovereign debt and endogenous default risk. To the best of my knowledge, this approach is new to the literature. The model consists of a small open economy that receives a stochastic stream of resource revenues. The government sells bonds to foreign investors which it can renege on in the future, at some cost, whilst international investors form expectations on the likelihood of sovereign default. This delivers an endogenous risk premium which is inversely related

[14]We also know that real exchange rate volatility can significantly stunt long-run productivity growth Aghion et al. (2009).

[15]See work by Calvo (1998); Calvo & Mendoza (2000); Mendoza (2006); Mendoza & Yue (2012) to name but a few.

[16]There has been a recent literature that uses a NK modelling environment to investigate optimal monetary policy in resource rich economies. Ferrero & Seneca (2015) for example, shows that conventional monetary policy may put too much weight on price stability at the cost of output stability following oil price shocks.

to the price of oil. The model is able to explain a large proportion of the business cycle fluctuations in interest-rate spreads in resource dependent developing economies.

The chapter then asks how specific structural features of developing economies, compared to advanced economies, affect the relationship between commodity prices and the optimal price of sovereign debt. These features include: a higher dependence on natural resource revenues, which manifests itself in part in covarying output and commodity prices; impatient consumers and governments; a higher degree of risk-aversion; and, a lower ability to substitute consumption inter-temporally. Including them in the model significantly improves the ability of the model to explain the key macroeconomic co-movements in a resource rich, developing economy context. Model simulations reveal an interesting policy insight. An endogenous risk premium that is driven by falling oil prices, provides an additional rationale for a *volatility* fund in which liquidity buffers are accumulated to manage debt repayments. These buffers should be larger the stronger the link between oil prices and the domestic economy is, the more impatient policy makers are and the more willing they are to substitute current for future consumption.

Chapter 2

Commodity price shocks, growth and structural transformation in low-income countries

2.1 Introduction

The recent collapse of world commodity prices has led to some analysts calling this the end of the so called commodities 'super-cycle'. Diminishing growth prospects for emerging market economies, especially China, combined with abundant supply has continued to put downward pressure on the prices of most commodities, with the oil price reaching its lowest levels by mid-2016 since the early 2000s according to the 2016 World Economic Outlook from the International Monetary Fund (IMF).[1] The question of how resource booms and busts affect growth is as pertinent today as it has ever been, particularly for commodity exporting low-income countries.

The empirical evidence on the effect of natural resource wealth on economic growth and development is mixed. Whilst some studies find that resource booms are a blessing, others find that they are a curse, and others still find that the outcome depends on a number of

[1] see IMF (2016)

factors.[2] Two key challenges have faced this literature from the start. The first, is that measures of resource wealth are likely to be endogenous to a whole host of economic and political factors. The second is that the channels through which any positive or negative economic impacts of a resource boom may operate are numerous and multifaceted. These include an appreciation in the real exchange rate, de-industrialisation, increased volatility, political patronage and rent-seeing, and violent conflict, to name but a few.

This chapter aims to shed light on both these fronts. It develops a new proxy for resource wealth, based on subterranean reserves, that is exogenous to economic conditions in low-income countries. A panel-VAR is used to orthogonalise movements in the commodity prices, specifically minerals and hydrocarbons, against global and local economic conditions. The short-run macroeconomic response for low-income, resource rich economies is investigated, as well as the de-industrialisation effects and a possible 'Dutch disease' channel.

Original work by Sachs & Warner (1995) and Sachs & Warner (1999) finds a negative relationship between resource wealth and long-run economic growth, whilst more recent work by Mehlum et al. (2006b) finds that this is particularly the case in countries with poor institutions. This result continues to hold even after controlling for a whole host of potential omitted variables, as shown by van der Ploeg (2011). Given the long time periods used, these findings can be thought of as relating to the long-run relationship between resource booms and economic growth.

The key challenge facing this early work was the endogeneity of potential measures of resource wealth. The standard approach has been to use resource abundance (the share of resources or commodity exports in GNP) as the independent variable. However, this measure is likely to be endogenous for two reasons. First, by construction, dividing by GDP means that higher GDP will correlate with lower resources as a share of GDP thus setting up a negative correlation. And second, this scaling exercise implies that resource wealth as a share of GDP is no longer independent of the economic policies and institutions that drive

[2] See van der Ploeg (2011) for a survey of the literature.

both economic growth and the scale of resource production (Manzano & Rigobon (2001)).

Controlling for this endogeneity results in the negative effect on growth coming into question. Brunnschweiler & Bulte (2008), for example, instrument for resource dependence, using resource abundance and institutions, and find that natural resource booms lead to higher growth. van der Ploeg & Poelhekke (2010) use economically recoverable reserves as a proxy for resource wealth and find that there is no evidence for either a curse or a blessing, unless one allows for an indirect negative effect via volatility.

Another strand of the literature, which is similar in vein to this chapter, focuses on the short-run effects of resource shocks on growth. Early work by Deaton & Miller (1995) finds a positive effect of commodity price booms on income levels in the Pan-Africa region using a VAR framework. More recent work by Raddatz (2007) makes use of a panel-VAR to study the impact of different external shocks, including commodity price shocks, on output volatility in low-income countries. They find that positive commodity price shocks significantly raise income in the short run, but that these shocks in general can only explain a small fraction of the output variance of a typical low-income country. In a slightly different approach, Collier & Goderis (2012) use a panel error correction model (ECM) to show that commodity booms have unconditional positive short-term effects on output, but potentially negative long-run effects in countries with poor governance. These studies take similar approaches to addressing the endogeneity issues that have plagued the earlier growth literature. The approach is based on the earlier work by Deaton & Miller (1995) who weight world commodity price indices by country specific commodity export shares in 1975, which are held constant over time.

The first contribution of this chapter is to construct an exogenous measure of resource wealth for low-income countries. The approach involves constructing measures of country specific commodity reserves which are then used to weight movements in global commodity price indices. By using reserves instead of exports, I argue that this new measure of resource wealth in low-income countries is more exogenous that those used in the literature to date.

Another concern is the assumed exogeneity of commodity prices and developing country economic conditions. Work by Hamilton (2003) and Kilian (2008) finds that much of the increase in crude oil prices since the 1970s, and particularity since 2002, has been due to increases in aggregate demand for all industrial commodities, rather than exogenous supply shocks. To the extent that global demand shocks may simultaneously affect both commodity prices and low-income country economic output, via export demand, foreign investment and capital flows, the endogeneity of commodity prices will be an issue.

Within the VAR literature, additional consideration must be given to the contemporaneous exogeneity assumptions imposed upon the system. Raddatz (2007), for example, imposes a lower triangular structure on the two blocks of a panel-VAR to study the impact of external shocks on developing countries. The first block contains high income country GDP, commodity based terms-of-trade, the international interest rate, and global disasters, with the second block containing low income country per capita GDP and aid flows. The key identifying assumption is that the flow of causality runs from rich country GDP to commodity prices to world interest rates, but not in reverse. However, this ignores the potential for monetary policy in rich countries to affect aggregate demand, as well as commodity prices.

The second contribution of this chapter is to develop a framework that allows for the high degree of simultaneity between global demand, commodity prices and low-income country growth. To this end, the panel-VAR approach serves two purposes. First, it allows us to control for the interdependence between global economic conditions and low-income country growth. And second, by using a block recursive causal ordering, it allows for a greater degree of simultaneity within the global variables block and the local variables block. The results suggest that a one standard deviation increase in commodity prices (around 17.7% on average) raises per capita income levels in low-income countries by 0.26%, government spending by 12.4% and investment by 4.4%. The effects are larger for less developed countries, economies with fixed exchange rate regimes and those that are more dependent on

commodity exports.

The final contribution of this chapter is the investigation of the impact on the structure of the extracting economies following commodity price shocks. Specifically, I ask whether commodity price shocks lead to de-industrialisation in low-income, resource rich countries. In an important contribution to the early literature, Corden & Neary (1982) formalised the de-industrialisation process that may result from resource booms. They distinguish between two channels: the 'resource movement' effect, which describes the process by which labour may be pulled out of the traditional manufacturing sector following a boom, and the 'spending effect', which describes the increase in demand for non-tradables following a boom. Using data from the UN on sectoral value added, I find that resource booms lead to significant de-industrialisation. A one standard deviation increase in commodity prices results in the share of value-added in manufacturing contracting by 0.17-0.22 percentage points (pp), whilst transport and other non-tradable service sectors expand by around 0.10-0.13 pp and 0.26 pp receptively. This link appears to operate through the real exchange rate, with resource shocks leading to appreciations which in turn cause a restructuring of the domestic economy.

Whilst these effects may appear small, it is important to bear in mind that they represent the effect of exogenous commodity price shocks that are not due to changes in aggregate demand or global financial conditions. These estimates suggest that commodity price movements alone may be less important in explaining the volatility of low-income country growth than other explanations, such as demand and productivity shocks.

The focus on resource rich, low-income countries is motivated on two fronts. First, recent discoveries of potentially vast hydrocarbon reserves in sub-Saharan Africa makes the question of how commodity shocks affect growth and structural transformation in these countries highly pertinent.[3] Second, low-income countries are often thought of as being

[3]According to estimates by Wood Mackenzie, recoverable natural gas reserves in Mozambique are 85 trillian cubic feet (tcf) and 18 tcf in Tanzania, with another 100 tcf yet to be found; In Uganda, Tullow Oil estimate that recoverable oil reserves are 1.1 billion barrels (bb), with some 1.4 bb potentially yet to be

highly susceptible to external shocks, particularly commodity price movements. Work by Collier & Gunning (1999), Dehn (2000) and Collier & Dehn (2001) has documented an important effect of commodity price shocks on growth in developing countries. Understanding both the dynamic and structural impacts these shocks have on low-income countries is vital.

The remainder of this chapter is structured as follows. Section 2 presents the empirical methodology and the two step estimation procedure. Section 3 describes the construction of the reserve-weighted commodity price index used to identify exogenous commodity price shocks for each country. Section 4 describes the data. Sections 5 and 6 present the estimation results. Section 7 explores the Dutch disease mechanism. Section 8 uses these results to simulate the effect of the recent oil price collapse on the Nigerian economy, and Section 9 concludes.

2.2 Empirical methodology

The main challenge in estimating the dynamic response of low-income, resource rich economies to commodity price shocks is in identifying movements in commodity prices that are truly exogenous to economic conditions in these counties. Movements in global commodity prices will reflect both demand and supply shocks. Work by Hamilton (2003) and Kilian (2008) finds that much of the increase in crude oil prices since the 70s, and particularity since 2002, has been due to increases in aggregate demand for all industrial commodities, rather than exogenous supply shocks. The problem is that movements in global aggregate demand, as well as international credit market conditions, are likely to simultaneously affect the economic conditions of low-income countries as well global commodity prices. In order to identify exogenous supply driven commodity price movements, I use a VAR model which orthogonalises the commodity price shocks against global aggregate demand and financial conditions.

found; in Madagascar, Madagascar Oil estimate recoverable oil reserves to be 1.7 bb, with a further 10.8 bb of oil and 167 tcf of gas in undiscovered reserves according to an assessment by USGS; Ghana discovered 4 bb of oil in its Jubilee field in 2007 and production is currently being scaled up; Kenya has also recently discovered potential oil and gas reserves.

The empirical approach consists of two steps. Given the interest in studying a group of low-income countries, the first step is to estimate a panel-VAR. The VAR model should, in theory, isolate only supply driven changes in commodity prices (or demand driven price changes that do not immediately affect low-income economies) and therefore be exogenous from the perspective of low-income, resource rich economies. The second step is to take these shocks and regress them on the various sectoral components of value added and operating surplus in these countries. The aim is to study both the aggregate macroeconomic response to commodity price shocks, as well as possible structural transformation processes that might take place as a result.

Clearly, the choice to focus on resource rich, low-income countries is an important one. As discussed, it is motivated on two fronts: (i) there has been a flurry of recent discoveries of potentially vast hydrocarbon reserves in parts of the developing world, particularly sub-Saharan Africa, and (ii) low-income countries are often thought of as being highly susceptible to external shocks, particularly commodity price movements. Understanding both the dynamic and structural impact these shocks have on low-income countries is vital. However, by restricting the sample to low-income countries raises issues around selection bias. Namely, by focusing on poor countries today the empirical results may be biased towards finding a negative relationship between commodity price shocks and economic growth. For this reason, rather than using current income classifications, I select the sample of low-income countries using per capita GNI data from the World Bank in 1970.

2.2.1 Structural panel-VAR

The first step involves estimating a structural VAR model on a panel dataset of macroeconomic variables of resource rich developing economies.[4] Using a VAR approach allows all of the variables in the system to be endogenous. This is important, given the simultaneity in global economic conditions, commodity prices and poor countries' economic conditions. Applying it to panel data allows for unobserved heterogeneity.

[4]I am grateful to Inessa Love and Lea Ziccino for the use of their panel-VAR programme in Stata. See Love & Zicchino (2006).

Given that I am interested in estimating the dynamic macroeconomic response of natural resource shocks, the key challenge is in identifying exogenous structural shock series for each country. I do this in two ways. The first involves using a recursive causal ordering on the system, or a Cholesky decomposition. Whilst this is a powerful method for identifying structural shocks in VAR models, it does come at the cost of imposing some prohibitively restrictive assumptions on the dynamics of the system. The second method uses a set of 'block recursive' type restrictions in which global macroeconomic variables are treated as a separate and independent block to the national variables, such as GDP, investment and government spending. Whilst I continue to impose contemporaneous exogeneity between the two blocks, I do not make any assumption about the causal ordering of variables within each block. This approach imposes much fewer restrictions on the system and is, for that reason, theoretically more appealing a priori.

I begin by describing the problem of identification of shocks in the panel-VAR model. Consider, the following first order structural VAR (SVAR) model:

$$\boldsymbol{B}\boldsymbol{x}_{i,t} = \boldsymbol{\Gamma}_0 + \boldsymbol{\Gamma}_1\boldsymbol{x}_{i,t-1} + \boldsymbol{\epsilon}_{i,t} \tag{2.1}$$

where $\boldsymbol{x} = [GDPH, LIBOR, CPIND, G, I, INFL, RGDPPC]'$ is a vector of endogenous variables for country i in year t. The variables included in the vector are: log changes in real GDP of high income countries ($GDPH$), real US Libor rate ($LIBOR$), log changes in a GDP weighted commodity price index ($CPIND$) which will be discussed in more detail below, log changes in real government expenditure (G), log changes in real fixed capital formation, (or investment) (I),[5] inflation as measured by the GDP deflator index ($INFL$), and log changes in real GDP per capita ($RGDPPC$). The log changes specification for most of the variables is required to ensure stationarity in $\boldsymbol{x}_{i,t}$.

An identification problem arises when attempting to estimate Eq. 2.1. What is actually

[5] G is General government final consumption expenditure (current US$); I is Gross fixed capital formation (constant 2005 US$). Both are taken from the World Bank's World Development Indicators.

estimated when taking the model to the data is the following reduced form VAR (RVAR):

$$\boldsymbol{x}_{i,t} = \boldsymbol{A}_0 + \boldsymbol{A}_1 \boldsymbol{x}_{i,t-1} + \boldsymbol{e}_{i,t} \tag{2.2}$$

where $\boldsymbol{A}_0 = \boldsymbol{B}^{-1}\boldsymbol{\Gamma}_0$, $\boldsymbol{A}_1 = \boldsymbol{B}^{-1}\boldsymbol{\Gamma}_1$, $\boldsymbol{e}_{i,t} = \boldsymbol{B}^{-1}\boldsymbol{\epsilon}_{i,t}$, $\mathbb{E}(\boldsymbol{e}\boldsymbol{e}) = \boldsymbol{\Sigma} = \boldsymbol{B}^{-1}\Omega\boldsymbol{B}^{-1'}$, and $\mathbb{E}(\boldsymbol{\epsilon}\boldsymbol{\epsilon}') = \boldsymbol{B}\boldsymbol{\Sigma}\boldsymbol{B}$. The immediate challenge in estimating Eq. 2.2 using panel data is that we are required to impose the restriction that the structure of the VAR, in particular the dynamics of the system, represented by the $\boldsymbol{A}$ matrices, are common across countries. This is a standard assumption in the literature (see Broda, 2001, 2004; Ahmed, 2003; Uribe & Yue, 2006) because, given the length of the time series dimension of the data (around 29 years), it would not be possible to estimate country-specific dynamics.

Including fixed effects, f_i, in the model would allow us to relax this assumption for the $\boldsymbol{A}_0$. The estimation equation becomes:

$$\boldsymbol{x}_{i,t} = \boldsymbol{A}_0 + \boldsymbol{A}_1 \boldsymbol{x}_{1,t-1} + f_i + \boldsymbol{e}_{i,t} \tag{2.3}$$

Since the fixed effects are correlated with the regressors due to lags of the dependent variables, the standard fixed effects estimator will be biased. I therefore use forward mean-differencing, also referred to as the 'Helmert procedure' (see Arellano and Bover, 1995). This procedure removes only the forward mean, i.e. the mean of all the future observations available for each country-year. This transformation preserves the orthogonality between transformed variables and lagged regressors, so the lagged regressors can be used as instruments and estimate the coefficients by system GMM[6].

Whilst the estimates of the RVAR are consistent, it is not possible to recover the coefficients of the SVAR, which is the object of interest. The RVAR forces the contemporaneous structural shocks of each variable in $\boldsymbol{x}_{i,t}$ into the errors $\boldsymbol{e}_{i,t}$ which consists of some combi-

[6]Various estimators based on GMM have been proposed to calculate consistent estimates of model equations, especially in fixed T and large N settings. Instead of using the standard first difference (FD) approach, in which each equation is instrumented using lagged levels as instruments for lagged differences in, we use forward orthogonal deviations (FOD). The estimation is carried out using the the interactive version of Stata's *gmm* command with analytic derivatives

nation of the true structural shocks related to each variable, $\boldsymbol{\epsilon}_{i,t}$. To isolate the structural shocks to any one of the variables in the system it is necessary to decompose the residuals in such a way that they become orthogonal.

In the following sections I describe two different approached to decomposing the reduced form residuals in order to retrieve the structural shocks. In each case the process involves: (i) estimating the panel-VAR by system GMM, (ii) making a set of identifying assumptions, and (iii) recovering the structural shocks from the estimated coefficients by imposing the implication of these identifying assumptions on the $\boldsymbol{B}^{-1}$ matrix.

2.2.1.1 Cholesky

The first approach uses a simple recursive causal ordering, or a Cholesky decomposition, of the $\boldsymbol{B}^{-1}$ matrix, which allows for the identification of the structural shocks $\boldsymbol{\epsilon}_{i,t}$ in the panel-VAR. Identification via a recursive causal ordering implies that the $\boldsymbol{B}^{-1}$ matrix linking the structural shocks and the reduced form errors is lower triangular.

$$\boldsymbol{e}_{i,t} = \boldsymbol{B}^{-1} \times \boldsymbol{\epsilon}_{i,t}$$

$$\begin{bmatrix} e_{GDPH_t} \\ e_{LIBOR_t} \\ e_{CPIND_{i,t}} \\ e_{G_{i,t}} \\ e_{I_t} \\ e_{INFL_{i,t}} \\ e_{RGDPPC_{i,t}} \end{bmatrix} = \begin{bmatrix} 1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ b_{71} & \cdots & 1 \end{bmatrix} \times \begin{bmatrix} \epsilon_{GDPH_t} \\ \epsilon_{LIBOR_t} \\ \epsilon_{CPIND_{i,t}} \\ \epsilon_{G_{i,t}} \\ \epsilon_{I_t} \\ \epsilon_{INFL_{i,t}} \\ e_{RGDPPC_{i,t}} \end{bmatrix} \tag{2.4}$$

where $\boldsymbol{\epsilon}_{i,t}$ is the vector of the structural shocks of interest from Eq. 2.1, $\boldsymbol{b}_{ij}$ are the parameters of interest from the $\boldsymbol{B}^{-1}$ matrix, and $\boldsymbol{e}_{i,t}$ are the reduced form errors from estimating the panel-VAR in Eq. 2.2. In using this approach, the ordering of the variables in $\boldsymbol{x}_{i,t}$ becomes important. The identifying assumption used in this case is that each variable

in $\boldsymbol{x}_{i,t}$ affects the variables that appear below it contemporaneously but only affects the variables above it with a lag. More specifically, it assumes that the contemporaneous causal order runs from the GDP of high income countries, to the US Libor rate, to the commodity price index for each country and so on. The most endogenous variable is real per capita GDP in each country which is affected by all of the variables appearing before it in the VAR. This includes government spending and aggregate investment. This causal ordering means, for example, that the residual from the first equation of the RVAR, $\boldsymbol{e}_{GDPH,t}$, contains only the orthogonalised shock of high income country GDP, $\boldsymbol{\epsilon}_{GDPH,t}$.

The ordering of the commodity price index ($CPIND$) after the global economic indicators ($GDPH$ and $LIBOR$) means that commodity prices are assumed to only affect rich country GDP and the US Libor rate with a lag. Given, that this assumption is unlikely to hold in practice I explore an alternative identification strategy in the following section.

Placing $CPIND$ before individual countries' economic indicators (G, I, $INFL$ and $RGDPPC$) implies that commodity prices are not affected by the present economic performance of the low-income country (government spending, investment, inflation and per capita GDP). This is the small country assumption. I appeal to two arguments in support of this assumption. First, none of the countries included in the sample are large enough exporters of any individual commodity to have an effect on global prices.[7] And second, the construction of the commodity price index $CPIND$ is such that any risk of contemporaneous correlation between it and national economic conditions is minimised. One concern is that the standard measures of resource abundance/wealth in the literature are endogenous to economic conditions, which raises the question as to whether a more suitable measure exists[8]. The approach taken here is to use historical commodity deposits rather than current commodity exports to weight each price index. The construction of the commodity

[7]The list of the 28 countries included in the panel is discussed in the data section and I believe that this assumption is likely to hold in the sample. No countries have been removed.

[8]That is, current per capita GDP of the low income country may be affected directly by the level of exports (and commodity exports in particular). If commodity exports are used in the denominator of the weighted price index, as is often the case in studies of natural resource shocks, this would likely violate the contemporaneous homogeneity assumption of the commodity price index $CPIND$ with respect to $RGDPPC$

price index is discussed in detail in Section 3.

2.2.1.2 Block recursive

The second approach involves splitting the VAR into two blocks and applying the recursive causal ordering to the blocks, rather than the specific variables. This is known in the literature as 'block recursive' identification. The first block contains the global variables ($GDPH$, $LIBOR$ and $CPIND$) whilst the second block contains the national variables (G, I, $INFL$ and $RGDPPC$). The VAR model becomes:

$$\boldsymbol{e}_{i,t} = \boldsymbol{D}^{-1} \times \boldsymbol{\epsilon}_{i,t}$$

$$\begin{bmatrix} \boldsymbol{e}_{I_t} \\ \boldsymbol{e}_{N_{i,t}} \end{bmatrix} = \begin{bmatrix} d_{11} & 0 \\ d_{21} & d_{22} \end{bmatrix} \times \begin{bmatrix} \boldsymbol{\epsilon}_{I_t} \\ \boldsymbol{\epsilon}_{N_{i,t}} \end{bmatrix} \tag{2.5}$$

where $\boldsymbol{e}_I = [\boldsymbol{e}_{GDPH_t}, \boldsymbol{e}_{LIBOR_t}, \boldsymbol{e}_{CPIND_{i,r}}]'$ is the vector of reduced form errors of the international variable block, and $\boldsymbol{e}_N = [\boldsymbol{e}_{G_{i,t}}, \boldsymbol{e}_{I_{i,t}}, \boldsymbol{e}_{INFL_{i,t}}, \boldsymbol{e}_{RGDPPC_{i,t}}]'$ is the vector of reduced form errors from the national block. The element d_{11} is a $[3 \times 3]$ matrix of parameters to be estimated, whilst $[d_{21}, d_{22}]$ is a $[4 \times 7]$ matrix of parameters. Identification via a block recursive strategy implies that the $\boldsymbol{D}^{-1}$ matrix is lower triangular but the $\boldsymbol{B}^{-1}$ no longer is.

The identifying assumptions made here are less strict than those in the Cholesky decomposition case in two ways. First, all of the national level variables are assumed to depend on each other contemporaneously. That is, I make no prior assumptions about which national level variables are pre-determined with respect to each other. And second, I allow the global variables, rich country GDP, the US Libor rate and the commodity price indices, to be contemporaneously inter-dependent with each other. The twelve zero restrictions in the top right corner of the B matrix are the sole restrictions placed on the system, namely that the three global level variables do not respond to the national level macroeconomic conditions contemporaneously. I justify the first eight restrictions ($b_{14} : b_{17} = 0$ and $b_{24} : b_{27} = 0$) using

the small country assumption, i.e. each individual country in the sample is not big enough to influence either global demand or global financial conditions. Given that I restrict the sample to low income countries, these restrictions seem plausible. The final four restrictions ($b_{34} : b_{37} = 0$) imply that the commodity price indices for each country in the sample do not react to national level economic conditions (specifically, national level GDP, investment, inflation and government spending). The same problems arise as in the Cholesky identification strategy. Again, the construction of the price index will be important here.

It is worth noting that while the Cholesky decomposition produces a unique solution for the $\boldsymbol{B}^{-1}$ matrix, the block recursive identification strategy does not. There are potentially an infinite number of rotations of the $\boldsymbol{B}$ matrix that would satisfy the zero restrictions placed on the block recursive model. The key, however, is that it is not possible to identify all of the structural shocks ($\boldsymbol{\epsilon}_{xt}$) under block recursion. In the current set-up however, it will be possible to identify the impact of individual shocks from one block on the variables of another.[9] Specifically, the interest will be in identifying the impact of commodity price shocks ($CPIND$), which is located in the first block, on the national level variables (G, I, $INFL$, $RGDPPC$) in the second block.

Finally, the model in Eq. 2.3 is estimated using a single lag. The use of a single lag structure is done for ease of interpretation of the results, however, higher order lag structures, up to four lags, are also estimated, as are a number of lag-order selection statistics for the model. I find that one lag is a reasonable choice for this model.

2.2.2 Seemingly unrelated panel regressions

Having identified the structural commodity price shocks in the panel-VAR, the second step involves taking the orthogonalised commodity price error ($CPIND_{i,t}$) and regressing these on the share of value-added in different sectors of the domestic economy. The aim is to in-

[9]Satisfying the rank condition is key here. In an n-variable VAR, an equation is identified if and only if at least one non-zero determinant of order $(n-1)$ by $(n-1)$ can be constructed from the coefficients on the variables not excluded from that particular equation, when they appear in other equations in the model. I check the rank condition and find that it is satisfied.

vestigate the impact of commodity price shocks on the sectoral performance of low-income, resource rich economies. The hypothesis is that, if a Dutch disease mechanism is at work, in which resource booms lead to a loss of international competitiveness, then these shocks, as proxied by shocks to the country specific commodity price index, should have a differential impact on the share of value-added in sectors that are more or less exportable. This follows from work by Rajan & Subramanian (2011) who find that aid inflows adversely affect manufacturing sectors that are more exportable and so susceptible to movements in the exchange rate.

I use a seemingly unrelated regression (SUR) framework which assumes that the error terms are correlated across equations. This is important as the residuals from the sectoral regressions are unlikely to be independent. I estimate the following model:

$$y_{i,j,t} = \alpha + \beta_1 G_{i,t} + \beta_2 I_{i,t} + \beta_3 INFL_{i,t} + \beta_4 RGDPPC_{i,t} + \beta_5 y^{ini}_{i,j,t} + \beta_6 CSHOCK_{i,t} + f_i + \eta_t + \epsilon_{i,j,t} \tag{2.6}$$

where G, I, $INFL$ and $RGDPPC$ are defined as before, $y_{i,j,t}$ is the share of real value-added in sector j of country i at time t in total real value-added, $y^{ini}_{i,j,t}$ is the initial share of real value-added, f_i are country fixed effects, η_t are year fixed effects, and $CSHOCK_{i,t}$ is the orthogonalised error ($\boldsymbol{\epsilon}_{CPIND_{i,t}}$) from the panel-VAR model for country i at time t, scaled by the standard deviation of the shocks[10]. Here, the structural error can be thought of directly as the exogenous component of commodity price shocks. The sign and size of the coefficient β_6 for each sector will highlight the existence or otherwise of potential de-industrialisation effects from positive commodity price shocks.

I exclude the global variables ($GDPH$ and $LIBOR$) from the regression as these will be endogenous. Whilst this will also be the case for the local level variables (G, I, $INFL$ and $RGDPPC$), I include these simply as controls as they will explain a large proportion of the

[10] As shown by Pagan (1984), a generate regressor can render the standard errors of the OLS estimates inconsistent. We correct for this using a two-step bootstrapping method to compute the standard errors.

variation in sectoral value added between countries. This should not be a concern, given that I am not making a causal interpretation of these coefficients, however I do check the robustness of the results to omitting all endogenous variables from the model (see appendix B for details).

2.2.3 Weighted price index

The construction of the reserve-weighted commodity price index ($CPIND$) deserves further explanation. The aim is to construct a country specific commodity price index which measures the exposure of each low income country to movements in various global commodity prices, but which does not violate the contemporaneous exogeneity assumption between the commodity price index and national level variables ($b_{34} : b_{37} = 0$).

I construct an annual commodity price index ($PIND_{c,t}$) using UNCTAD commodity price data for nine minerals[11] and EIA data for oil and gas prices, where 2010 = 100. The analysis is restricted to minerals and hydrocarbons for a number of reasons.

First, much of the theory of the resource curse tends to focus on point resources. By focusing on minerals and hydrocarbons, this chapter remains strictly within the realm of this existing theoretical work. In addition, the commodities I choose to focus on are predominantly windfall resources; that is, their discovery, excavation and exploitation is less endogenous to natural, political or economic conditions in the country. This is less likely to hold for a number of commodities, particularly agricultural products.

Second, the resource curse theory emphasises that the commodity sector is price-taking so that the effects of domestic demand do not feature in the analysis. This is important as it allows us to focus on the relative price challenge, i.e. the real exchange rate. Again, this would not be the case for many agricultural commodities which are subject to local and regional demand dynamics.

[11] These are: gold, silver, tungsten, iron, copper, lead, nickel, tin and zinc. Some organic sediments such as coal are not included in the UNCTAD monthly price series.

And finally, a key feature of the chapter is the construction of an exogenous commodity index measure for each country. To do this, I make use of subterranean reserve data. Clearly the concept of reserves does not apply to many commodities, including agricultural products.

Changes in commodity prices may affect low income commodity exporters differently depending on the importance of these commodities in the economy. The problem with simply weighting the commodity price index by the share of that commodity in exports or GDP is that this scaling exercise implies that this variable is no longer independent of economic policies and institutions, and is potentially endogenous to domestic economic conditions. Instead, I construct a weight for each commodity, in each country, which is equal to the value of commodity reserves as a proportion of GDP in 1985.[12] The key is that this 1985 reserve-GDP weight is independent of economic policy and institutions from 1985 onwards and is thus exogenous to local economic conditions. This weight is held constant over time and constructed as follows:

$$w_{c,i}^{1985} = \frac{P_c^{1985} r_{c,i}^{1985}}{Y_i^{1985}} \tag{2.7}$$

where P_c^{1985} is the price of commodity c in 1985, $r_{c,i}^{1985}$ is the volume of reserves of commodity c in country i in 1985, and Y_i^{1985} is nominal GDP in country i in 1985. One problem with using 1985 reserves directly, is that the discovery of subterranean natural resources is likely to be endogenous to local economic and political conditions. New discoveries post-1985 are more likely to occur in countries that have favourable economic conditions. The risk is that, by ignoring these newly discovered reserves, I over-estimate the impact of commodity price shocks, which are weighted by reserves, on growth.

To deal with this, I construct an estimate of commodity reserves in each country in 1985

[12]The data I use comes from the IntierraRMG database on mining intelligence. The database includes annual production, at the mine level, from 1985 to 2013 as well as reserve data at the mine level for 2013. This means that the earliest I can construct reserve data for is 1985.

by summing up the annual production over the period and adding this to the 2013 estimate of sub-soil commodity reserves from the IntierraRMG database. The definition of $r_{c,i}^{1985}$ is therefore given by:

$$r_{c,i}^{1985} = r_{c,i}^{2013} + \sum_{t=1985}^{2013} y_{c,i}^{t} \tag{2.8}$$

where $r_{c,i}^{2013}$ is the volume of reserves of commodity c in country i in 2013, and $y_{c,i}^{t}$ is the production of commodity c in country i at time t. I then apply these weights to the annual price index $pind_{c,t}$ for each commodity c in each year t. This forms the 1985 reserve-GDP weighted commodity price index for each commodity produced in each country, which is summed over all commodities to produce a reserve-GDP country commodity price index ($CPIND_{i,t}$) as follows:

$$CPIND_{i,t} = \sum_{c=1}^{13} w_{c,i}^{1985} pind_{c,t} \tag{2.9}$$

This allows the effect of commodity price movements to be larger for countries in which commodities represent a larger share of GDP, as measured by the size of their commodity reserves in 1985.

2.3 Data

2.3.1 Sources

The analysis in this chapter makes use of data from a number of different sources: United Nations (UN) Statistics, United Nations Industrial Development Organisation's (UNIDO) INDSTAT2 database, UN Conference on Trade and Development (UNCTAD), International Monetary Fund (IMF) World Economic Outlook (WEO), IMF International Financial Statistics (IFS), World Bank (WB) World Development Indicators (WDI), The Penn World Tables 6.2 (PWT), The Natural Resource Governance Institute, US Energy Information Agency (EIA) and Intierrar-RMG.

I categorise a country as being resource rich using the IMF's definition. According to the

IMF, a country is defined as resource rich if it has either natural resource revenue or exports at least 20% of total fiscal revenue or total exports, on average, over the period 2006-2010.[13] There are 51 countries classified as resource rich. In addition, I add 12 countries who have recently discovered substantial natural resource reserves but where production may not yet have begun.[14] I also include 9 countries where, according to the Natural Resources Governance Institute, minerals hold great potential for future fiscal revenue.[15] This gives a total of 72 resource rich countries.

I classify countries as low, middle or high-income using data on real GNI per capita using the World Bank Atlas method. Low-income countries are those whose real GNI per capita in 1970 was below the equivalent of 1045 USD in the 2015 fiscal year. Of the 72 resource rich countries, 46 are classified as low-income according to this definition.

Data on the sectoral breakdown of value added, employment and gross fixed capital formation comes from the United Nations (UN) database on National Accounts Estimates of Main Aggregates and UNIDO. The National Accounts data records the value added by industry according to International Standard Industrial Classification of All Economic Activities (ISIC), revision 4, 1-digit codes. The data is only available for the main groups and some categories are grouped together. The data covers 7 sectors[16] in 211 countries from 1970 to 2012.

Rich country GDP, which covers advanced economies[17] only, as well as data on the US Libor rate and the US deflator, comes from the IMF's WEO database from 1980 to 2013.

[13] Gabon and Equatorial Guinea are included because of CEMAC monetary union membership; while Liberia, Niger, Cote d'Ivoire and Uzbekistan are included in spite of incomplete data availability.

[14] These are: Sierra Leone, Afghanistan, Madagascar, Mozambique, CAR, Uganda, Tanzania, Togo, Kyrgyz Re-public, Sao Tome and Principe, Ghana and Guatemala.

[15] These are: Cambodia, Colombia, Egypt, India, Myanmar, Morocco, Philippines, South Sudan and Zimbabwe.

[16] The sectors are: agriculture, hunting, forestry, fishing (ISIC A-B); mining, manufacturing, utilities (ISIC C- E); manufacturing (ISIC D); construction (ISIC F); wholesale, retail trade, restaurants and hotels (ISIC G-H); transport, storage and communication (ISIC I); other activities (ISIC J-P).

[17] Composed of 35 countries: Australia, Austria, Belgium, Canada, Cyprus, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hong Kong SAR, Iceland, Ireland, Israel, Italy, Japan, Korea, Luxembourg, Malta, Netherlands, New Zealand, Norway, Portugal, San Marino, Singapore, Slovak Republic, Slovenia, Spain, Sweden, Switzerland, Taiwan Province of China, United Kingdom, and United States.

Data on the real exchange rate, inflation, population, T-bill rate, real per capita GDP in PPP terms for each low income country also comes from the WEO, also from 1980 to 2013.

Purchasing power parity (PPP) conversion factors and nominal exchange rates for the calculation of the real exchange rate come from the Penn World Tables version 6.2.

Annual production of natural gas and oil comes from the US Energy Information Administration (EIA) and covers 208 countries from 1980 to 2013. The EIA also provides a time series of gas prices from 1922, and crude oil prices from 1986.

Data on individual mineral production and reserves comes from the IntierraRMG database on mining intelligence. The database includes annual production, at the mine level, from 1985 to 2013 as well as reserve data at the mine level for 2013. This data is aggregated up to the national level, and reserves in 1985 are calculated by adding the cumulative annual production to the reserves remaining in the ground in 2013, as discussed above.

I use commodity price data from UNCTAD from 1960 to 2013.[18] The minerals for which I have production, reserves and prices for are: gold, silver, platinum, palladium, aluminium, iron, copper, lead, nickel, tin, zinc, diamonds and tungsten.[19]

2.3.2 Summary statistics

I use the full sample of low-income, resource rich countries in the sample, for which there is sufficient data. Of the original 46 countries, 16 are dropped from the analysis due to lack of data. Table B.1 provides a summary of the data. Of the 30 remaining countries, half are in Sub-Saharan Africa, 8 are in Latin America or the Caribbean, 4 are in North Africa or the Middle East and 3 are in Asia (East or South). The majority, 12, are LMICs today (according to their per capital GNI in 2013), with 8 UMICs and 2 HICs. 8 of the countries

[18] Price data for palladium and platinum are not available from UNCTAD so I use data from the RMG database instead.

[19] The diamond price is calculated by dividing production quantity by production value - thus specific to each country.

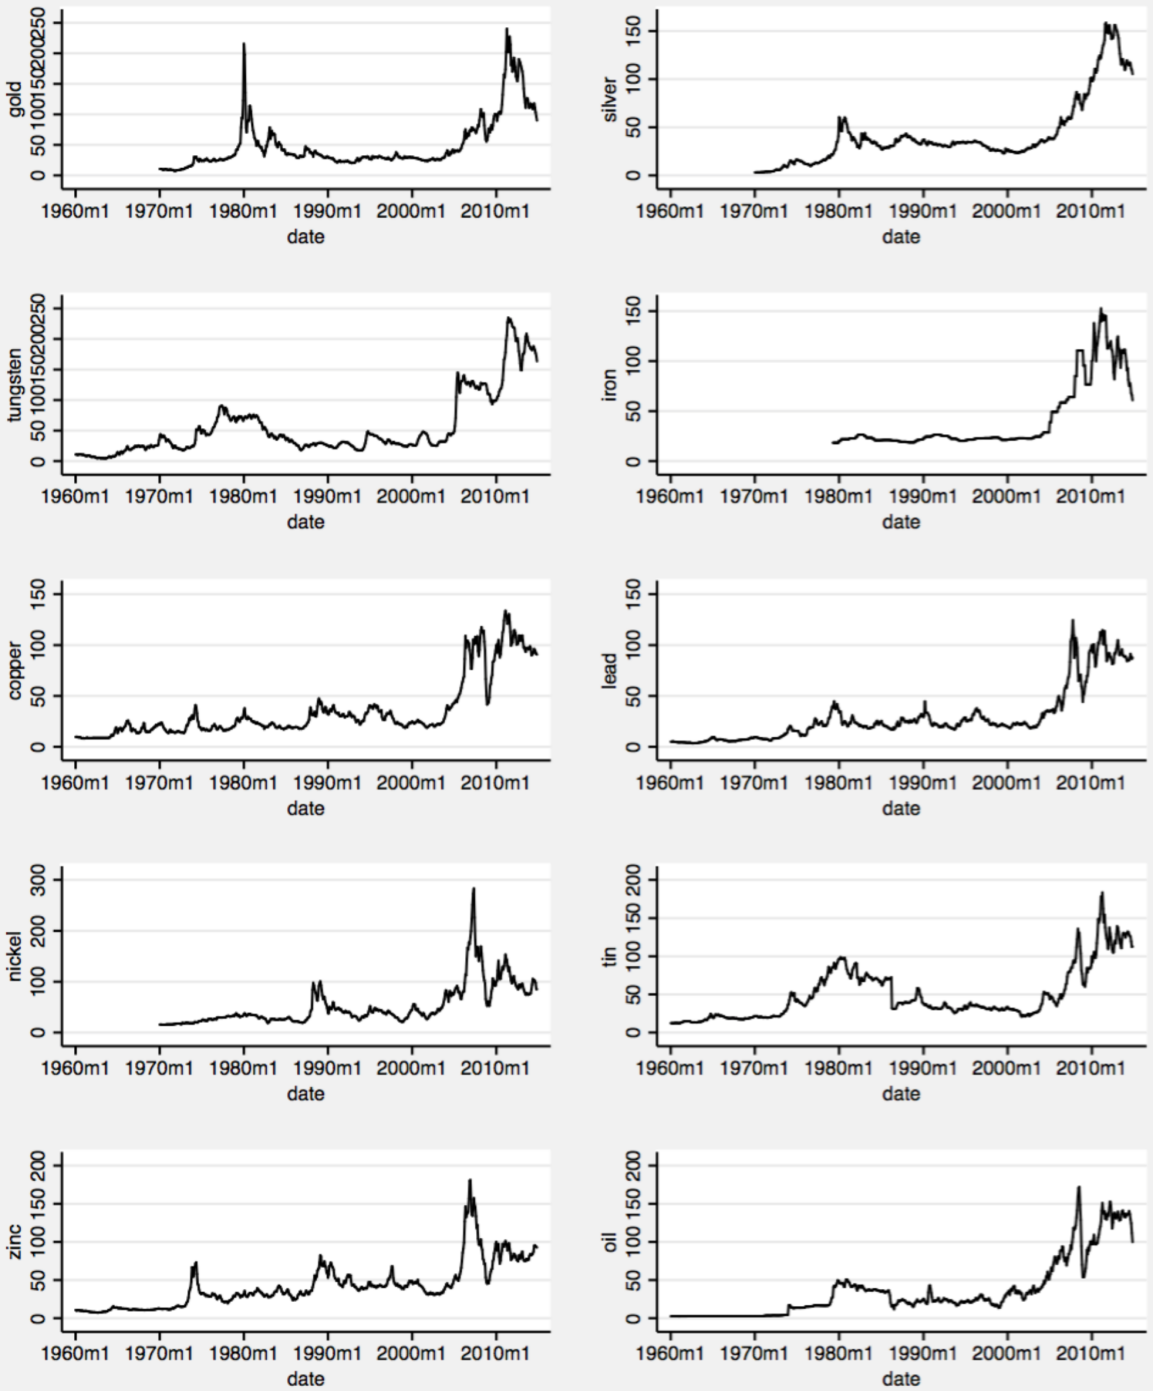
remain LICs today.

The first thing to note is that the majority of countries in the sample are no longer low-income countries. That is to say, many of the resource rich countries who were poor in 1970 have in fact managed to foster sufficient economic growth to move themselves out of low-income country status. Clearly, the countries that have grown faster over the period are now richer, as is reflected in the average growth rates by income group. However, those countries that have remained poor have not only grown at a slower pace over the period, but this growth has been much more volatile than any other group. Second, the countries that have remained poor have also experienced the largest declines in government expenditure on average. Those that have managed to achieve middle-income country status have seen government expenditure fall by far less over the period. Only the high-income country group have had positive growth in government expenditure. Finally, it would appear that the low growth rates observed in low-income countries do not appear to have been due to low levels of investment. Investment, as measured by the growth in real fixed capital formation, has actually been relatively high in the low-income country group at around 7.1% per annum on average. However, this investment has been extremely volatile over the period.

The monthly evolution of commodity prices over the period 1960-2013 for the nine mineral commodities, plus oil,[20] used in this chapter is given in figure 2.1. The first thing to note is the significant amount of volatility in each of the series over the time period. Second, there appear to be three periods in which the commodity price series display common shocks. These are around 1980, 1990 and 2008, all of which correspond to major global economic crises. It is important that, when using commodity prices as a proxy for commodity shocks, these potentially demand driven movements in commodity prices are dealt with appropriately.

[20] I do not have monthly prices for natural gas, only annual prices.

Figure 2.1: World commodity prices



Monthly commodity price index calculated using commodity price data from UNCTAD for the following commodities: gold, silver, tungsten, copper, lead, iron, nickel, tin, zinc and oil. Monthly price index for gas not available.

2.4 Results

This section presents the results from estimating the panel-VAR model in Eq. 2.3. The model is estimated using annual data for the 30 low-income, resource rich countries in the dataset, covering the period 1985-2013. The average number of years for which data exists for each country is 24, giving 704 observations in total. The model in Eq. 2.1 could be estimated using OLS, however, as discussed in Section 2, given the inclusion of fixed effects and a lagged dependent variable, common dynamic panel bias becomes a problem (Nickell (1981)).[21] To overcome this problem, I estimate the coefficients by system GMM using forward-mean differencing. Given the use of some observations (and their lags) as instruments in the GMM estimation the sample size is reduced slightly to 643.

The following two figures present the IRFs for various shocks identified using the Cholesky and block recursive strategies. Only selected IRFs are presented here and confidence intervals are left out for ease of viewing. Note however, that all of these IRFs are statistically significant at the 5% level or better.[22] As most of the variables in the model are in log differences, the IRFs show the log changes in each variable. The IRFs are scaled by the standard deviation of the shocks, and so they represent responses of variables to shocks of the order of magnitude seen in the data. Each panel plots the dynamic response of one variable to a one standard deviation shock in another.

In interpreting these results, it is worth bearing in mind that, if the identifying restrictions hold, these IRFs plot the dynamic response of each variable in the system to an exogenous shock, or impulse, of each variable.

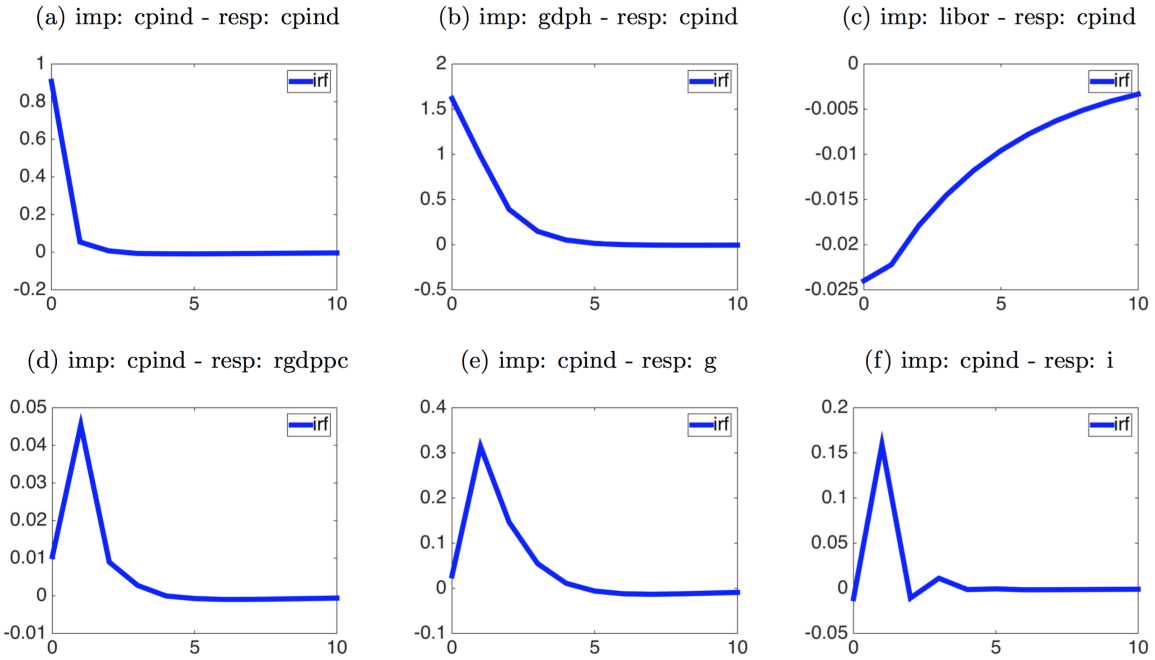
[21] While it is true that the bias in dynamic panels approach zero as T becomes large, Judson & Owen (1999) find significant bias even when T = 30. In my dataset, I have 24 years of data on average for each country.

[22] The IRFs from the Cholesky model with 95% confidence bands, before scaling, are presented in appendix B (figure B.1). The confidence bands are generated by Monte Carlo simulation with 10,000 simulations. The process involves randomly generating a draw of coefficients of A0 and A1 from Eq. 2.2 using the estimated coefficients and their variance-covariance matrix and re-calculate the impulse- responses. This process is repeated 100 times and the 5th and 95th percentiles of the distribution of coefficients are used to generate the confidence interval for the IRFs.

2.4.1 Cholesky

Figure 2.2 presents the orthogonalised and cumulative impulse response functions (IRFs) from applying the Cholesky decomposition. Panel (a) shows the own IRF for the country specific reserve-GDP weighted commodity price index ($CPIND$). It plots the time profile of a one-period, one standard deviation increase in the commodity price index (just over 17.7% on average). We see that the bulk of the effect of shocks on the price of commodities produced by low-income, resource rich, countries tend to last no more than one year on average.

Figure 2.2: Panel-VAR IRFs - Cholesky identification strategy



Panels (b) and (c) show the dynamic response of the commodity price index ($CPIND$) to a one standard deviation impulse to advanced economy real GDP ($GDPH$) and the US Libor rate ($LIBOR$) respectively. These two panels correspond to the dynamic response of commodity prices to shocks in aggregate global demand and global financial conditions. As expected, a one standard deviation increase in global aggregate demand (or around 4.2%), leads to an immediate and statistically significant increase of the average commodity price index for low-income countries of around 1.5 standard deviations. The dynamic response of commodity prices to global demand shocks lasts for around 3 years. A one standard

deviation shock to the 3-month US Libor rate (or an increase of around 1.25 percentage points), which proxies for the actual cost of borrowing for LICs, leads to an immediate fall in commodity prices of 0.024 standard deviations. This is likely due to the fact that in normal times, the US Libor rate will move with the Federal Funds rate. Any positive shock to this rate will result in lower demand holding supply constant, and so lower commodity prices.

Panel (d) shows the response of real per capita GDP ($RGDPPC$) to a shock in the commodity price index ($CPIND$). A one standard deviation increase in the commodity price index results in an increase in real per capita GDP in the following period of around 0.043 standard deviations. Given that the average standard deviation of income growth is 4.3%, this response is equivalent to an increase in real GDP per capita of around 0.18%. Growth remains above its baseline level for around two to three years following a commodity price shock, before returning to the pre-shock level.

This is a relatively small effect. It suggests that commodity price movements alone may be less important in explaining the volatility of low-income country growth than other explanations. These could include demand, productivity and climatic shocks, international financial flows and conflict episodes.

The remaining panels, (e) and (f), show the response of government spending (G) and fixed investment (I) to the same commodity price shock. We see that the commodity price shock results in higher government expenditure and investment growth.[23] From panel (e) we see that a one standard deviation increase in $CPIND$ results in a statistically significant increase in real government expenditure (G) of around 0.3 standard deviations (or 9.3%) in the following period. Government expenditure remains above its baseline level for at least 4 years following a commodity price shock. Finally, from panel (f) we see that the same commodity price shock results in an increase in real capital formation (I) of around 0.15

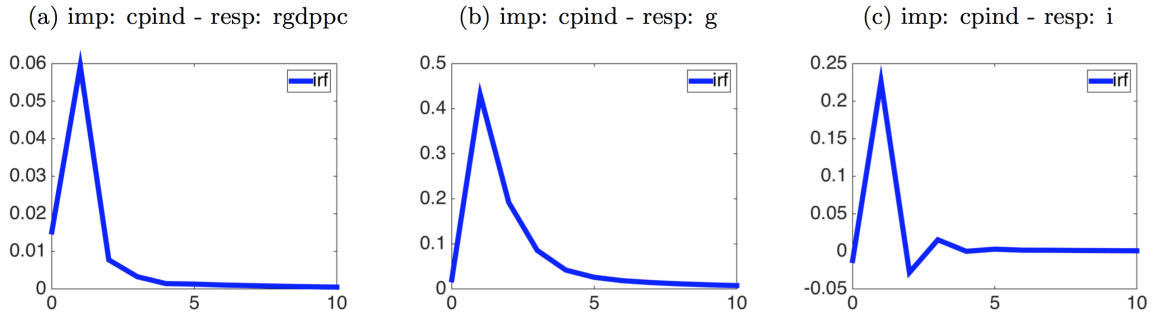
[23]The average standard deviation of the residual is 31.0% from the G equation, and 19.3% from the I equation.

standard deviations (or 2.9%). Investment then returns quickly to its baseline level in year 2, but remains volatile as a result for at least four years.

2.4.2 Block recursive

Figure 2.3 presents the IRFs using the block recursive restrictions rather than Cholesky to retrieve the structural shocks to the panel-VAR system. As discussed in the methodology section, given the nature of the identifying restrictions placed on the model, it is not possible to retrieve all of the structural errors from the panel-VAR. Here, I focus on the response of national level variables, which appear in the second block, to shocks to the commodity price index, which appears in the first block.

Figure 2.3: Panel-VAR IRFs - block recursive identification strategy



I continue to scale the response functions by the standard deviations in the data. The results suggest that whilst both identification strategies broadly agree on the direction of the IRFs, the block recursive strategy delivers larger response functions for real income, government spending and investment in low-income countries. Given the more plausible restrictions placed on the $\boldsymbol{B}^{-1}$ matrix under block recursiveness, the difference in the IRFs is evidence that the simple recursive ordering approach underestimates the impact of commodity price movements on macroeconomic conditions in resource rich, low-income countries.

Panel (a) of figure 2.3 shows that, under block recursion, a one standard deviation shock to commodity prices leads to an immediate increase in per capita GDP growth of around 0.06 standard deviations, equivalent to 0.26%. This is larger than under Cholesky and is

due to commodity shocks now being stripped of any endogeneity with respect to global economic conditions.

Panels (b) and (c) of figure 2.3 show the response of government spending (G) and fixed investment (I) to a commodity price shock under block recursiveness. The direction of the IRFs are all in agreement with the results from the Cholesky identification strategy, and the magnitudes of the responses are all larger. From panel (d) we see that a commodity price shock results in a statistically significant increase in real government expenditure (G) of 0.4 standard deviations (or 12.4%) in the following period, compared to 0.3 under Cholesky. From panel (e) we see that the same commodity price shock results in an increase in real capital formation (I) of around 0.23 standard deviations (or 4.4%), compared to 0.15 under Cholesky.

2.4.3 Robustness

One concern may be the inclusion of OPEC members in the analysis who are likely to set their oil production levels strategically in response to the oil price.[24] I therefore re-estimate the IRFs without OPEC members. The results are presented in figures B.2 and B.3 in appendix B. The impact of $CPIND$ shocks on real income growth fall slightly under both identification strategies; from 0.043 to 0.038 under Cholesky, and from 0.06 to 0.042 under block. The impact on government spending and investment are also reduced. The intuition here is that OPEC members may be more able to guide the oil market to their collective advantage, resulting in more beneficial effects of commodity price movements on local economic conditions.

Another interesting experiment is to test the sensitivity of the headline results to restricting the sample to low income countries. I re-run the panel-VAR estimations on an extended sample which includes middle-income, resource rich countries.[25] The results are

[24]The recent revival of OPEC discussions following the low oil price is an example of coordinated action by members in an attempt to guide the oil market. The OPEC members included in the dataset are: Algeria, Ecuador, Gabon, Indonesia, Iran, Nigeria and Saudi Arabia; although Gabon, Ecuador and Indonesia all suspended the membership at some point over the period from 1985 to 2013.

[25]The World Bank classifies countries as middle income if GNI per capita in 2015 is between $1,025$

presented in figures B.4 and B.5 in appendix B. The IRFs under Cholesky present broadly the same picture in this extended sample (see figure B.4). Under the block restrictions, the impacts of commodity price shocks on local variables are again reduced. The impact of commodity price shocks on real income falls from 0.06 standard deviations to 0.04 (see figure B.5), whilst the impacts on government spending and investment also fall to 0.3 and 0.15 respectively. On one hand we might expect more developed economies to be better placed to take advantage of resource windfalls. On the other hand, however, the effect of commodity windfalls is likely to be diluted in more developed countries, whose economies are less dependent on resource exports.

Standard macroeconomic theory suggests that the choice of exchange rate regime is important for the ability of an economy to respond to external shocks. Flexible exchange rates allow for a faster adjustment to real shocks whilst fixed exchange rates insulate the domestic economy from monetary shocks. Under a flexible exchange rate regime for example, a rise in a county's terms of trade would typically lead to an expansion in aggregate demand, putting upward pressure on the interest rate and causing capital inflows. This in turn would lead to an appreciation in the exchange rate and a reduction in net exports, thus mitigating the effects of the initial shock. Under fixed exchange rates the economic impacts of real shocks are magnified. I test this prediction using our panel-VAR. In figures B.6 and B.7 I classify countries as fixed or flexible exchange rate regimes respectively and estimate the impulse response functions separately.[26] The results support the theory. I find that the response of real income, to commodity price shocks, whilst continuing to be short lived, is magnified under fixed exchange rates, by up to four times on average. A similar result exists for government spending and investment.

Finally, I explore the link between countries' exposure to commodity price shocks and

and 12,475. The countries added to the analysis are: Kyrgyz Republic, Lao PDR, Mongolia, Norway, Sao Tome and Principe, South Sudan, Timor-Leste, Uzbekistan, Venezuela, Vietnam, Yemen, Albania, Angola, Azerbaijan, Iraq, Khazakstan, Suriname, Turkmenistan.

[26] I use data from the IMF's Annual Report on Exchange Arrangements and Exchange Restrictions (AREAER) database to classify countries. I omit countries whose exchange rate regimes that fall between the fully flexible or hard fix regimes.

the magnitude of the dynamic responses of macroeconomic variables to these shocks. We may expect the size of the IRFs to depend on the importance of commodity exports to each economy. In effect, countries with larger resource exports should have larger IRFs. I investigate this by splitting the sample of countries into two groups: (i) those with no natural resource exports in 2013 and (ii) those with positive levels of resource exports in 2013. I exclude countries which had extremely large resource exports in 2013, specifically above $US10 million.[27] I compare the magnitude of the response of real per capita GDP growth to shocks in the commodity price index. The results, reported in figures B.8 and B.9 in appendix B, appear to offer some support to this hypothesis. This is particularly visible in the block recursive IRFs presented in figure B.9 where the size of the response of real GDP per capita to $CPIND$ shocks increases as the exposure increases. An exogenous one standard deviation increase in commodity prices results in an increase in real per capita GDP of 0.08 standard deviations in panel (a) to 0.2 in panel (b).

2.4.4 Sectoral value added results

This section investigates the impact that resource booms have on the sectoral structure of resource rich, low-income economies. The aim here is to test the hypothesis that resource booms lead to a surge in demand for non-tradables and so a reallocation of economic activity away from more exportable sectors. Whether this is due to a simple spending effect or a combination of this and a real exchange rate effect will be investigated further in the following section.

I present the results from estimating the SUR fixed effects model in Eq. 2.11. The block recursive orthogonalised structural shocks for the commodity price index ($CPIND_{i,t}$) from the panel-VAR model in the previous section are included as a regressor in the model with up to three lags. The dependent variable ($y_{i,t,j}$) is the share of real value-added in sector j of country i in year t in total real value-added. The commodity shock variable is scaled by its standard deviation, and so the coefficient estimates show the effect of an exogenous

[27] The reason for this is that these countries tend to be large commodity exporters, and thus not representative of the small open developing economy.

commodity price shock equivalent to one standard deviation, or a 17.7% increase. I cluster the standard errors by country, to account for autocorrelation, and year since the treatment, the price shocks, are aggregated at the country level.

Table 2.1 presents the results from estimating the model using data on real value-added from the UN National Accounts database. This data breaks total value-added into seven sectoral categories which correspond roughly to the ISIC 1-digit codes. Columns 1-7 present the results of the model for each sector. The sectors are:

1. Agriculture, hunting, forestry, fishing (ISIC A-B) (agric)
2. Mining, Utilities (ISIC C-E) (min)
3. Manufacturing (ISIC D) (manuf)
4. Construction (ISIC F) (constr)
5. Wholesale, retail trade, restaurants and hotels (ISIC G-H) (whole)
6. Transport, storage and communication (ISIC I) (trans)
7. Other Activities (ISIC J-P) (other)

The first point to note is the slight reduction in sample size, compared to the panel-VAR, for the various regressions at the 1-digit sectoral level. This is due to missing data on value-added at the sectoral level for some countries included in the analysis. The estimates for real value-added cover 30 countries with around 16 years of data for each, giving 570 observations in total.

Positive commodity price movements are associated with contractions in the share of value added in the manufacturing sector, expansions in the transport and other[28] sectors, but no effect on agriculture. A one standard deviation, or 17.7%, exogenous increase in

[28]The other category includes: financial intermediation (ISIC J); real estate, renting and business activities (ISIC K); public administration and defence (ISIC L); compulsory social security (ISIC M); education (ISIC N); health and social work (ISIC O); other community, social and personal service activities (ISIC P); and, private households with employed persons (ISIC Q).

Table 2.1: SUREG FE regression - all sectors

VARIABLES	(1) agriculture	(2) construction	(3) manufactures	(4) minerals	(5) transport	(6) wholesale	(7) other
ln_govexp	-0.423	-0.113	-0.469***	-0.361	0.246*	-0.0882	0.708***
	(-1.572)	(-1.482)	(-3.350)	(-1.077)	(-1.915)	(-0.540)	(-4.037)
ln_inv	-0.331	1.883***	-1.045***	-0.0307	-0.386**	-0.242	0.0669
	(-0.791)	(-8.313)	(-4.612)	(-0.0515)	(-2.217)	(-0.708)	(-0.199)
infl_defl	-0.000264	-0.000174**	0.000315*	-0.000172	-0.000123	0.000188	0.000437***
	(-0.819)	(-2.382)	(-1.865)	(-0.681)	(-0.337)	(-1.084)	(-3.125)
ln_rgdppc	-6.899***	-1.476***	4.712***	12.43***	-0.195	0.397	-6.432***
	(-5.958)	(-3.694)	(-5.901)	(-5.317)	(-0.380)	(-0.513)	(-4.124)
cshock	-0.0183	0.0217	-0.172**	0.0107	0.0962**	0.0541	0.0707
	(-0.190)	(-0.6)	(-2.070)	(-0.0535)	(-2.275)	(-0.852)	(-0.595)
L.cshock	-0.0978	0.0391	-0.0943	-0.285	0.132***	0.0926	0.255**
	(-1.085)	(-1.082)	(-1.161)	(-1.419)	(-2.824)	(-1.466)	(-2.232)
L2.cshock	0.0154	-0.0212	-0.0717	-0.397**	0.123***	0.0902	0.252**
	(-0.194)	(-0.574)	(-0.954)	(-2.126)	(-2.863)	(-1.42)	(-2.113)
L3.cshock	0.00777	0.0461	0.0135	-0.394**	0.0292	0.0687	0.176
	(-0.0907)	(-1.299)	(-0.175)	(-2.016)	(-0.719)	(-1.105)	(-1.356)
initial rvashare_agric	0.569***						
	(-6.444)						
initial rvashare_constr		-0.0526*					
		(-1.725)					
initial rvashare_manuf			0.489***				
			(-18.76)				
initial rvashare_min				0.461***			
				(-3.916)			
initial rvashare_trans					0.851***		
					(-7.471)		
initial rvashare_whole						0.285***	
						(-3.8)	
initial rvashare_other							0.277***
							(-2.87)
Observations	577	577	577	577	577	577	577

Robust z-statistics in parentheses
*** p¡0.01, ** p¡0.05, * p¡0.1

Shock = cshock; Country and year fixed effects included; Low-income, resource-rich countries; standard errors clustered by country and year

the country specific commodity price index results in an immediate reduction in the manufacturing sector of around 0.17 percentage points (pp). The share of value-added in the transport sector increases by 0.10pp immediately, and a 0.13pp after two years, whilst the share of value-added in the other sectors, which consists of a number of non-tradable goods and services, increases by around and 0.26 pp after a year or two. Interestingly, the mining sector also contract, albeit after two to three years.

These results lend support to previous empirical findings of a 'de-industrialisation' effect of resource booms in low-income countries. The sectors that appear to benefit most from commodity booms are transportation, finance, real estate and public administration. This may be due to the complementaries that exist between these sectors and the natural resource sectors in low-income countries.

The manufacturing sector as a whole, which is typically thought of as a tradable sector suffers. This may be due to a combination of the 'resource movement' and 'spending effect' channels, in which the tradable nature of manufacturing output leaves the sector susceptible to de-industrialisation pressures following a resource boom. I explore one of the channels through which this de-industrialisation effect may operate in the following section; namely, the Dutch disease.

Because national level control variables on the right hand side of Eq. 2.11 may be endogenous to sectoral value-added, I have re-estimated the SUR models excluding them as a robustness check. The results are presented in table B.2 in appendix B. There are some small, but ultimately inconsequential, changes to the results. The negative effect of shocks on manufacturing increases to -0.19 pp. Finally, I exclude OPEC members from the sample and re-estimate the sectoral regression. The results are in table B.3 in appendix B.

Again, the results do not change much, although the effect on manufacturing is increased to -0.22 pp. This is consistent with the earlier result that OPEC members appear to see larger macroeconomic responses to commodity price shocks.

2.5 Dutch disease mechanism

So far the results have shown that whilst resource booms are associated with higher levels of income growth, government spending and investment over the short-run, they also lead to significant reallocation of economic activity away from tradable sectors. This section explores the possibility of a Dutch disease mechanism at work, in which the real exchange rate appreciates following resource booms thus leading to de-industrialisation. A formal exposition of this link was first presented by Corden & Neary (1982) who argue that resource booms lead to an appreciated real exchange rate through two channels. A spending effect in which an increase in domestic income levels leads to higher demand and spending, leading to higher prices in the non-tradable sector (prices for tradables are fixed in world markets). A resource movement effect, which the booming resource sector attracts investment and so pulls capital and labour away from other parts of the economy. Prices decline faster in the tradable sector, where prices are fixed in world markets. Both effects lead to a fall in non-resource, non-tradable output and real exchange rate apprecaition as the price of non-tradables increases relative to tradables to clear the domestic market. This section prevents evidence that the contraction in manufacturing presented in the previous section does indeed operate through an exchange rate channel.

The approach taken here is relatively straightforward. I regress the share of manufacturing real value-added in total value-added on a measure of the real exchange rate. The key step is to instrument for changes in competitiveness using the orthogonal commodity price shock variable from the block recursive panel-VAR model. I use this to test whether the de-industrialisation effect of resource booms operates through a real exchange rate channel.

Following Rodrik (2008), I construct a measure of the real exchange rate (e) using data on nominal exchange rates (E) and purchasing power parity conversion factors (PPP), both from the Penn World Tables (PWT) version 6.2 as follows:

$$\ln(e_{i,t}) = \ln\left(\frac{E_{i,t}}{PPP_{i,t}}\right) \tag{2.10}$$

where an increase in E is a depreciation. A key issue here is that as countries develop, and income levels rise, there are natural appreciation pressures on the real exchange rate. This is known as the Balassa-Samuelson effect[29] and is well documented in the literature. I control for this effect by including national level regressors, including real per capita GDP. I estimate the following equation using two-stage least squares (2SLS) and generalized method of moments (GMM):

$$\begin{aligned} rva_sharemanuf_{i,t} =& \alpha + \beta_1 G_{i,t} + \beta_2 I_{i,t} + \beta_3 INFL_{i,t} + \beta_4 RGDPPC_{i,t} \\ &+ \beta_5 rva_sharmanuf_{i,t}^{ini} + \beta_6 RER_{i,t} + f_i + \eta_t + \epsilon_{i,t} \end{aligned} \tag{2.11}$$

where $rva - sharemanuf_{i,t}$ is the share of real value i added in total value-added in country i in year t, f_i and η_t are country and year fixed effects, G, I, $INFL$ and $RGDPPC$ are the country level control variables as before, and $RER_{i,t}$ is the log of the real exchange rate, instrumented for using the commodity price shock variable ($CSHOCK_{i,t}$), the inflation rate ($INFL_{i,t}$) and log of real per capita GDP ($RGDPPC_{i,t}$) to account for the Balassa-Samuelson effect.

The results are presented in table 2.2. In column (1) I exclude all national level controls (G, I, $INFL$, $RGDPPC$) in the main equation and estimate the model using 2SLS. Column (2) includes national level controls and column (3) re-estimates this model using GMM[30]. In each case, the real exchange rate variable is instrumented for using commodity price shocks. All three models agree in the direction of the effect: a one standard deviation fall in the real exchange rate (i.e. an appreciation) caused by positive commodity price shocks results in a reduction in the share of manufacturing value-added of between 5 and 8 percentage points.

[29]The Balassa-Samuelson effect is the observation that richer countries have systematically higher consumer prices and so more appreciated real exchange rates.

[30]GMM uses assumptions about specific moments of the random variables instead of assumptions about the entire distribution, which makes GMM more robust than maximum likelihood (ML), at the cost of some efficiency. The main assumption underlying GMM is that the sample moments are provided by the data and are set equal to the population moments

Table 2.2: IV regression - real exchange rate channel

VARIABLES	rva_sharemanuf	rva_sharemanuf	rva_sharemanuf hline
ln_rer	8.181***	5.086**	4.899**
	-3.462	-2.386	-2.231
inirva_sharemanuf	2.094***	1.729***	1.683***
	-6.027	-3.917	-3.46
ln_govexp		0.111	0.0826
		-0.615	-0.441
ln_inv		-0.34	-0.356
		(-1.180)	(-1.255)
infl_defl		-0.000256	-0.000271
		(-0.803)	(-1.457)
ln_rgdppc		1.537	1.819
		-1.427	-1.582
Constant	-14.67***	-15.29*	-15.64
	(-4.203)	(-1.675)	(-1.608)
Observations	580	580	580
R-squared	0.9	0.932	0.933
Estimation	2SLS	2SLS	GMM
Controls	NO	YES	YES
Year and Country FE	YES	YES	YES

Robust z-statistics in parentheses
*** p¡0.01, ** p¡0.05, * p¡0.1

2.6 Nigeria and the oil price collapse

This chapter has argued that exogenous commodity price movements can be used as a proxy for resource shocks in resource rich, low-income countries. Over the short run, booms tend to lead to modest increases in growth, government expenditure and investment in these countries. However, they also result in significant short run structural transformation away from the exportable sectors towards more non-tradable activities; perhaps as a result of movements in the real exchange rate that undermine competitiveness.

But how large are these effects in practice? To put these results into context, this section analyses the recent oil price collapse at the beginning of 2015. I focus specifically on Nigeria as a major oil exporter in my sample, and construct the implied change in the country specific price index after the shock. I calculate the impulse effect on growth, government spending and investment as a result of the shock, using the results from the block recursive identification scheme.

The first step is to identify how the recent oil price collapse compares to the normal variation in the data. The UK Brent Crude Price fell from 101.8 to 53.9 USD per barrel in the 6 months between 2014q3 and in 2015q1, equivalent to a 47.1% collapse. Over roughly the same period, rich country GDP growth slowed to 1.8% and the US Libor rate fell by .08pp. The slowdown of aggregate demand and looser financial conditions account for around 18 percentage points (pp) of the oil price collapse. I treat the remaining 29% drop in oil prices as the exogenous shock. I also assume that any movement in the oil price is exactly reflected in the movement of its reserve-weighted commodity price index.[31]

Table 2.3 reports the results of simulating a one period exogenous collapse in the price of oil of 29% for Nigeria, equivalent to a 1.6 standard deviation shock. As a result of the oil price shock, Nigeria suffers a contraction of real per capita GDP of 0.4% in the same period; government spending and investment shrink by 20.3% and 7.2% respectively. Table 2.3 also reports actual data for the Nigerian economy in 2014 and 2015 based on IMF WEO data. Between 2014 and 2015 real per capita GDP shrunk by 0.1%, government spending by 5.9% and investment by 3.8%. Whilst the direction of movement of the key variables are in line with the model predictions, the magnitudes of the simulated effects are somewhat larger than the actual data suggests.

Table 2.3: Nigeria projections vs outturn

Simulation				Actual Outturn			
Period	RGDPPC	G	I	Year	Real per capita GDP	Real inv	Real gov exp
0	-	-	-	2014	3.5	12.7	0.2
1	-0.4	-20.3	-7.2	2015	-0.1	-5.9	-3.8

2.7 Conclusion

Over the past decade extractive industries in Africa have emerged as a powerful engine of economic growth, driven in part by surging demand for natural resources in emerging mar-

[31] I believe this is a sensible assumption as oil reserves make up almost all of Nigeria's natural resource reserves. Nigeria produces a small amount of natural gas but reserves are insignificant compared to oil.

kets. The continent's hydrocarbon and mineral potential has been a powerful magnet for foreign investment with new exploration revealing vast reserves, particularly in sub-Saharan Africa. Recent discoveries of natural gas in Mozambique and Tanzania, and oil in Ghana, Uganda and Madagascar are cases in point. The challenge facing the region now is how to transform these temporary resource windfalls into permanent economic and human development successes.

This chapter presented an empirical investigation of the dynamic macroeconomic response of resource rich, low-income countries to global commodity price fluctuations; specifically minerals and hydrocarbons. It provides new evidence on the relationship between resource booms, economic growth and structural transformation in these economies. The endogeneity problems inherent in the cross-country literature on resource booms make it challenging to identify a causal effect of resource wealth on economic development. This work improves upon previous estimates by identifying exogenous, supply driven commodity price shocks for resource rich, low-income economies. It then uses these shocks to investigate the potential Dutch disease effects of resource booms at the sectoral level.

I used a panel-VAR approach to identify exogenous shocks to global commodity prices and to estimate the dynamic, short-run response of these shocks on economic growth, government spending and domestic investment. The potential for short-term Dutch disease effects is investigated using a seemingly unrelated regression (SUR) framework. Finally, I presented preliminary evidence that the channel through which the link between resource booms and sectoral growth operates is the real exchange rate.

Focusing on resource rich, low-income economies the results suggest that positive commodity price shocks raise per capita income levels, investment and government spending in the short to medium term, as well as lowering inflation. These effects, however, are quantitatively small, suggesting that commodity price movements alone may be less important in explaining the volatility of low-income country growth than other explanations, such as demand and productivity shocks. Resource booms also result in significant restructuring of

the domestic economy. I show that an empirically robust link exists between resource wealth and the relative growth of tradable or exportable industries. The share of value-added in non-tradable goods and services, such as transport, finance and real estate, expands in response to positive resource shocks whilst value-added in manufacturing contracts. This mechanism appears to operate through the real exchange rate, with resource shocks leading to appreciations which in turn cause a restructuring of the domestic economy. The effects are larger for less developed countries, economies with fixed exchange rate regimes and those that are more dependent on commodity exports

Chapter 3

Surfing a wave of economic growth

3.1 Introduction

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

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

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

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

We measure economic activity using two detailed and geographically disaggregated datasets on night-time lights and population. The first records the amount of light emitted at night around the globe at approximately a $1km^2$ resolution annually, from 1992-2013. This has been used by many recent studies as a geographically disaggregated proxy for economic activity (Chen & Nordhaus (2011); Henderson et al. (2011, 2012); Donaldson & Storeygard (2016)); covering institutions (Michalopoulos & Papaioannou (2013, 2014)), political favouritism (Hodler & Raschky (2014)), infrastructure investment (Jedwab et al. (2017); Jedwab & Moradi (2016)) and poverty (Smith & Wills (2016)) amongst others. The second, from LandScan, records global population at a $1km^2$ resolution annually, by interpolating sub-national population counts using satellite images of buildings, roads and

[2] A note on terminology. A 'surf break' is a location. The quality of a break is determined by geography and is generally stable over time, unless there has been some human intervention (see Section 3.4). A 'wave' is an individual pulse of energy that propagates across the surface of the ocean, and which may be the fleeting subject of a surfer's sport. Groups of waves are 'swells'. A wave's quality depends on the wind at its source and end, and the break where it is ridden (see Section 3.3.1).

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

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

land-cover, amongst other things.

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

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

The second experiment asks, does discovering a new surf break (or seeing one disappear) alter nearby economic growth? We do this using two approaches. The first conducts event studies on three recent discoveries of surf breaks, due to competitions organised by Surfer Magazine and the World Championship Tour; and two disappearances, due to construction of a coastal road and dredging of a river-mouth. We find that lights grew (or fell) 2.2 percentage points per annum faster after discovery (or disappearance), relative to the

[5]An earlier working paper version of this chapter, publicised in The Economist (The Economist (2016)), BBC World The BBC (2016) and the Financial Times (The Financial Times (2017)) amongst others, used a less conservative polynomial model (see Section 3.3.2). In this version of the paper, we include three natural experiments instead of one, and focus on the most conservative results.

global average. The second conducts an event study around a technological discovery: Rip Curl's 2007 invention of battery-heated wetsuits. We find that lights grew 2.7 percentage points faster than the global average in the 5km surrounding 83 cold-water breaks above 55 degrees latitude. This suggests that discovering a new natural amenity, or destroying an existing one, may have a large short-term effect on growth.

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

Collectively these three experiments show that natural amenities raise economic growth. Natural amenities, like high-quality surf breaks, are usually not paid for directly. Instead, they attract tourists and other types of transient as well as more permanent visitors, who travel and spend money to be able to take advantage of these amenities. This augments the productivity of physical capital, by providing a basis for recreational industries and tourism. This then acts as a pull factor for other sectors and industries that have up or downstream linkages with tourism. With more investment, increased incomes, lower transaction costs and productivity spillovers, the agglomeration of economic activity near surfing breaks results in higher economic growth. Experiment II shows that the increment to growth is larger when the amenity is initially discovered, and Experiment I shows that the effect on growth can persist for decades as surfing originally became popular in the 1960s.

High quality natural amenities may also help solve the coordination problem behind locating sectoral clusters (Rodrik (1996, 2004)). For example, along a stretch of coast a

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

This work contributes to the 'geography hypothesis', which has a long history in explaining the pace of economic growth. This includes channels like suitability for agriculture (Marshall (1890)), exposure to disease (Diamond (1999); Sachs (2000, 2001)), coastal access (Gallup & Sachs (2000)), narrowing of export markets (Malik & Temple (2009)) and frequency of natural disasters (Hsiang & Meng (2015)). It is typically contrasted with the 'institutions hypothesis' (North (1989, 1990); Acemoglu et al. (2002)), which argues that institutional technology is the major driver of growth[6]. By using spatial and temporal variation in the quality of natural amenities at a fine scale we find that geography does play a role in explaining local economic growth. This is solely due to spillovers as surfing does not contribute directly to economic activity. Faber & Gaubert (2015) find a similar result, by studying the effects of tourism in Mexico.[7] We also find that the importance of geography varies by institutional quality, measured using political stability and ease of doing business indices, which provides some evidence to reconcile these competing hypotheses. At a local level both geography and institutions matter, and so does the way they interact.

It is also well-established that natural geography plays a role in determining the lo-

[6]Garcia-Ponce & Wantchekon (2014) put the two together and argue that the geographic characteristics of some countries in Africa may have resulted in different types of institutions post-colonialism, which in turn have an affect on growth rates

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

cation of economic activity. The New Economic Geography literature draws a distinction between 'first-nature geography', concerning natural advantages, and 'second-nature geography', concerning endogenous forces of agglomeration and dispersion (see Fujita et al. (2001); Redding (2010)). Natural advantages, like rivers, ports and resource endowments, have been shown to cause agglomeration as they directly reduce trade and input costs (Ellison & Glaeser (1997, 1999); Ellison et al. (2010); Redding (2010); Redding & Rossi-Hansberg (2016)). Natural amenities, like coastlines, mountains and lakes, have also been found to anchor the location of high-income suburbs (Lee & Lin (20157)). Path dependence is important, as natural advantages have persistent effect after they become obsolete, like portage sites near rivers (Bleakley & Lin (2012, 2015)), and existing agglomeration can prevent natural advantages being exploited (Michaels & Rauch (2016)). This study finds that the benefits from surf breaks accrue to nearby towns, suggesting that natural advantages benefit existing areas of economic activity. We also find that when breaks are discovered or disappear, economic growth hastens or slows accordingly.

Finally, this work also contributes to the recent literature on the local effects of natural resources (see reviews by Cust & Poelhekke (2015) and van der Ploeg & Poelhekke (2017)). These tend to focus on the effects of non-renewable resources with market prices, like oil, gas and minerals (both short run, e.g. Aragón & Rud (2013); Caselli & Michaels (2013); Allcott & Keniston (2014); and very long run, e.g. Dell (2010)). Complementing those studies we find evidence that non-market, renewable resources can also have a positive impact on the local economy, in both the short and long run.

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

happens when the surf is good?'. Section 6 concludes by discussing policy implications and extensions.

3.2 Data and measurement

We use two main datasets to study how surf breaks affect local economic growth: on the location and characteristics of over 5000 surf breaks and on night time light emissions at a $1km^2$ resolution. We use local population estimates to identify nearby towns and cities and study population changes; measures of urban areas and political/economic activity to add granularity to results; and data on atmospheric pressure and wave heights to study how good, large waves interact with our headline results.

3.2.1 Surf breaks

WannaSurf (www.wannasurf.com) is an online 'world surf spot atlas'. It records the location, quality, type, accessibility, coastal and oceanic characteristics of 5,288 surf breaks around the world, which we collect using Python (figure 3.1). Of these we drop 137 for which the data on quality is either missing or rated 0 stars ("choss"), leaving 5,151 breaks in our dataset.

The main characteristics used in this study are surf break location and quality. The GPS coordinates of each break are provided using the WGS84 datum. Accuracy ranges from 3m-200m depending on the original datum used to record the location. Each break is assigned one of five quality ratings, which are verified by independent experts and range from "sloppy" (1-star) to "totally epic" (5-star) (see figure 3.1). Quality describes the average physical quality of the waves at the location, and in most cases is fixed over time (we study rare examples of changing qualities in Experiment II). It does not capture crowds, ease of access, surrounding views, etc., which are recorded separately by WannaSurf. In the first instance data on a new break is nominated by one of the website's 78,000 registered users. The data is then checked and monitored for accuracy by one or more WannaSurf Regional Correspondents. For reference, of the eleven locations on the surfing World Championship Tour seven are 5-star, three are in popular surfing areas with a dense concentration of 4-star

Figure 3.1: Global surf breaks



Overview of surf break locations in WannaSurf database.

breaks, one is 3-star (in central Rio de Janeiro, to promote the sport), and crowds range from "empty" to "ultra-crowded".

Measurement error may be a problem if the quality ratings on WannaSurf are inaccurate. Surfers may bias ratings down to discourage crowding at their favourite breaks. This is not a concern for four reasons. First, we use quality ratings that are verified by independent WannaSurf experts, and ignore the poll-based 'user ratings'[8]. Second, WannaSurf removes legitimately secret locations. Third, a downward bias would overstate the growth potential of low-quality breaks, making our estimates more conservative. Fourth, while wave quality is crucial, so too is shared knowledge of it, which is directly measured by WannaSurf.[9]

The data may also suffer some selection bias. Ideally we would have data on wave quality at regular intervals across the world's ∼620,000km of coastline. Instead our breaks

[8] see http://www.wannasurf.com/help/faq/index.html

[9] Oceanologists are not yet been able to predict break quality using geographic and climactic models, due to the complex interactions that determine it at a metre-by-metre scale (see appendix section C.1; though they do forecast wave heights relatively successfully, for example www.swellnet.com.au). Rather than using geographic variables to instrument for break quality, this chapter has the advantage of measuring quality directly.

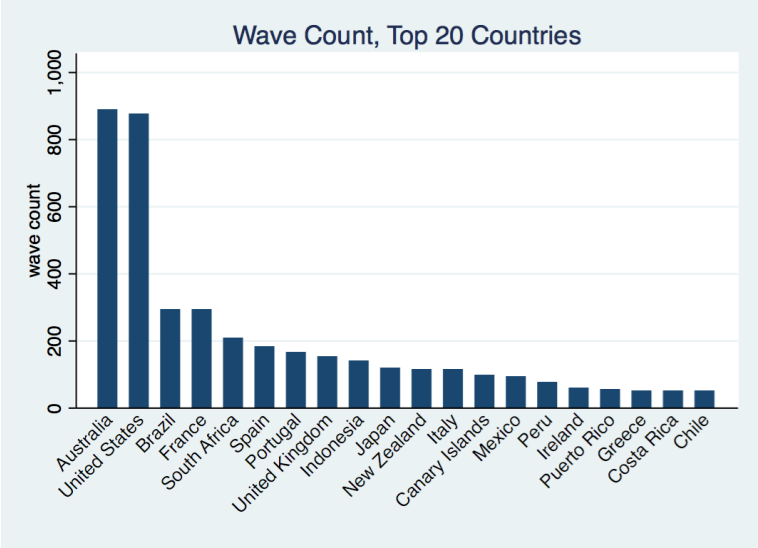
Table 3.1: Surf break quality

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

Breakdown of WannaSurf surf breaks by quality.

are distributed across 146 countries, but are concentrated in Australia (888 breaks) and the US (878) (see figure 3.2). Break quality varies across countries due either to exogenous variation or to selection.[10] Low-quality breaks also tend to have more light than high-quality breaks, which we attribute to surfers entering accessible low-quality breaks into the database, but ignoring isolated ones. We discuss robustness to selection bias in Section 3.3.3.3 and appendix sectionC.4.3.

Figure 3.2: Surf breaks - top countries



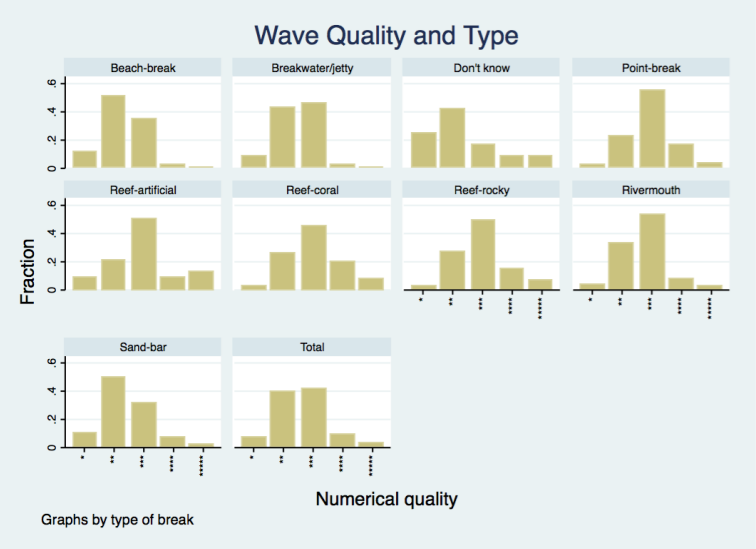
WannaSurf data covers 146 countries, though data on surf breaks is concentrated in Australia and the US.

WannaSurf records a variety of other characteristics for each break. These include variables on accessibility ("Distance", "Easy to find?", "Public access?", "Crowd"), difficulty

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

("Experience"), the type of wave ("Frequency", "Type", "Direction", "Length", "Bottom", "Power") and oceanic conditions ("Good swell direction", "Good wind direction", "Swell size", "Best tide"). Of these we use the "Type" variable, which indicates whether the shoreline is a beach, a reef, a river-mouth, a headland (point-break) or a breakwater, to test whether other omitted geographic variables influence our results (see figure 3.3 and table C.7).

Figure 3.3: Surf breaks by geographic type



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

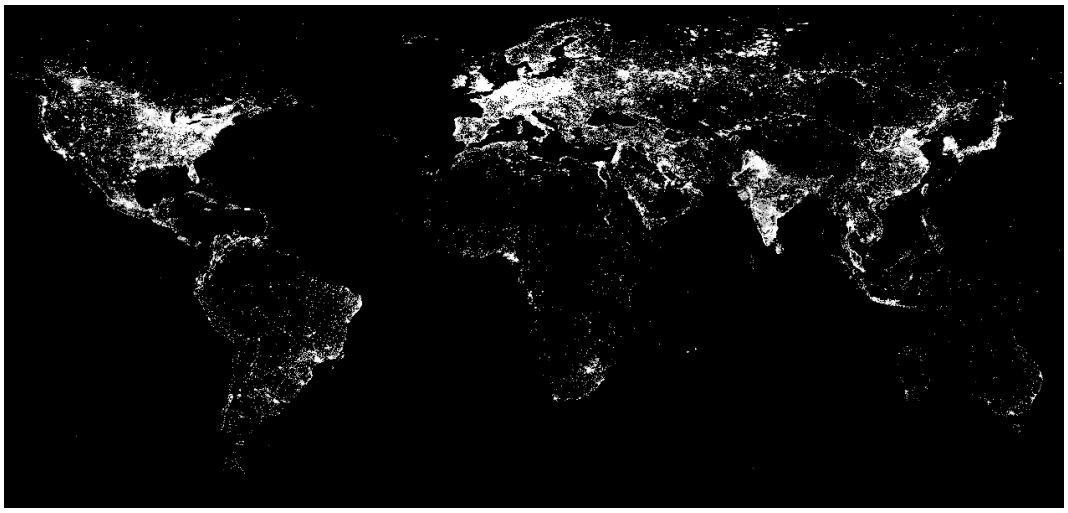
3.2.2 Night-time lights

The Defence Meteorological Satellite Program's Operational Linescan System (DMSPOLS) uses satellites to record the average annual night-time light intensity around the world, from 1992-2013 (figure 3.4a). The data is provided at a resolution of 30x30 arcseconds (approximately 1 square kilometre near the equator), and ranges from 0 to 63. The data is constructed by overlaying all daily images over the course of a year, discarding those that are obfuscated by cloud cover, lightning, aurora, etc. for a given pixel.

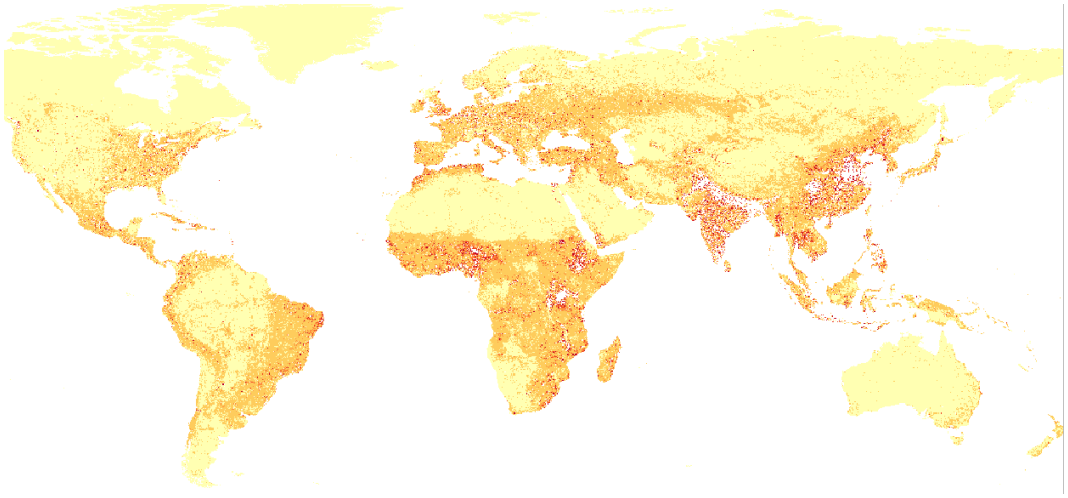
There is a strong link at the national level between the growth of GDP and mean light intensity (Doll (2008); Henderson et al. (2012); Michalopoulos & Papaioannou (2014)). This

Figure 3.4: Night-time lights and population data

(a) Data on night-time lights (DMSP-OLS)



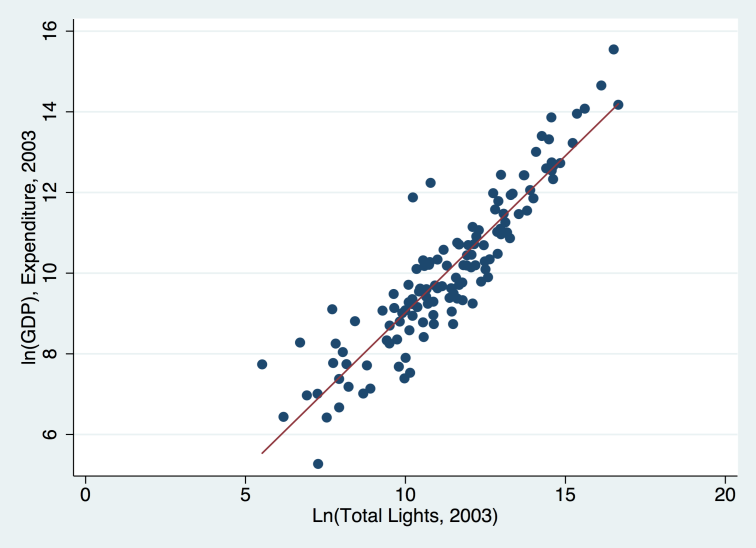
(b) Data on population (LandScan)



is illustrated in figure 3.5, which plots the log of the sum of light readings by country against two measures of log PPP-adjusted GDP: based on expenditure and production (Penn World Tables 8.1). The associated regressions yield an adjusted r-squared of .82 and .80 respectively. We make use of the high spatial resolution of the data to study economic activity at a sub-national level, as has been done in a number of other studies (Chen & Nordhaus (2011); Michalopoulos & Papaioannou (2013); Hodler & Raschky (2014); Jedwab et al. (2017); Jedwab & Moradi (2016); Smith & Wills (2016)).

We use the night-time light data to measure real economic growth in two ways: illumi-

Figure 3.5: PPP-adjusted GDP vs Night-time lights



PPP-adjusted GDP vs Night-time lights (in logs), 2003 (see Smith & Wills (2016)).

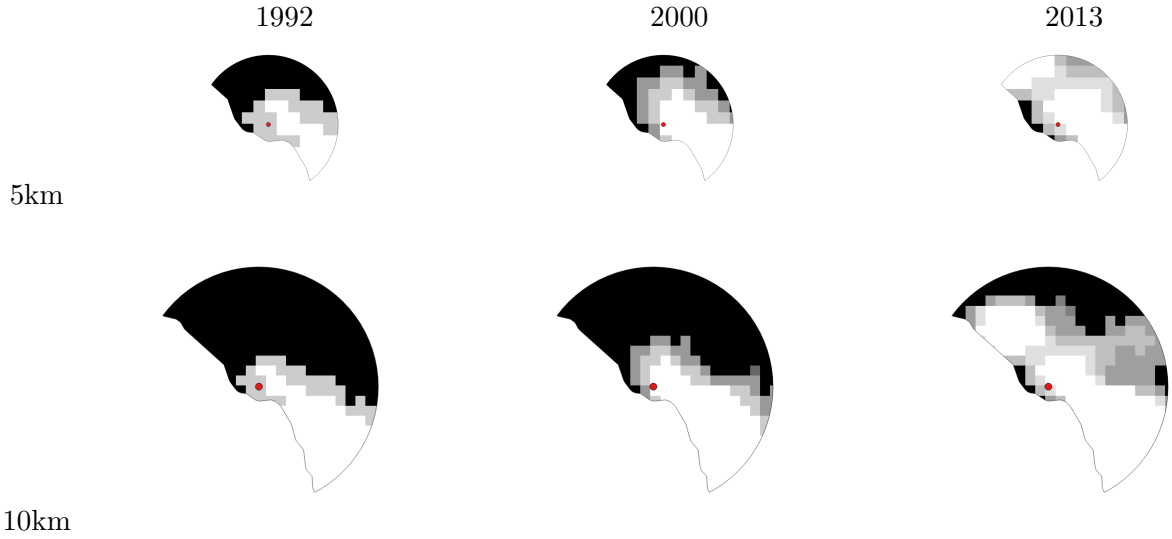
nation in the immediate vicinity of a surf break and illumination in nearby towns.

Illumination in the immediate vicinity of a break is measured using luminosity in concentric circles of various radii. We draw these circles at 5km, 5-10km and 10-50km around each break, clipped along the coastline (see figure 3.6), and take the total illumination on the remaining land in each circle for each year. Coastlines are defined by a shape file from VDS Technologies, a private mapping firm, which may be subject to some small measurement error but it will be uncorrelated with break quality. Circles may overlap, so we include all overlapping areas in all breaks they are close to. This is the most conservative approach because high quality breaks will raise light growth around nearby low quality breaks, making the null hypothesis more difficult to reject. An alternative approach, based on treating individual pixels with the quality of nearby waves, is described in appendix section C.4.1.

Illumination in nearby towns is measured by endogenously locating towns by their population density, using data outlined below. A town is defined by a perimeter enclosing cells with a population density of 300 persons per square kilometre or more (see figure C.6).[11]

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

Figure 3.6: Illumination growth in southern Morocco



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

Each break is linked to two towns: the closest town and the largest town within 50km radius (based on total population).

Night-time light data is subject to some issues that are relevant for this study. First, 'topcoding' refers to pixels with a light reading of 63, beyond which we cannot distinguish levels of economic activity. This is not important in our data as it occurs for less than five 'novelty' breaks in the centre of cities, like a wave park in Kuala Lumpur and a river-wave in central Munich. Second, light data includes significant luminosity readings from gas flares, which do not reflect economic activity. To control for this we drop all cells with gas flare activity according to the provider of the lights data (the NOAA Earth Observation Group), and trim any observations over water. Third, light data is affected by overglow (or 'blooming'), where light is recorded in pixels away from its origin, and is magnified over terrain like water and snow (Doll (2008)). Townsend & Bruce (2010) find that overglow is linearly proportional to lit area, which is consistent with a physical model for atmospheric scattering. We clip our observations around the coastline, and in appendix section C.2 confirm that overglow changes linearly with light, so will not bias our study of light growth. Fourth, the satellites used to construct the data change in 1994 (F11), 1995 (F13), 1997

(F14), 2000 (F15), 2003 (F16), 2007 (F17), and 2010 (F18) - which we average in years with multiple satellites, and the effectiveness of the sensors diminishes over time. To control for this we use year fixed effects.

3.2.3 Population

Oak Ridge National Laboratory produces the LandScan dataset which provides annual mid-year spatial population counts at a 30x30 arcsecond resolution from 2000-2013 (figure 3.4b). It is similar to the Gridded Population of the World data from NASA's Socioeconomic Data and Applications Center (SEDAC), which measures population at a 30 30 arc second resolution in 1990, 1995 and 2000, and has been used by Dell (2010) and Alesina et al. (2016) amongst others. LandScan records estimates of the ambient, or daily average, population in each grid cell. Importantly, this excludes intermittent population such as tourists or temporary relief workers, and may not reflect things like seasonal migrations or refugee movements. This allows us to focus explicitly on spillovers to the permanent population. The estimates are based on aggregate data for second order administrative units compiled by the International Programs Center of the US Bureau of Census. The aggregate data is distributed throughout the grid according to a likelihood model that uses inputs including data on elevation, land cover, roads, night lights, coastlines, settlements, and high resolution satellite imagery, among other sources.

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

3.2.4 Urban and rural classification

SEDAC also provides an 'Urban Extents Grid', from its Global Rural-Urban Mapping Project (GRUMP), which uses 1995 population estimates to classify each square of a 30x30 arc second global grid as either urban or non urban. The classification is based on contiguous lit squares (as of 1995) and settlements known to hold at least 5000 people, and agrees with urban extents based on DHS surveys (Dorélien et al. (2013)).

3.2.5 Political stability and business environment

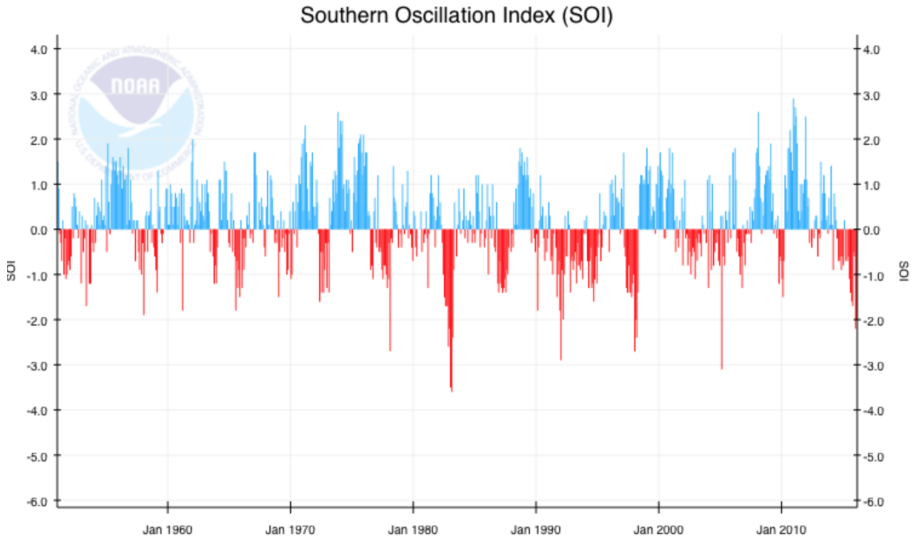
The World Bank provides data on political stability in its Worldwide Governance Index, and on the ease of doing business in its Doing Business Survey. We collect countries into four groups with similar numbers of breaks based on their 2014 scores on each. Table C.5 groups countries into bins by political stability, and table C.4 groups them by the quality of their business environment, with bin 1 being "very low" and 4 being "high".

3.2.6 El Niño and wave heights

The National Oceanic and Atmospheric Administration records the monthly Southern Oscillation Index (SOI), based on the difference in the standardised sea level pressure between Tahiti and Darwin, Australia (see figure 3.7). An El Niño event is defined as any where $\text{SOI} < -0.7$ for three or more consecutive months, which gives events in 1992, 1993, 1994, 1997, 1998, 2002, 2006, 2009, and 2010.

Australia's Commonwealth Scientific and Industrial Research Organisation (CSIRO) provided data on the monthly mean significant wave height from its Centre for Australian Weather and Climate Research (CAWCR) Wave Hindcast on a global 24 arcminute grid, from 1992-2013 (Hemer et al. (2011)). The mean significant wave height is defined as the average height (trough to crest) of the upper third of the waves in the wave-field. They also provided data on the monthly wave height anomaly, which is the deviation from the annual cycle at each grid point.

Figure 3.7: NOAA Southern Oscillation Index.



The NOAA's monthly Southern Oscillation Index.

3.3 Experiment I: Cross section

> *"Surfers are the trendsetters, then the other tourists follow"* - Tarik Senhaji, Director General, Moroccan Sovereign Wealth Fund (2016)

3.3.1 Identification

In experiment I we ask the question: 'Do good surf breaks contribute more to economic growth than bad breaks?'. Surfing offers a clean natural experiment because the quality of surf breaks is exogenously determined by a careful calibration of climatic, bathymetric and geographic conditions, which we exploit in the following three experiments.

Waves are created by wind acting on the surface of the ocean far out to sea. If the wind is strong and consistent enough the waves can travel up to 15,000km across the ocean, until they break against the shallow sea floor. The quality of a surfing wave depends on a complex and finely-calibrated interaction of time-variant factors, like the strength, direction and duration of winds where the wave originates and breaks; and time-invariant factors,

like the sea floor bathymetry over the distance the wave travels, and the shape of the coastline where it breaks. These factors determine the quality of a surf break, which will be zero for the vast majority of the world's coastline (for more detail see appendix section C.1).

Therefore, while a surf break's existence may be related to economic activity (due to trade advantages of coastlines, etc.) the quality of the break is essentially random. Experiment I exploits this randomness to determine whether areas near high-quality waves grow faster than those near low-quality waves. The control group is areas near 1-star waves, which is a relatively high hurdle as they are also coastal and of sufficient interest to surfers to appear in the WannaSurf database. This identification strategy may still be subject to issues like endogeneity, selection bias and measurement error, which we consider in our robustness tests in Section 3.3.3.3.

An example may be helpful. In less than a mile along the north shore of Oahu, Hawaii, lie three of the world's most famous, 5-star reef-breaks: Backdoor, Pipeline and Sunset Beach (see figure 3.8). Every December the world surfing tour concludes there, as champion surfers test themselves against long-range north Pacific swells that break against a succession of shallow reefs. In between Pipeline and Sunset lies Kammieland: another reef break on exactly the same beach, facing the same direction, and receiving the same swell, which is only rated 2-stars due to a complex interaction of water flowing over that reef and others. The rest of the breaks on this uninterrupted stretch of sand do not even warrant a name.

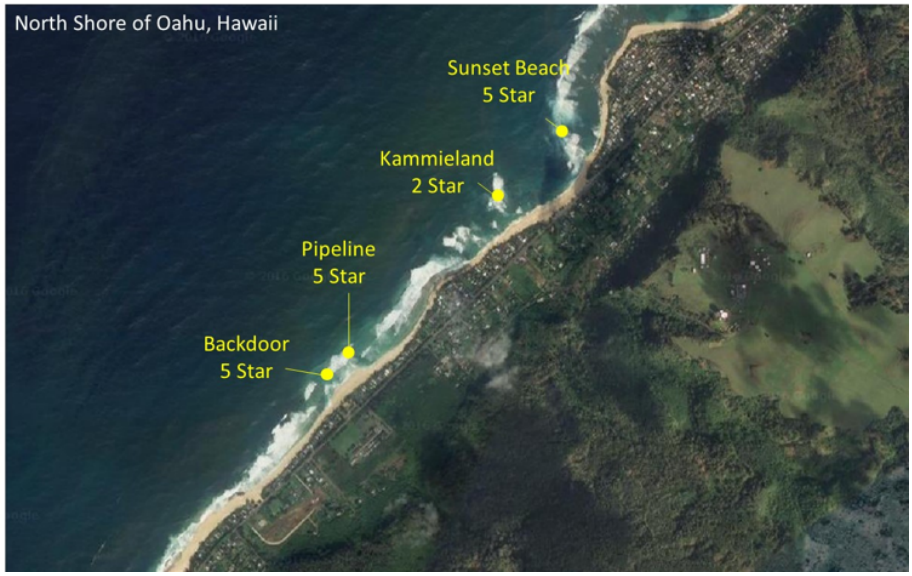
3.3.2 Estimating equations

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

$$\ln(Y^d_{i,2013}) - \ln(Y^d_{i,1992}) = \alpha + \beta Q_i + F_i + Z_t + \epsilon_{i,t} \tag{3.1}$$

where $Y^d_{i,t}$ is the light/population within d kilometres of break i in year t, $Qi \in [1, 5]$ is the star quality rating of break i which we treat as both continuous and an indicator, F_i is

Figure 3.8: North Shore Oahu, Hawaii



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

zone or country fixed effects, and Z_t is year fixed effects. When the quality indicator takes the form of a 0-1 dummy variable, where 1-star waves are coded as the base category, the regression becomes a difference-in-difference specification. The counterfactual is the change in $\ln(Y_{i,t}^d)$ from 1992-2013 for a 1-star break. This is a relatively high hurdle because 1-star breaks are on the coast and are sufficiently known by surfers to appear in WannaSurf, so our estimates are conservative. We deal with potential spatial correlation by clustering the standard errors by country and zone,[12] and using spatially-robust heteroskedastic and autocorrelation consistent (HAC) standard errors following Conley (1999, 2010); implemented using Hsiang (2010) and Fetzer (2014)).[13]

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

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

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

$$\begin{aligned} \ln(Y_{i,t}^d) &= \alpha + \beta(t)Q_i + \gamma(t) + W_i + Z_t + \epsilon_{i,t} \\ &\text{where} \\ \beta(t) &= \beta_1 t + \beta_2 t^2 + \beta_3 t^3 + \beta_4 t^4 \\ \gamma(t) &= \gamma_1 t + \gamma_2 t^2 + \gamma_3 t^3 + \gamma_4 t^4 \end{aligned} \tag{3.2}$$

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

3.3.3 Results

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

3.3.3.1 Surfing and economic activity

Economic activity near high-quality breaks grew significantly faster than near low-quality breaks from 1992-2013. Table 3.2 (columns 1-3) shows that a 1-star increase in break quality increased lights in the surrounding 5km by 3.6-3.8 log points over our sample ($\sim$0.17 percentage points per annum). This was primarily driven by 3- and 4-star breaks, which grew by 8.6-8.8 log points more than 1-star breaks ($\sim$0.41-0.42 percentage points per annum; columns 4-6). 5-star breaks saw a large but imprecisely estimated increase. Higher

economic activity near good breaks was not just redistributed from surrounding areas, with positive and diminishing spillovers found at least 50km away (see tables C.2 and C.3).

Table 3.2: Total growth regression, 5km

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

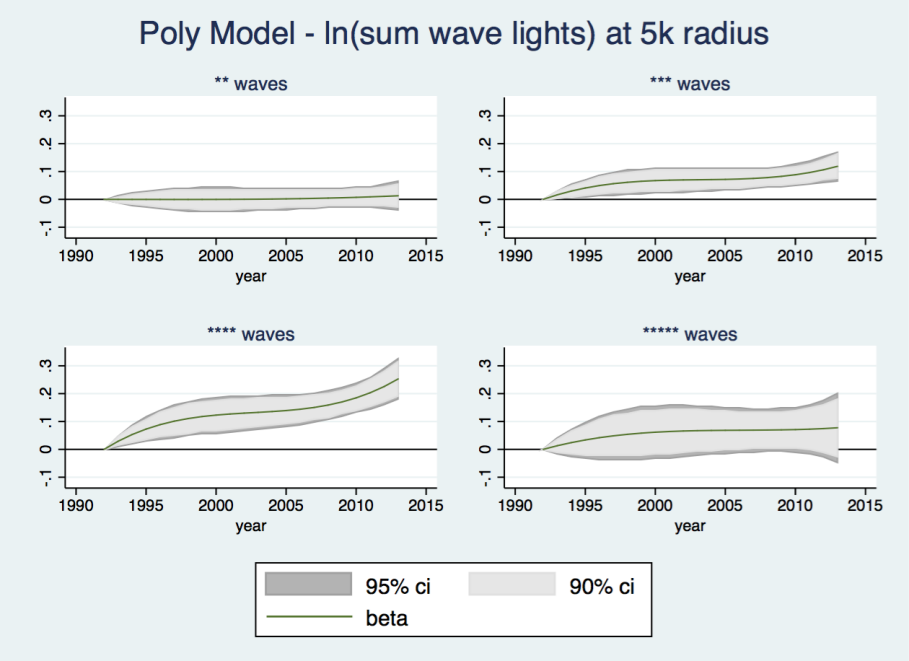
Robust p-values in parentheses
*** p<0.01, ** p<0.05, * p<0.1

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

Figure 3.9 supports these results by showing a consistently positive trend in the level of lights near 3- and 4-star breaks, relative to 1-star breaks, from 1992-2013 (equation 3.2). The effects in this figure are even larger because it applies break fixed effects to light levels, rather than country fixed effects to light growth; though we focus on the table to be conservative.

Together these results indicate that natural amenities can have a persistent effect on economic growth, because modern surfing has been popular since the 1960s. As with other

Figure 3.9: Growth regression, PDL model 5km



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

natural amenities, people do not pay to enjoy surf breaks so additional growth must come through spillovers in the markets for physical capital and labour. Surf-breaks can augment physical capital by providing a basis for tourism and recreational industries. They also augment productivity through the agglomeration of firms and industries around the surf and tourism industries. They may also attract entrepreneurs willing to accept lower profits, if they can regularly surf good waves (see review of the literature on natural amenities by Waltert and Shlapfer, 2010). In a standard Ramsey growth framework this would encourage investment, as places with more natural capital have a higher steady-state level of output. Additional investment causes faster output growth, and the results above suggest that this happens on a timescale of decades.

Surfing also boosts economic growth in nearby towns (see table 3.4). Defining towns by population density (see figure C.6), and controlling for convergence by including the initial level of lights, we find that the closest town to 3-star and 4-star breaks grew by 11

and 13 log points more respectively from 1992-2013, relative to 1-star breaks (∼0.54 and 0.63 percentage points faster per annum). The largest town within 50 km of 3- and 4-star breaks grew an additional 7.1 log points respectively over the period (∼0.34 percentage points faster per annum).[14] 4-star breaks had the largest effect, which is consistent with 5-star breaks requiring a lot of experience to surf (see figure C.7). These results suggest that natural amenities particularly benefit areas of existing economic activity, so that path dependence matters for the location of economic growth.

Table 3.3: Growth regressions in developing countries

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

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

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

Surfing's effect on economic growth is concentrated in developing countries (see table 3.3). We proxy the supply of tourist services by the country's ease of doing business, and demand for tourism by its political stability (with and without Australia and the USA).

[14]The results are robust to replacing the initial level of town lights with country fixed effects for the largest nearby town. The results are no longer significant if both the initial level of town lights and country fixed effects are included.

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

These results can be placed in perspective by converting them into approximate dollar values. The average value of a unit of light can be found by taking global GDP (US$ 2013 PPP, World Bank WDI) and dividing by the global sum of lights in that year. Taking the areas within 5km and 10km of each break, and correcting for overlap, we find that 3-star breaks contributed on average an extra $1.01 million per break per year to the surrounding 5km (2013 PPP) and $2.45 million to the surrounding 10km, relative to 1-star breaks from 1992-2013 (based on the results in table 3.2). Similarly, 4-star breaks contributed $0.53 million and $1.2 million respectively, which is less as these breaks are on average more rural. In total 3- and 4-star surf breaks add $1.36 billion to their surrounding 5km globally per year, and $4.00 billion to their surrounding 10km, relative to 1-star breaks.

3.3.3.2 Surfing and population

The permanent population near high-quality surf breaks fell relative to low-quality breaks from 1992-2013, up to 50km away (table 3.5). The effect was strongest out to 10km from 3- and 4-star breaks, which also saw the largest relative increase in economic activity. As the LandScan data explicitly excludes tourists, this result is consistent with tourists moving towards areas with good surf breaks, raising property prices and driving away locals.

The aggregate fall in the permanent population appears to mask a reallocation of people from rural to urban areas near high-quality breaks. The population in the closest town, and the largest town within 50km of a high-quality break, rose from 1992-2013 (table 3.6). This

was particularly true for the 3- and 4-star breaks. At the same time the rural population within 50km of a high-quality break fell. This suggests that the increase in economic activity near high-quality surfing breaks was accompanied by urbanisation.

3.3.3.3 Robustness

There is a possibility that omitted geographic characteristics both increase the quality of surfing breaks and drive economic growth. We use quality as the treatment because this will be less correlated with economic activity than the existence of a wave. However, there may still be omitted characteristics like openings in reefs and river-mouths, which are favourable for surfing and may also facilitate fishing and trade through boating lanes. To address this we re-run the analysis excluding the breaks that are of disproportionately high quality: reefs (30% of sample), point-breaks (or headlands, 13%) and river mouths (2%) (see figure 3.3 and table C.7). Table C.8 shows that 3- and 4-star breaks see similarly-sized effects to the main specification, though the estimate for 4-star breaks is no longer significant due to the smaller sample size.

Circles around surf breaks may overlap. As noted in Section 3.2.2 we address this by conservatively including overlapping areas in all nearby breaks making it more difficult to reject the null (as good breaks would increase lights at nearby bad breaks). In addition, we re-run our analysis using pixels, rather than breaks, as the unit of observation (see appendix section C.4.1). This allows us to treat each pixel individually with the maximum quality break within 5km, and we find similar results to the baseline specification. When treating each pixel with a count of breaks of each quality within 5km, we find that more low-quality breaks reduce light growth, which disappears for higher-quality breaks.

The sample is dominated by Australia and the USA, which are large, developed countries and account for 17% of breaks each (see figure 3.2). This is a problem if idiosyncratic trends in Australia and the USA bias the overall results, or if lights are a poor proxy for growth in large developed countries. Excluding Australia and the USA from the sample confirms, and in some cases strengthens, the results in the original specification (see table

Table 3.4: Growth in nearby towns

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

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

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

Table 3.5: Population regressions

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

Robust p-values in parentheses

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

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

Table 3.6: Population regressions, closest & largest town

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

Robust p-values in parentheses

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

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

C.8).

Finally, the WannaSurf database exhibits some selection bias, where the lowest quality breaks are more likely to appear in the data when they are in urban areas (see figure C.10). This may bias our results down if urbanisation drives faster growth in lit areas, or up if lights in rural areas converge to those in towns. We control for this in three ways. First, we re-run the analysis using 2-star breaks as the baseline, which suffer less from selection because they have a similar urban share (49%) as the 4-star breaks that drive our main results (48%). We find that the areas near 3- and 4-star breaks also grow faster than those near 2-star breaks (see table C.9). Second, table 3.4 focuses only on towns to avoid any selection bias in rural areas, and controls for convergence using the initial level of lights. The results confirm that lights grow faster in towns near high-quality breaks. Third, we also conduct the pixel level analysis controlling for the initial level of lights in table C.10. The results again suggest that 3- and 4-star breaks have a positive and significant effect on light growth, though the effect for 5-star breaks becomes negative.

3.4 Experiment II: Discoveries

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

3.4.1 Identification

Experiment I compared how quickly lights grew near high-quality surf breaks, relative to low-quality breaks. Experiment II asks the question: 'Does discovering a new surf break alter economic growth?'. It uses, as a counterfactual, areas with no surf breaks at all. To

do this we use a series of event studies to understand how economic activity is affected by discovering a new surf break, seeing an existing one disappear, or gaining the ability to surf a new break because of technological innovation. These can be thought of as a series of event or regression discontinuity studies.

The first approach in Experiment II uses a sample of three discovery and two disappearance events, drawn from a survey of the surfing literature (see table 3.7).[15] We require these events to have four characteristics. First, they must involve high-quality surf breaks. Second, discovered breaks must become known, and disappearances must have been known, to the global surfing community via mainstream media. Third, this must happen at a clearly defined time during our sample period, 1992-2013. Fourth, the event must be exogenous to local economic activity.

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

The other two discoveries come from the first times the surfing World Championship Tour was held at a 'secret spot': in 2006 at La Jolla, Oaxaca, Mexico;[16] and in 2007 at El Gringo, Arica, Chile.[17] Conversely, breaks disappeared in 2005 at Jardim do Mar, Madiera, Portugal, when a rock wall was built to protect a new coastal road; and in Mundaka, Bis-

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

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

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

cay, Spain, when a rivermouth was dredged for boats leading to the cancellation of a 2005 world tour event. Both disappearances were by-products of local investment, making any estimates of economic decline very conservative.

Table 3.7: Discoveries and disappearances

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

The second approach in Experiment II uses Rip Curl's release of the first battery-heated surfing wetsuit in 2007, which made surfing in extremely cold waters much more accessible (Longman (2016)). We study whether this affected economic activity near cold-water surf breaks by treating 2007 as a discovery event for the 83 locations in our database at latitudes above 55 degrees north (none are below 55 degrees south; see table C.12).

3.4.2 Estimating equations

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

$$\ln(Y_{i,t}^{50}) - \ln(\bar{Y_t^{50}}) = \alpha_i + \beta_i t + \alpha_i^D I_i^D + \beta^D I_i^D D_i t + \epsilon_{i,t} \tag{3.3}$$

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

estimates the average marginal effect on the time trend of the break being discovered or disappearing, β^D.

3.4.3 Results - discoveries and disappearances

Table 3.8: Discoveries - diff-in-diff regressions

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

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

Difference-in-difference in average log lights around (1) break discoveries and disappearances using actual event dates (for details see table C.11); (2) break discoveries and disappearances using dates drawn from a uniform U(1992,2013) distribution; (3)-(4) cold water breaks after battery-heated wetsuits were invented in 2007; and (5)-(6) cold water breaks using 1997 as the invention date.

Discovering a new 4- or 5-star break increased (and a disappearance decreased) annual light growth in the surrounding 50km by an average of 2.2 percentage points, relative to the global average (see table 3.8). For robustness we draw alternative event dates from a uniform $U(1992, 2013)$ distribution and find no effect (column 2). Appendix section C.5 illustrates the effects for each individual break by plotting de-trended lights, $\left(\ln(Y_{i,t}^{50}) - \overline{\ln(Y_t^{50})}\right)$. Figure C.11 shows that de-trended light growth more than doubled after each discovery, increasing by 4.7 percentage points at Skeleton Bay, 1.8 at La Jolla and 2.1 at El Gringo, while figure C.12 shows that each disappearance caused light growth to slow by around 2 percentage points (significant at the 5% level). The disappearances are notable because the

road-building and river-dredging that caused them was intended to raise economic growth. This shows that depleting the quality of natural amenities can be an economically important side-effect of infrastructure investment.

3.4.4 Results - Cold water breaks

The invention of battery-heated wetsuits in 2007 increased annual light growth in the 50km surrounding cold-water surf breaks by 2.67 percentage points, relative to the global average (see table 3.8). This result is significant at more than the 1% level. For robustness we test if the invention had hypothetically occurred a decade earlier, in 1997, and find no effect. In table C.13 we re-run the analysis in equation 3.3 but estimate only one $\beta_i = \beta$, $\alpha_i^D = \alpha^D$, and $\beta_i^D = \beta^D$ for all i. This shows that prior to the wetsuit invention, light near cold water breaks was growing 1.6 percentage points p.a. slower than the global average. After the invention it grew by 1.25 percentage points faster.

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

3.5 Experiment III: Large waves

> *"Getting ready for swells is one thing, but getting ready for a season of swells, like El Niño, is a whole different thing. It's exciting to think that we might possibly score this year'* - Timmy Reyes, professional surfer, 3 November 2015 (Surfer Magazine (2006))

3.5.1 Identification

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

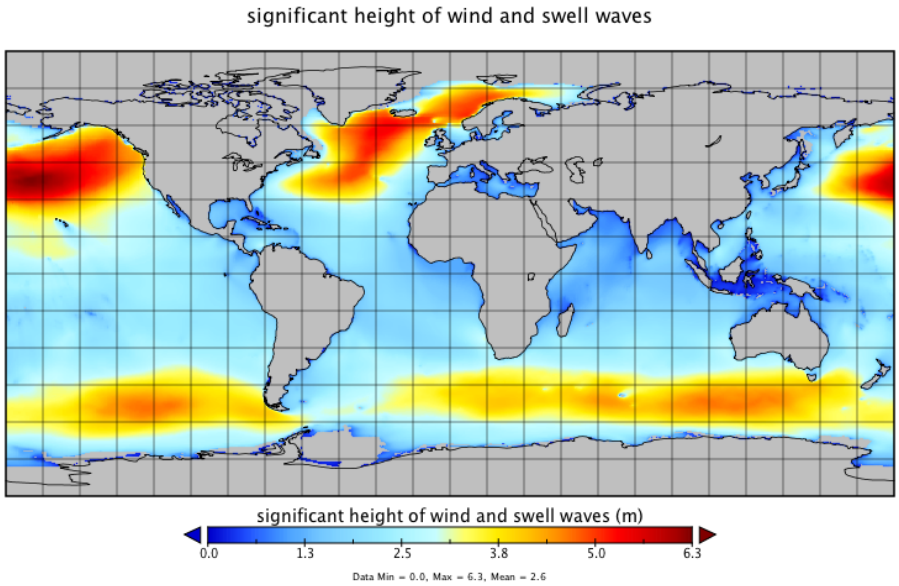
The El Niño Southern Oscillation (ENSO) fluctuates between El Niño and La Niña states every three to seven years (Butt (2009)).[18] [19] In a typical year trade winds blow across the Central Pacific from east to west, pushing warm surface water towards Australasia. Here warm air rises, creating low air pressure and precipitation. The surface water is replaced with cool upwellings near Chile and Peru, which creates high pressure and dry conditions there. During a La Niña episode these trade winds are stronger than average, reinforcing the effects. During El Niño events the trade winds weaken, if not reverse. Warm surface water stays in the east, causing low-pressure systems in the far North Pacific to track much further south. Broadly, the result is larger waves in the North and Eastern Pacific (including California and Hawaii); and smaller waves in Australia and the North-West Atlantic (Butt (2009); Housman (2016)). The advantage of El Niño swells, particularly for surfers in the North-East Pacific, is that they are generated far away and are often coupled with clement local conditions (Butt & Russell (2004); Butt (2009); Housman (2016)).

[18] Exogenous ENSO changes have been used as natural experiments in a range of economic studies on a variety of countries, sectors and commodities (e.g. Handler & Handler (1983); Brunner (2002); Hsiang et al. (2011); Ubilava (2012); Iizumi et al. (2014); Hsiang & Meng (2015)).

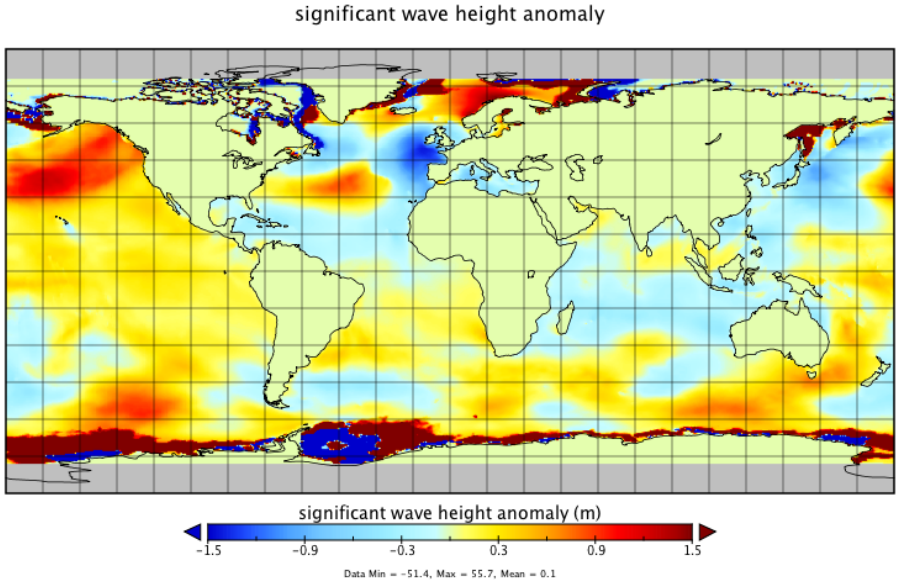
[19] Cashin et al. (2017) find El Niño events cause economic activity to fall in Australia, Chile, Indonesia, India, Japan, New Zealand and South Africa, while the US and Europe benefit.

Figure 3.10: Wave heights

(a) Significant wave heights



(b) Significant wave height anomalies



measurements in (m) during the January 1992 El Niño event.

We use a triple-difference approach to estimate the effect of large waves, during El Niño years, at high-quality breaks. The interaction of all three variables is needed. Large waves alone are not sufficient for good surfing conditions, as they will be unruly if generated nearby and accompanied by strong winds and storms. El Niño events produce long-range

swells and good weather in some parts of the world, but reduce wave heights in others (see figure 3.10b). They also affect the economy through other channels, like weather and climate (Cashin et al. (2017)). Finally, neither large waves nor good weather will matter if the break itself is low quality and cannot accommodate large swells.

3.5.2 Estimating equations

We conduct this experiment using a triple-difference specification:

$$\begin{aligned}\Delta \ln(Y_{i,t}^d) =& \alpha + \beta wha_{i,t} + \sum_{j=2}^{5} \beta_j wha_{i,t} Q_i \\ &+ \gamma ENyr_t + \sum_{j=2}^{5} \gamma_j ENyr_t Q_i \\ &+ \delta wha_{i,t} ENyr_t + \sum_{j=2}^{5} \delta_j wha_{i,t} ENyr_t Q_i \\ &+ W_i + Z_t + \epsilon_{i,t}\end{aligned} \tag{3.4}$$

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

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

3.5.3 Results and discussion

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

First we run the regression in equation 3.4 excluding any terms involving $ENyr_t$, to find the average effect of wave heights on light growth (columns (1) and (2) of table 3.9). We find that, on average, unusually large waves had no effect on light growth for most break qualities, and a negative effect for "Normal" (2 star) breaks. This result may be attributed to a 'stormy seas' effect: years with larger than average waves may also experience worse than average weather. As noted above, for large waves to be attractive to surfers, they must be accompanied by calm local conditions.

To isolate the marginal effect of waves generated by El Niño events we run the full specification in equation 3.4 at radii of 5km and 50km around each break (see columns (3) and (4) of table 3.9). We find that light growth was considerably higher during El Niño years for all locations. This is consistent with better weather and improved agricultural yields in North America and Europe, as found by Cashin et al. (2017). However, satellite sensitivity also significantly improved in 1994, 1997 and 2010 when new satellites were launched, which also happen to be El Niño years. Therefore these improvements in sensitivity may incorrectly have been attributed to $ENyr_t$ rather than year fixed effects, overestimating the true effect. However, this is not our focus.

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

These results are robust to excluding the US, which is a major beneficiary of El Niño

Table 3.9: Growth regressions, El Niño years

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

Robust p-values in parentheses

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

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

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

3.6 Conclusion

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

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

Collectively these results show that natural amenities play an important role in eco-

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

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

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

The chapter also suggests further avenues for research. One is to study the local impacts of natural amenities at a firm level, to disentangle the mechanisms through which they affect growth. This would distinguish between increased tourist activity, permanent em-

ployees accepting lower wages and entrepreneurs accepting lower profits. Another extension would be to estimate the economic contribution of other natural amenities in a spatially disaggregated model, with obvious candidates including scuba-diving, rock-climbing, and UNESCO natural heritage sites. Future work may also directly study the feedback effects of economic growth onto the quality of natural amenities, through pollution and overcrowding. Finally, amenities like surf breaks may be a useful as instruments when studying the impact of economic growth on other variables at a local level.

Chapter 4

Pricing sovereign debt in resource rich economies

4.1 Introduction

After a series of high profile defaults in the late-90s, including by Russia (1998), Ecuador (1999) and Argentina (2001), sovereign default issues moved to the forefront of economic policy discussions. Interest was again renewed by the, still ongoing, debt crisis in Greece and other peripheral eurozone countries following the financial crisis of 2007/08. The collapse in commodity prices in 2014/15 has sparked a new wave of debt distress across the resource-dependent developing world[1] leading to the de-facto default by Mozambique earlier this year. Given the continued slump in commodity prices and the sharp rise in sovereign debt levels in many developing economies, the question of how countries can best manage volatile oil revenues is as pertinent as ever.

As discussed in chapter 1, the conventional wisdom for countries managing oil and gas windfalls relies on the permanent income hypothesis (PIH) of Friedman (1957). Governments should smooth consumption by borrowing before a windfall and building up foreign assets in a sovereign wealth fund, or repaying existing debt, during the windfall. Whilst this

[1] This includes ongoing debt crises in Nigeria, Gabon and Azerbaijan, and worryingly high debt levels in Angola, Ghana, Kenya and South Africa

is often good advice for advanced economies, it can be dangerous for developing countries. First, the inherently volatile nature of commodity prices has led some to propose the use of *volatility* funds, which would be larger the more permanent the windfall is (Bems & de Carvalho Filho (2011)). Second, if a country faces high borrowing costs or restricted access to international finance, investment will be too low and oil revenues may be better used to invest in the domestic economy. van der Ploeg & Venables (2012) show that an *investment* fund, which prioritises domestic asset accumulation, may be useful in this respect. In addition, given that developing countries often face significant absorptive constraints, this fund could be used to manage the timing of these investments.[2]

This chapter investigates a specific source of volatility in resource rich economies - the link between the oil price and sovereign risk premiums. Empirically, there is a strong negative relationship between oil prices and interest rates in resource rich developing economies. I develop a stochastic general equilibrium model of a small open economy that receives a stream of income from production and a separate stream of resource revenues. The government trades in sovereign bonds with risk-neutral, competitive investors and chooses the level of consumption and debt position, as well as whether or not to default on its existing debt, where defaulting carries some cost. International investors form expectations on the likelihood of default in any given period and charge a risk premium on sovereign borrowing to account for this. This endogenous risk premium, which is affected by the stock of debt, as well as the stochastic processes in the model, then feeds back into the government's optimal debt position. The result is a negative correlation between oil prices and interest rates. The model is able to explain a large proportion of the business cycle fluctuations in interest-rate spreads in resource dependent, developing economies, particularly the strong countercyclical movements of interest rate spreads with oil prices. I also show that a standard model without oil revenues is unable to explain a significant proportion of the variability in the data.

[2]See Collier et al. (2010); van der Ploeg & Venables (2012); Wills (2015) for a review of the theoretical underpinnings of these different 'types' of funds

Using the model, I attempt to answer two important questions. How do commodity price movements affect the pricing of sovereign debt in resource rich economies? And, what does this mean for the optimal management of volatile oil revenues in developing countries? Falling oil prices present an important source of default risk for resource dependent sovereigns. Managing this risk is important for developing countries, particularly those that are capital scarce. The model delivers an endogenous link between oil prices and country risk premiums and is, to the best of my knowledge, new to the literature[3]. This mechanism provides an additional rationale for a *volatility* fund in which liquidity buffers are accumulated to manage debt repayments during negative oil price shocks.

This work is closely related to a large literature on sudden capital flow reversals in emerging markets, which emerged following the Latin American and East Asian financial crises of the 1990s. Original work by Calvo (1998) investigates the mechanisms of sudden capital reversals in developing economies, arguing that concerns of bankruptcy and financial disruption can often become self-fulfilling with large deleterious effects on the domestic economy. More recent work has highlighted the importance of debt-inflation dynamics and endogenous credit constraints in explaining sudden reversals in emerging markets. Calvo & Mendoza (2000) show that when the collection of information on developing countries is costly and imperfect, globalisation can reduce the expected-utility gain of paying information costs, thus leading to actions by a small number of informed investors causing large swings in portfolio rebalancing by uninformed investors. Mendoza (2006) develops a quantitative analysis of this mechanism in a dynamic, stochastic general equilibrium (DSGE) framework. He shows that these episodes, which are sparked by shocks to fundamentals, can result in credit constraints that bind when countries are highly indebted, causing significant output declines. Mendoza & Yue (2012) extend this framework further to include working capital constraints. Here, default results in capital market exclusions and triggers

[3]Work by van der Ploeg & Venables (2011) shows the importance that an exogenous risk premium on international borrowing has on the saving and investment decisions for developing economies. Whilst some recent work investigates the empirical nature of this question (Arezki & Brückner (2012)), to the best of my knowledge, this chapter is the first to address this in a comprehensive macroeconomic context. Work by van der Ploeg & Venables (2011) shows the importance that an exogenous risk premium on international borrowing has on the saving and investment decisions for developing economies.

an efficiency loss as these inputs are replaced by imperfect local substitutes.

There is also a growing literature on the empirics of default premiums in emerging economies. Volatile interest rate spreads have been found to be an important driver of output volatility in developing economies. Neumeyer & Perri (2005) and Uribe & Yue (2006) provide evidence that interest rate spreads can explain a sizeable proportion of output volatility and vice versa. Shocks to trend output are also empirically important in emerging markets. Aguiar & Gopinath (2006) develop a model in which the likelihood of default is affected more after a trend shock than compared to a transitory one.

The seminal works of Eaton & Gersovitz (1981) and Grossman & Van Huyck (1988) were the first to study the link between default risks and optimal borrowing behaviour in a dynamic macroeconomic framework. Eaton & Gersovitz (1981) develop a model of sovereign debt in which the possibility of permanent market exclusion following default is sufficient to generate an endogenous debt limit; below which it is always optimal to service outstanding debt but above which default is optimal for the sovereign. Grossman & Van Huyck (1988) show that even non-permanent market exclusion can present a sufficient risk to the sovereign, yielding a similar result (see also Kletzer & Wright (2000)). Finally, Park (2017) shows that in a model with capital accumulation default can also occur in 'good' times[4].

Following this work, the literature on sovereign default has expanded to tackle a number of different questions. Arellano (2008) develops a small open economy model of default under limited commitment. Her model delivers a strongly countercyclical process for interest rate spreads and output. Default is more likely in recessions because this is when it is more costly for a risk averse borrower to repay non-contingent debt. Adam & Grill (2017) extend the existing Ramsey policy literature to a setting with non-contingent sovereign debt and continuous default costs. They find that, when government bond markets are incomplete, and under certain parameterisations of default costs, partial default is Ramsey optimal

[4]The model delivers a U-shape in the capital stock: at both low and high levels of capital, the economy has an incentive to default on its debt. Default in good times occurs when the economy has over-invested in capital during booms

when a country's wealth is sufficiently low.[5]

There also exists a large body of research that explores the link between commodity prices and macroeconomic performance. Hamilton (2003) and Barsky & Kilian (2004) find that whilst there does exist a strong negative link between oil prices and growth in the US, the relationship is complex, often subject to reverse causality and almost certainly non-linear. For commodity exporters, higher commodity prices have positive short-term effects on output and growth but adverse long-term effects (see Deaton & Miller (1995); Raddatz (2007); McGregor (2017)). This is particularly true in countries with poor governance (Collier & Goderis (2012) Mehlum et al. (2006b)). An extensive literature on the 'resource curse' (Sachs & Warner (1999)) in developing economies also exists, in which resource booms lead to worsening economic and development performance over the long run (see van der Ploeg (2011) for a review of the literature).

The remainder of this chapter is organised as follows. Section 4.2 describes the evolution of interest rate spreads in resource rich, developing economies and presents some evidence for a counter-cyclical relationship between spreads and oil prices. Section 4.3 sets out the stochastic debt model. Section 4.4 solves the model under a baseline calibration and discusses some of its key properties. Section 4.5 derives an analytic expression for the price of sovereign debt and estimates a numerical approximation to this under the baseline model. Section 4.6 assess the model's performance along some key dimensions. Section 4.7 concludes.

4.2 Emerging market spreads

There is growing evidence that emerging market output volatility may be related to the ability for emerging markets to access international financial markets in times of need. In particular, a large proportion of the volatility in output is due to highly countercyclical interest rates and countercyclical default risk. Work by Neumeyer & Perri (2005) presents

[5]Others include the inclusion of long-term debt (Chatterjee & Eyigungor (2012)), CRRA lenders Lizaro (2013), misspecified default probabilities (Pouzo & Presno (2016)) and bailout risk Fink & Scholl (2016).

evidence that in contrast to advanced economies, emerging market business cycles are more volatile and interest rates are countercyclical and lead the economic cycle. Uribe & Yue (2006) show that the majority of the movements in emerging market spreads are due directly to changes in the country specific spread, or the risk-premium. To date, the literature has remained relatively silent on these issues in a resource rich setting.

These relationships are presented in figure 4.1. I plot quarterly real GDP in logs and the interest-rate spread for four oil exporting, developing economies: Colombia, Indonesia, Mexico and Russia.[6] The data are seasonally-adjusted and de-trended.[7] They cover the period 1995q1-2016q3. Interest rate spreads in these four economies have been volatile, as has GDP growth. The various economic downturns are visible for each country, with an upward spike in the interest rate spread coinciding with, and often leading, the economic slump. All experienced sharp downturns between 2009 and 2010 in the aftermath of the financial crisis, apart from Indonesia which recovered relatively quickly.

Critically, these countries are highly resource-dependent and rely heavily on oil exports as a source of foreign exchange. An empirical regularity in these economies is a negative correlation between oil prices and interest rate spreads. Spreads remain low during the 'good' times when oil prices are high, and rise in the 'bad'. This negative relationship can be seen clearly in figure 4.2, which plots the real world oil price against the interest rate spreads for these four economies. The price of crude oil has been highly volatile, with a pronounced boom and bust cycle during the 2007/08 financial crisis and a continued downward spiral since the end of 2014. It has fluctuated between a low of around $US 25 in 2001q4 and a peak of $US 124 in 2008q2. Spreads have also been volatile and negatively correlated with the oil price. This negative relationship is confirmed in the data where the correlation coefficient between the mean interest rate spread and the oil price is $-.14$[8].

[6]These four countries comprise the full set of resource rich, developing economies who are net exporters of oil and for which we have quarterly GDP and interest-rate data.

[7]I decompose the real GDP and spread series into a trend and cyclical component, using a simple HP filter with smoothing parameter of 1600. The interest rate spread is the difference between the yield on 10-year government bonds in each emerging economy and the US

[8]The individual country correlations are: Colombia = $-.33$, Indonesia = $-.03$, Mexico = $-.09$ and Russia = $-.12$

Clearly, oil price volatility is an important contributor to the pricing of risk and debt in these economies.

Figure 4.1: Emerging market economic performance

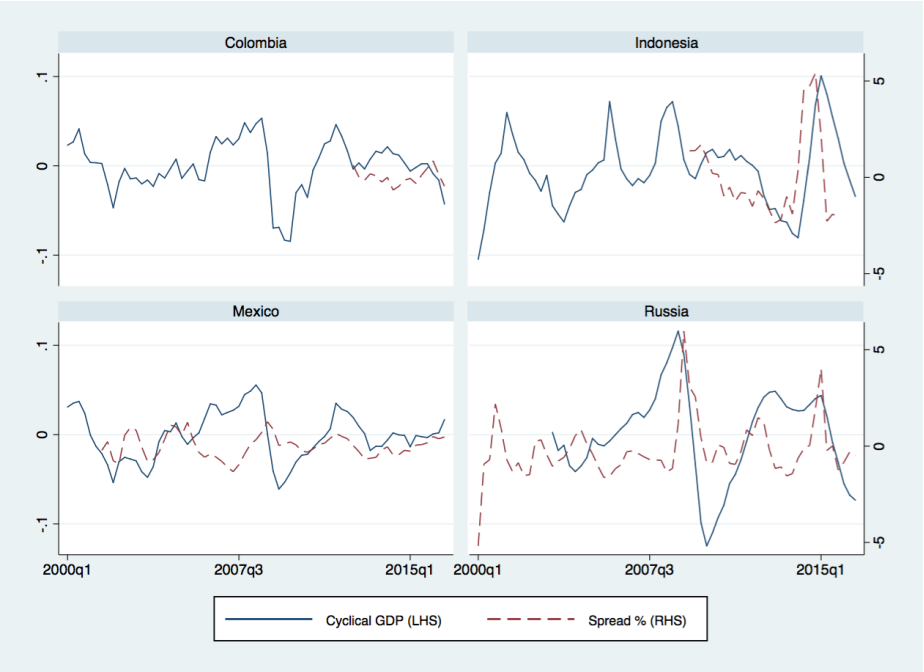
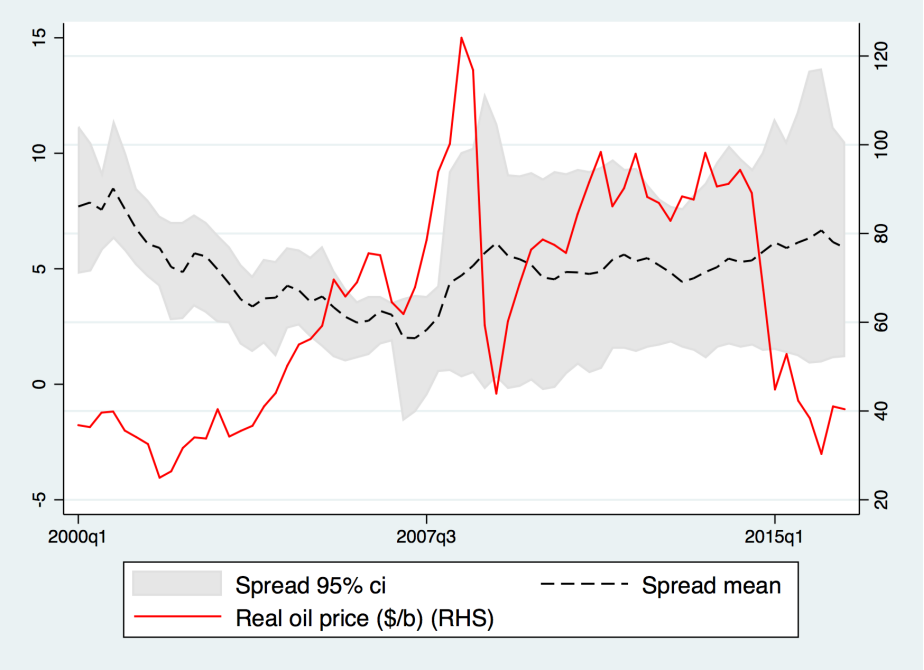


Figure 4.2: World oil price



4.3 The model

The model here is based on original work by Eaton & Gersovitz (1981) on international lending which has been extended by Arellano (2008) to study the default experience of Argentina. Consider a small open economy that receives two stochastic streams of income; one from output, y, and one from tax revenues from the oil sector, which are driven by a volatile oil price, p, as follows:

$$\begin{aligned} y_t &= \mu_y + z_t \\ z_t &= \rho_z z_{t-1} + \epsilon_t^z \\ \epsilon_t^z &\sim N(0, \sigma_z^2) \end{aligned} \tag{4.1}$$

$$\begin{aligned} p_t &= \mu_p + h_t \\ h_t &= \rho_h h_{t-1} + \epsilon_t^h \\ \epsilon_t^s &\sim N(0, \sigma_h^2) \end{aligned} \tag{4.2}$$

where μ_y and μ_p are the means of output and the oil price respectively, and ϵ_t^z and ϵ_t^h are the i.i.d shocks to each series. The representative agent seeks to maximise CES utility as follows:

$$U_t = \left[c_t^{1-\sigma} + \beta \left(\mathbb{E}_t U_{t+1}^{1-\sigma} \right) \right]^{\frac{1}{1-\sigma}} \tag{4.3}$$

where, as is standard in this class of utility function, the coefficient of relative risk aversion is given by $\sigma \geq 0$, $1/\sigma$ is the inter-temporal elasticity of substitution (IES), and $\beta = \frac{1}{1+\rho}$ is the discount factor. I follow the exposition of Caldara et al. (2012) who nest a simple CES utility in a general exposition of recursive preferences in a DSGE framework.

In every period the government chooses to either default on existing debt, D, or to not default and continue servicing the debt stock, N. This is different to the proportional

default approach taken in some of the literature. In the default state consumption is simply equal to income plus oil revenues under default, $c^D = y^D + o^D$, where y^D and o^D are as follows:

$$y_t^D = \begin{cases} y_t, & \text{if } y_t < \overline{y} \\ \hat{y} = \psi \overline{y}, & \text{if } y_t > \overline{y} \end{cases} \tag{4.4}$$

$$o_t^D = \begin{cases} p_t \bar{s}, & \text{if } p_t < \overline{p} \\ \hat{o} = \psi \left(\overline{p} \bar{s} \right), & \text{if } p_t > \overline{p} \end{cases} \tag{4.5}$$

where $\psi \in (0, 1)$ represents the (common) output and oil revenues cost of default (higher ψ means lower cost of default), $\overline{y} = \mathbb{E}\{y\}$ is expected output and $\overline{p} = \mathbb{E}\{p\}$ is the expected oil price, $\bar{s}$ is the annual flow of extracted oil, assumed to be fixed. The default cost is assumed to be a fixed proportion of income for sufficiently high levels of income.

If the government defaults it is then assumed to be in autarky (i.e. the country is excluded from capital markets) with an exogenous probability of re-entry in the future. If either output or the oil price is above their respective means, then actual output and oil revenues under default is truncated down by a fixed proportion of the expected variables, as governed by the cost parameter, ψ. The result is an asymmetric cost of default schedule. The output and oil revenue processes are truncated because it is assumed that default entails some direct cost. For output these costs may arise due to the inability of firms to access credit for investment or imports, reduced foreign investment, costly legal processes, political upheaval and civil protest, etc. For oil revenues, defaulting may simply result in a punitive default procedure in which the taxation of oil revenue, which is an immediate source of foreign exchange earnings for the government, are used to pay preferred creditors immediately.[9] The main motivation for doing this is to bring the default probability implied by the model in line with the data (see Arellano (2008)).

[9]See figure D.4 in the appendix for graphical exposition of this output truncation process simulated over 1000 periods. As we can see, the output realisation under default has a ceiling imposed; high output realisations are truncated. A similar process applies to oil revenues under default.

In the no-default state consumption, c^N is equal to stochastic income plus oil plus new borrowing, net of interest payment on existing borrowing as follows:

$$c_t^N = y_t + p_t\bar{s} - q_t(d_{t+1}, y_t, p_t)d_{t+1} + d_t \tag{4.6}$$

where, d_{t+1} represents one-period bonds sold at price $q(\cdot)$, and d_t is the repayment due to previous borrowing (for indebted countries, $d < 0$).

Creditors are assumed to be risk-neutral. They lend to sovereigns in the current period by buying bonds, d_{t+1}, at price $q_t(\cdot)$. In the following period, the creditors receive the face value of the bond if the sovereign does not default, and nothing if the sovereign defaults. The probability of default is given by:

$$\delta(d_{t+1}, y_t, p_t) = Pr\big[V^D(y_t, p_t) > V^N(d_{t+1}, y_t, p_t)\big] \tag{4.7}$$

so the probability of the creditor getting paid (i.e. no default) is $1 - \delta(d_{t+1}, y_t, p_t)$.

The bond pricing schedule, $q_t(d_{t+1}, y_t, p_t)$, is given by:

$$q_t(d_{t+1}, y_t, p_t) = \frac{1 - \mathbb{E}_t\{\delta(d_{t+1}, y_t, p_t)\}}{1 + r_t} \tag{4.8}$$

where r_t is the risk-free interest rate and $\mathbb{E}$ is the expectations operator. The bond price depends negatively on the default probability, $\delta(d_{t+1}, y_t, p_t)$, and the risk-free interest rate, r_t. A higher default probability reduced the price investors are willing to pay to hold the bond. A higher output realisation, y_t or lower debt position in the following period, d_{t+1}, feeds through to lower probability of default and so a higher bond price.

Every period the sovereign's default decision can be represented by:

$$V(d_t, y_t, p_t) = \max_{D,N}\{V^D(y_t, p_t), V^N(d_t, y_t, p_t)\} \tag{4.9}$$

where output, y_t follows the stochastic process defined above, and d_t denotes the asset stock of the sovereign ($d_t < 0$ signifies indebtedness). Default is chosen by the sovereign if $V^D > V^N$.

The present value of being in the default state, V^D, is given by:

$$V^D(y_t, p_t) = \left[u(y_t^D + o^D)^{1-\sigma} + \beta \mathbb{E}_t V^D(y_{t+1}, p_{t+1})\right]^{\frac{1}{1-\sigma}} \tag{4.10}$$

whilst the present value of being in the non-default state, V^N, is given by:

$$V^N(d_t, y_t, p_t) = \max_{d_{t+1}} \left[u(c_t^N)^{1-\sigma} + \beta \mathbb{E}_t V^N(d_{t+1}, y_{t+1}, p_{t+1})\right]^{\frac{1}{1-\sigma}} \tag{4.11}$$

where $\beta = \frac{1}{1+\rho}$ is the discount factor and $\sigma \geq 0$ is the inverse of the IES. The expression for the expected future value will differ depending on whether the government chooses to default or not. These are given by:

$$\mathbb{E}_t V^D(y_{t+1}, p_{t+1}) = \sum_{y=1}^{yn} \sum_{p=1}^{pn} \left[\theta V(0, y_{t+1}, p_{t+1}) + (1-\theta) V^D(y_{t+1}, p_{1+1})\right] \Pi\left(y_{t+1}, p_{t+1} | y_t, p_t\right) \tag{4.12}$$

and

$$\mathbb{E}_t V^N(d_{t+1}, y_{t+1}, p_{t+1}) = \sum_{y=1}^{yn} \sum_{p=1}^{pn} \left[V(d_{t+1}, y_{t+1}, p_{t+1})\right] \Pi\left(y_{t+1}, p_{t+1} | y_t, p_t\right) \tag{4.13}$$

where θ is the exogenously determined probability of re-entry into credit markets after default. Finally, $\Pi(\cdot)$ is the transition probability matrix associated with each stochastic process and is given by:

$$\Pi\left(y_{t+1}, p_{t+1} | y_t, p_t\right) = \Pr\left(y_{t+1}, p_{t+1} | y_t, p_t\right) \tag{4.14}$$

Under the assumption of independence between the shocks series to output and the

oil price I define the transition probability of each series following Tauchen (1986). When the shocks are allowed to have a non-zero covariance, the transition matrix is different as it needs to take these cross-correlations into account. In this case the derivation of the transition probability is described more fully in section D.1.2.

4.4 Quantitative analysis

4.4.1 Solving the Model

The above equations are sufficient to characterise a simple default model. Due to the non-linear nature of the model, it is necessary to use numerical methods to solve for the optimal policy functions. Here I use a value-function iteration approach. This is implemented by discretising the state-space for debt, output and oil revenues[10], using a Markov chain for the stochastic output and oil price processes, and solving for the policy functions using a guess and verify procedure.

The recursive solution to the model consists of a set of policy functions for the representative agents' choice of consumption, $c(d_t, y_t, p_t)$, as well as default decision, (D, N), and optimal asset holdings in the following period, d_{t+1}; and a bond pricing function, $q(d_{t+1}, y_t, p_t)$, which reflects the sovereign's default probabilities and is consistent with the creditors' expected zero profit condition. For a given choice of parameters, the numerical solution method involves: taking an initial guess for the bond price schedule and the value functions, under default and no-default, and solving for the optimal policy functions and probability of default. The bond pricing schedule is then updated using the probability of default and the preceding steps repeated until the problem converges.[11]

A simplified RBC version of the model in section 4.3 is described in section D.1.1 of the technical appendix D.1 in which I derive expressions for a number of analytical properties of the model. I show in proof 1 that the Euler equation of the RBC debt model collapses to

[10] I use 11 output states, 11 oil price states and 125 debt states

[11] See section D.2.1 in the appendix for a description of the algorithm.

that of a standard RBC model when default is not an option for the sovereign. Further, I show in proof 2 that there is a single parameterisation of the fixed default cost and risk-free interest rate such that the optimal debt choice collapses to that of a simple RBC model with perfect capital markets.

The optimal choice of sovereign debt depends crucially on whether we allow for default. Without default the level of debt is decreasing in the elasticity of the sovereign's interest rate with respect to the level of indebtedness. That is, the more the interest rate charged on borrowing, or conversely the bond price, responds to the level of outstanding debt, the less the sovereign chooses to borrow. This is described in lemma 1 in the technical appendix D.1. When default is possible, the sovereign may choose to borrow over and above that required for pure consumption smoothing. In particular, consumption tilting may occur when the sovereign is impatient relative to the risk-free interest rate, or when the debt-elasticity of the bond price is low. This relationship is described in lemma 2 in the technical appendix D.1.

4.4.2 Base calibration

The model includes the following nine base parameters: $\rho, \beta, \sigma, \theta, \psi, r, r^*, \hat{r}, \bar{s}$. Table 4.1 lists the choice of parameters along with their baseline calibration.

Table 4.1: Model Parameters - base

Parameter	Value	Description
ρ	.020	Discount rate (set for consumption smoothing)
$\beta = \frac{1}{1+\rho}$	.980	Discount factor
σ	2.0	Coefficient of relative risk aversion
θ	.25	Probability of re-entry into capital market
ψ	.95	1 - output cost of default
$\bar{s}$	.060	Annual production of oil (oil share in GDP)
r^*	.004	risk-free interest rate
$\hat{r}$	.016	Risk premium (exog)
r	$r^* + \hat{r} = .020$	Default-risk-free rate

The risk-free interest rate is set to $r^* = .004$ per quarter. When compounded over four quarters, this is equivalent to the 1.6% mean annual real yield on 10-year US government bonds between 1995q1 and 2016q3. I include an exogenous risk premium for emerging market 10-year bonds of $\hat{r} = .016$, which is equivalent to the 6.5% per year mean annual interest rate spread of the four economies in our sample over the US. The total interest rate facing emerging markets, without default risk, is therefore set to $r = r^* + \bar{r} = 0.020$ per quarter, or 8.2% per year.

The rate of time preference is calibrated to yield consumption smoothing behaviour in our model, that is $\rho = r = .020$. By consumption smoothing here, we mean that the model delivers a flat consumption profile in the steady state. A rate of time preference parameter that is greater (smaller) than the prevailing risk-free rate of interest would create an incentive for households to consume less (more) today and save more (less) for tomorrow thus tilting consumption to the future (present). This assumption ensures the model has a well-defined steady state. The coefficient of relative risk aversion is set to $\sigma = 2$ as is standard in the RBC literature. The asymmetric default cost parameter is set to $\psi = .95$.

The probability of re-entry following default is set to $\theta = .25$ every quarter. Using a monte carlo simulation approach, I find that, over 1000 runs, the average number of years that a country would be excluded from financial markets using this re-entry probability, is just over 2 years. This is somewhat lower than some of the evidence suggests. For example, Richmond & Dias (2009) find that countries are typically excluded for around 5.7 years following a sovereign default episode.

How dependent the domestic economy is on oil revenues is a key parameter in the model. Given the model set-up, oil revenues are important as they provide vital foreign exchange earnings which can then be used to pay down foreign denominated debt. One way to calibrate this parameter would be to estimate the share of oil production in GDP. Using data from the World Bank's WDI database, I find the average share of oil and gas rents in total GDP across our sample of countries over the period 1995q1 to 2016q3 to be

6%. Another way to calibrate this parameter would be to estimate the share of total export earnings which are due to oil exports. Using UN Comtrade data, I find the average share of petroleum in total exports to be 35%. Table 4.2 gives these share estimates for each country. What is clear from the data is the large variation in oil shares (GDP or exports) across the four countries.[12]

I calibrate the oil share parameter using the GDP share of $\bar{s} = .060$ but present simulation results for calibrating this parameter using the oil export share in the appendix (see section D.2.2). In general a higher oil share increases the spread persistence as well as the strength of the negative correlation between interest rate spreads and the oil price.

Table 4.2: Oil and gas shares

country	GDP share	Export share
Colombia	.039	.445
Indonesia	.037	.265
Mexico	.036	.114
Russia	.128	.579
All	.060	.351

4.5 Pricing sovereign debt

The model described in section 4.3 and solved in section 4.4.1 presents a theoretical framework in which the ability of a resource rich economy to borrow is dependent on the cyclical properties of the commodity price. The result is that international capital markets, when setting the price of a sovereign's debt, internalise the effect that commodity price fluctuations might have on the likelihood of default.

We can be more specific about this relationship. In the simplified RBC model setup described in section D.1.1 in the technical appendix D.1 I show that with no opportunity

[12] It is worth noting that while the mean GDP share has remained relatively stable over time, the export share has increased steadily. See figure D.1 in the appendix

of default and a stochastic income process, the price of sovereign debt is simply inversely related to the interest rate and the amount of debt outstanding. When default is open to the sovereign, and a fixed default cost associated with this decision, the relationship between the debt price and the level of debt becomes non-linearly related to the level of debt as well as the default cost and the likelihood of the sovereign to default in any given period. Using this simple model, I derive an expression for the steady-state price of sovereign debt in expression D.24 in the technical appendix D.1, which I reproduce here:

$$q = \frac{1}{1+R} = \frac{1-\lambda}{R\left(\delta'(\lambda-d)-\delta+1\right)} + \frac{1}{d\delta' R} \tag{4.15}$$

The steady-state price of sovereign debt is a complex, non-linear function of the fixed cost of default, λ, the interest rate charged to the sovereign, R, the level of debt, d, the probability of the sovereign to default, δ, and the sensitivity of this probability to the level of debt, δ'.

When default is not possible for the sovereign we set $\delta = \delta' = 0$ and $\lambda = 0$, and the bond pricing expression collapses to $q = \frac{1}{R}$. This simple expression states that the steady state bond price of sovereign debt is inversely related to the interest rate facing the sovereign, R. When default is possible, this inverse relationship is augmented by the cost of default, λ, and the optimal default decision by the sovereign, δ. The more sensitive the default decision of the sovereign is with respect the level of debt (that is, the larger is δ'), the lower the price investors are willing to pay to hold this debt, and the higher the effective interest rate. This relationship is described in lemma 4 in the technical appendix D.1. Note that the interest rate, R, is itself a function of the level of debt as well as the output and oil price states as described in the expression D.5 in the technical appendix D.1.

Crucially, the bond price is inversely related to the level of debt and positively correlated with output and the oil price, but in a non-linear way. Given that we do not have explicit functional form representations of the interest rate function, $R(\cdot)$, and the default decision, $\delta(\cdot)$, the shape of this non linear bond pricing schedule can only be approximated using

the numerical solution algorithm. For the base calibration of our model, the bond pricing schedule is given in figure 4.3. As we can see from the figure, the bond price is a downward sloping, non-linear function of the stochastic output and oil price process.

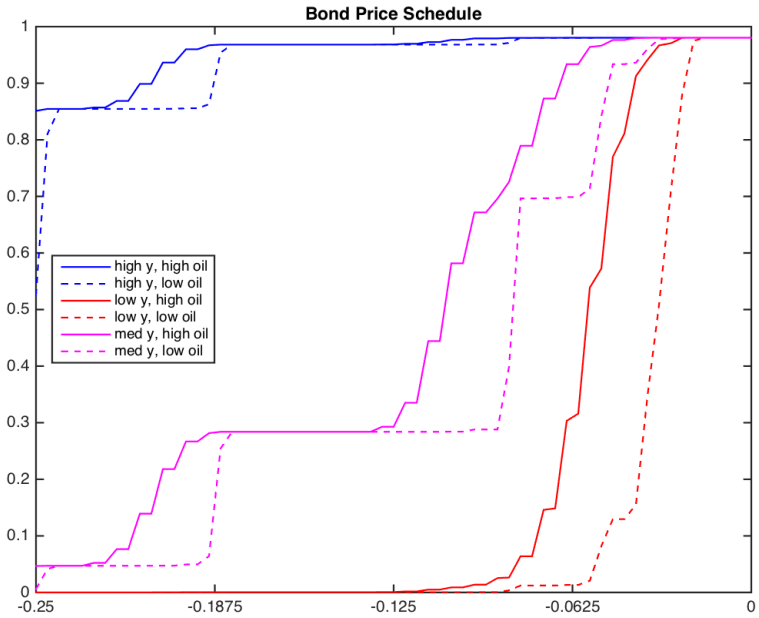
Figure 4.3 shows how the equilibrium bond price (vertical axis) varies with the level of indebtedness (horizontal axis) for different output and oil price states. The solid lines represent the schedule from Eq. 4.8 for 'high', 'medium' and a 'low' output states as a function of the current debt level given a 'high' oil price[13]. The bond price is an increasing function of foreign assets (negative assets signify a net debt position). The sovereign has a choice in any single period: (i) default on debt today, increasing consumption and so utility today, but being excluded from asset markets and facing a default cost, or (ii) do not default today, lowering consumption and utility today due to interest payments, but maintaining the ability to borrow and avoiding the negative consequences of default. When setting the bond price, international investors take the sovereign's default decision into account, decreasing the amount they would be willing to pay for this debt the more likely it is that the government will choose to default. The larger the outstanding debt (i.e. a more negative asset position), the more likely the sovereign is to default and so the lower is the bond price. The pricing schedule for the 'high' output state is above that of the 'low' output state simply because higher output realisations mean a lower chance of default. This in turn means investors are more willing to hold government debt and so are willing to pay a higher price for it.

The dashed lines represent the same bond pricing schedules but for the 'low' oil price state. When oil prices are low the pricing schedules shift down. This is due to the oil revenue entering additively into the sovereign's budget constraint: a lower oil price reduces revenues, increasing the likelihood of default and lowering the equilibrium bond price. We get a similarly shaped bond pricing schedule in a model without oil. That is, the bond price is still a downward sloping, non-linear function of the output states. However, without oil in the model the bond pricing schedule has significantly larger steps in it. See figure D.5 in

[13]I use the 3rd, 6th and 9th positions for the high, medium and low states respectively.

the graphical appendix.

Figure 4.3: Bond pricing schedule



4.6 Model performance

4.6.1 Independent shocks

The model includes two stochastic processes: innovations to output and the oil price. I calibrate the output series, y, using quarterly data from the four developing economies and the oil price series, p, using quarterly data on on the real crude oil price. I deflate the nominal GDP series for each country in local currency using the national GDP deflator, and the crude oil price using US CPI data. I then take the natural log of these real GDP series and the real oil price. All data are quarterly. Finally, I pass the data through a Hodrick-Prescott (HP) filter with a smoothing parameter of 1600.

Data on nominal GDP, national deflators and US CPI all come from the Global Financial Database. The data on oil prices is the quarterly average of the daily West Texas oil price in USD per barrel taken from the IMF's IFS database.

To retrieve the structural parameters required for the model, I estimate the following auto-regressive process for the cyclical components of output and the oil price:

$$x_t = \rho x_{t-1} + u_t \tag{4.16}$$

where x_t is the cyclical component of the series in question at time t and u_t is assumed to be an i.i.d error. The coefficient estimates for ρ are used to calibrate parameters ρ_z and ρ_h in equations 4.1 and 4.2 respectively, while the sample standard errors of the residuals, u_t, are used to calibrate the volatility parameters σ_z and σ_h.

Table 4.3 gives the calibration of the four parameters governing the stochastic processes of the model: $\sigma_z, \sigma_h, \rho_z, \rho_h$. For the output process, the parameters are the mean of the estimated individual country parameters. The quarterly output process has a persistence term of 0.841 and is subject to an i.i.d shocks process with standard deviation of 0.015. Oil prices have a persistence term of 0.725 and shock standard deviation of 0.136.[14]

Table 4.3: Model parameters - stochastics

Parameter	Value	Description
ρ_z	.841	Persistence of output
ρ_h	.725	Persistence of oil price
σ_z	.015	Std. dev of output shocks
σ_h	.136	Std. dev of oil price shock

Using the above calibration I solve the model for the sovereign's policy functions using the value function iteration algorithm. I then simulate the model over 1000 periods and assess the ability of the model to reproduce the moments observed in our data. Tables 4.4 and 4.5 compare the moments of the simulated data to the data moments for our four resource rich economies between 1995q1 and 2016q3.

[14]It is worth noting that the shock processes to these series are not very persistent. In fact the half-life of these shocks can be calculated using: $h = T\log(2)/\log\left(\frac{x_t}{x_T}\right)$. This gives a half-life of 4.1 quarters for innovations to output and 2.2 for innovations to the oil price.

The first two columns in table 4.4 present the mean AR(1) coefficients and standard deviations for output, the oil price, consumption and the interest rate spread from our panel dataset. Consumption is the most persistent series, followed by output, then the oil price and finally the interest rate spread. The oil price is the most variable. The next two columns present these same data moments from the simulated model. The first thing to note is that the model successfully reproduces the correct signs on all of the moments. In addition, the ordering of the persistence and standard deviation terms between the four series are also correctly predicted. These are promising results.

Table 4.4: AR(1) model: autocorrelations and standard dev

	Data		Model	
	AR(1)	std	AR(1)	std
output (HP filt)	.841	.032	.717	.025
oil price (HP filt)	.725	.198	.650	.195
cons (HP filt)	.898	.026	.834	.018
spread	.631	.020	.188	.003

Table 4.5: AR(1) model: cross correlations with real interest rate spread

	Data	Model
output dev	-.138	-.106
oil price dev	-.279	-.094

The model correctly reproduces the autocorrelations and standard deviations of output and the oil price. This is to be expected as the relevant model parameters were chosen to match these two processes. It also does a good job on consumption, which is highly persistent and has a low variance. However the model substantially under predicts the persistence and standard deviation in the interest rate spread.

The first column in table 4.5 presents the average correlation between the real interest

rate spread and the deviations of output and the oil price in the data respectively. Both output and the oil price are negatively correlated with the spread in these economies. That is, periods of above trend output and oil prices are associated with lower interest rate spreads. A key feature of spreads in resource rich, emerging economies, is strongly countercyclical movements with oil prices. The second column presents the same moments from the model. The first thing to note is that the model produces the correct signs on these correlations, which is reassuring. However, whilst the model gets quite close to the output correlation, it fails to generate a sufficient oil price correlation.

4.6.2 Covarying shocks

Up to now I have assumed independent output and oil price shocks. It seems plausible that for resource rich developing economies, positive terms-of-trade shocks (that is, increases in the price of oil relative to imports) might lead to increases in domestic output. In the model of section 4.3, output, y, is defined as total non-oil output while the resource sector is defined as, $p\bar{s}$, where the oil price, p is subject to shocks. So positive co-movements in this setting would be between oil prices and non-oil output.

Deverajan et al. (1990) provide a parsimonious model for a single economy producing two goods (exports and domestic non-tradables) in a three good world (which includes imports). Their model is useful for thinking about the effects of a change in the export price (oil in this case) on the domestic economy[15]. In the model, a positive terms-of-trade shock has two effects: an *income* effect and a *substitution* effect. The fall in the relative price of imports means that domestic consumers face a positive income effect as their real incomes rise, thus increasing the demand for non-tradables. The substitution effect reduces demand for non-tradables as their price relative to imports has risen. The overall effect on domestic output, excluding exports, is ambiguous. If consumers cannot easily substitute between imports and non-tradables, a likely assumption in developing economies, then the

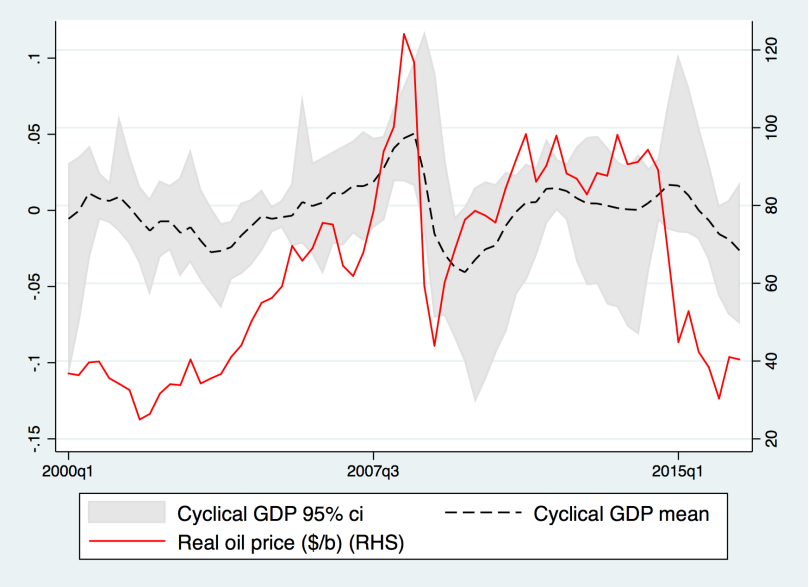
[15]Some structure may be useful here. On the demand side, consumers substitute (imperfectly) between imports, M, and domestic non-tradables, D, while on the production side the economy produces, D, and exports, X, which are imperfectly transformed between. A positive terms-of-trade shock implies a fall in the price of imports, P_M, relative to exports, P_X.

income effect will dominate and domestic output will increase. The result is that increases in domestic output are likely to coincide with positive oil price shocks.

So far, this simple analysis has focused only on the immediate inter-sectoral adjustment of the economy following a terms-of-trade shock, and has hinged on the low elasticity of substitution assumption. Other channels however are likely to operate, all of which would strengthen this positive correlation between oil prices and domestic output. These include: cheaper imported inputs, increased inward investment flows and larger foreign exchange earnings. These effects will be stronger the more important are exports relative to total domestic output, which is typically the case for resource dependent economies.

In resource rich economies this effect typically manifests itself as a positive co-movement of commodity prices and growth. This co-movement can be seen in figure 4.4 which plots the oil price against the cyclical component of GDP in our four countries. The mean correlation coefficient between real GDP and the oil price in our four economies is .20[16].

Figure 4.4: Covarying oil price and growth



[16]The individual country correlations are: Colombia = .41, Indonesia = −.06, Mexico = .14 and Russia = .29

I model these covarying output and oil price series in a simple 2-variable VAR model, over the period 1995q1 - 2016q3 for each country, as follows:

$$\mathbf{x_t} = \mu + \rho \mathbf{x_{t-1}} + \epsilon_\mathbf{t}$$
$$\begin{bmatrix} Y_t \\ P_t \end{bmatrix} = \begin{bmatrix} \mu^Y \\ \mu^P \end{bmatrix} + \begin{bmatrix} \rho^Y & \rho^{YP} \\ ... & \rho^P \end{bmatrix} \begin{bmatrix} Y_{t-1} \\ P_{t-1} \end{bmatrix} + \begin{bmatrix} \epsilon_t^Y \\ \epsilon_t^P \end{bmatrix} \tag{4.17}$$

$$\text{where } \Sigma_{YP} = \text{Var}(\epsilon^Y, \epsilon^P) = \begin{bmatrix} \sigma_Y^2 & ... \\ \sigma_{YP}^2 & \sigma_P^2 \end{bmatrix} \text{ and } [\epsilon^Y, \epsilon^Y]' \sim N(0, \Sigma_{YP}).$$

where Y_t and P_t are the cyclical components from the HP filter applied to real GDP of each country and the real oil price respectively. The parameter estimates for ρ^{YP} for each country are given in table 4.6. As we can see, there is a significant variation in these estimates.

Table 4.6: Cross AR(1) estimates - ρ^{YP}

	Coef.	Std. Err.
Colombia	.035**	.014
Indonesia	.077***	.019
Mexico	−.007	.008
Russia	.074***	.010
Mean	.045	n/a
Range	[−.007, .077]	n/a

Given this lack of tight priors about the strength of the correlation between output and oil price shocks, I explore the dynamics of the model for a reasonable range of values around the central estimate of $\rho_{zh} = .045$[17]. I choose upper and lower values for the covariance parameter, $\rho_{zh} \in [.01, .10]$, solve and simulate the model for 50 points in between these bounds. In each case, the model is then simulated over 1000 periods and the various data

[17] cross-correlation, ρ^{YP}, is estimated to be 0.045 in the data, the range is from −.007 to .077

moments retrieved. I focus on the estimates of the spread persistence and the its correlation with the oil price from the simulated data, as these are the key data moments we are trying to explain.

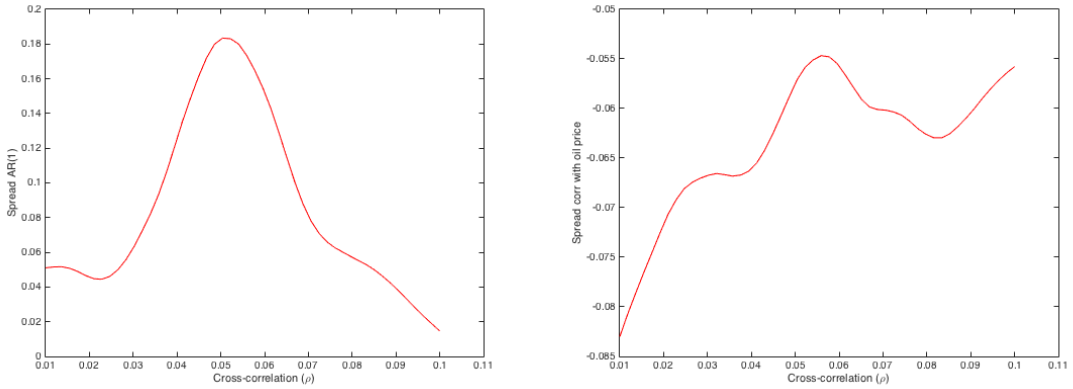
Figure 4.5a presents the AR(1) parameter for the interest rate spread across the range of parameters, while figure 4.5b presents the spread's correlation with the oil price[18]. The auto-regressive properties of the interest rate spread are reduced for very low and very high correlations of output and oil price shocks. For high and low correlations, the model displays a low spread auto-correlation, whilst for moderate correlations the spread is relatively persistent. The reason for this inverse-u shape is that there are two competing effects influencing interest rate spreads as the two shocks become more correlated. The first is an increased risk of default which drives up spreads. When the shocks are independent they act as partial insurance mechanisms against each other, and so default is less likely, ceteris paribus, and spreads remain low. As the shocks become more correlated, periods of low output combined with low oil prices become more frequent, thus increasing the likelihood of default and raising the risk premium. The second effect is the rising cost associated with default as default becomes more frequent. As the shocks become more correlated and the risk of default increases, the sovereign finds a lower level of indebtedness optimal at a given point in time. This in turn reduces the likelihood of default and so too the spread persistence.

The correlation of the spread with the oil price is broadly decreasing with the correlation of output and oil price shocks. The intuition here is relatively straightforward. The more correlated are the shocks, the more likely low oil prices coincide with low output, and so the more likely default becomes. As the default probability increases so too does the interest rate spread and so the higher the correlation of the interest rate spread with oil price shocks.[19]

[18] Given the discretised state-space and the numerical solution method, the data moments estimated from the simulations display significant estimation noise. I therefore display a smoothed version of these estimates in the figures.

[19] A similarly upward sloping relationship exists between the correlation of the shocks and the correlation of the spread with output shocks.

Figure 4.5: Sensitivity to correlation - $\rho_{zh} \in [.01, .10]$



(a) AR(1) parameter

(b) Correlations

4.6.3 Impatience and oil dependence

The baseline model assumed pure consumption smoothing. That is, the impatience parameter is set equal to the risk-free real interest rate, $\rho = r = .020$, or 8.2% per year. While this assumption may hold for developed economies, it is unlikely to be the case in developing economies where current generations discount the future more heavily due to stronger growth and greater uncertainty.

The permanent income theory of consumption yields the basic inter-temporal approach to the current account: in a small open economy, under perfect consumption smoothing, temporary income shocks will be offset by changes in national savings via the current account (Obstfeld & Rogoff (1995); Sachs (1982)). Consumption tilting, which arises due to differences between the subjective discount rate of the domestic agent and the prevailing world interest rate, yields behaviour in which a country shifts its consumption toward the present or the future, independently to the balance of prevailing shocks to income. In general, agents in developed economies are observed to be relatively patient. Braeu (2010) finds that households in Canada tilt consumption toward the future, while Cashin & McDermott (2002) find that the dynamics of international capital flows to Australia during the 1990s, when the country was a net capital importer, were broadly consistent with utility maximisation under consumption smoothing.

More recent work has examined the degree of impatience in consumption in developing country settings. Zhuang et al. (2007) and Harrison (2010) for example, find social discount rates that are in excess of prevailing market interest rates in poor countries. These high social discount rates are applied in practice by development organisations, such as the World Bank and the Asian Development Bank, who typically use annual discount rates in the range of 10-12% when evaluating projects in developing countries. Some developing country governments apply discount rates as high as 15% in their project appraisals.

Another consideration, is the degree to which governments attempt to smooth changes in fiscal policy. The basic Ricardian view, as described by Barro (1979), is that given the amount of government spending, the balance of government financing between taxes and debt should have no real effects. If taxes are distorting however, it will be optimal to smooth tax rates over time, thus implying that temporary shocks to government spending will be debt financed. The basic theory also implies that government deficits should move independently from the outstanding debt-income ratio as well as with the level of government spending. This inter-temporal tax-smoothing model appears to hold in developed economies such as the US.

A growing literature explores tax smoothing behaviour in emerging and developing country settings. These studies typically find evidence of both tax smoothing and tax tilting. Tax tilting is when a government chooses fiscal deficits that are either larger or smaller than those generated from the trade-off between tax smoothing and debt-sustainability objectives. Cashin et al. (2003) find that taxes in Pakistan remained fairly unresponsive to anticipated changes in expenditure, but that deficits were systematically larger over the 1970-90 period than would be expected from a series of optimal tax smoothing fiscal decisions. Tax tilting in developing countries may also be due to increased political risk. Pastén & Cover (2015) study the case of Latin America between 1984 and 2009, and find that a higher risk of losing power in the future increases the rate at which a government discounts the future resulting in higher deficits today.

The second issue is the degree to which countries depend on oil revenues as a source of foreign exchange. Our baseline model is calibrated to the mean hydrocarbon share, $\bar{s} = 0.060$. However, the level of dependence on oil of each of the four economies in our sample is highly variable. Mexico has the lowest dependence, with the share of GDP coming from the oil and gas sector standing at around 3.6% on average. In Russia, this figure stands at just under 13%.

This section tests the sensitivity of the model to the level of impatience and the dependence on oil. I vary the rate of time preference in the model in the range, $\rho \in [0.01, 0.05]$, per quarter, which works out at a discount rate roughly in the range of 4-20% of per annum. I also vary the oil share parameter in the range, $\bar{s} \in [0.02, 0.6]$.[20]

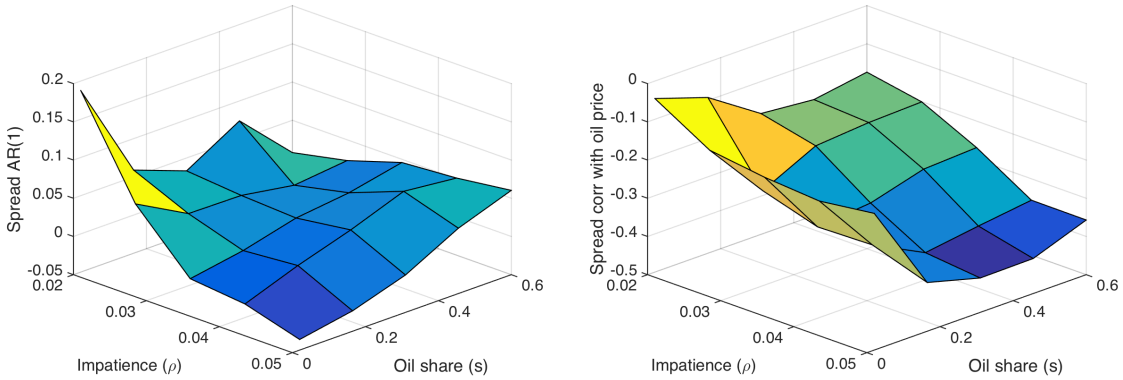
Figure 4.6a presents the AR(1) parameter for the interest rate spread for different combinations of impatience and oil dependence. The non-linear relationship between these deep parameters and the data moments is again clear. Broadly speaking, a higher level of impatience reduces the persistence of the interest rate spread whilst a higher oil share increases it. The more impatient the sovereign is, the more they tilt consumption to the present and push debt repayments to the future. This in turn increases the probability of default, ceteris paribus, thus increasing the interest rate spread. The more important oil revenues are to the sovereign, the more likely oil price shocks are to have an impact on the ability to meet debt repayments, and in turn on the likelihood of default.

The correlation between the spread and the oil price is given in figure 4.6b. There is a relatively strong positive relationship between increasing impatience and oil dependency and the strength of the negative correlation between interest rate spreads and movements in the oil price. Impatience and oil dependency have a similar effect here. They both act to bring consumption forward and push debt repayments back, thus increasing the default

[20] Given the computational weight of solving and simulating the model for different parameter sets jointly, I choose 5 evenly spaced points covering each parameter range.

probability when oil prices (and output) are low and raising risk premiums.

Figure 4.6: Sensitivity to impatience, $\rho \in [.01, .05]$, and oil dependence, $\bar{s} \in [0.02, 0.6]$



(a) AR(1) parameter

(b) Correlations

4.6.4 Recursive preferences

How agents make inter-temporal trade-offs and how they treat risk are likely to be important in any analysis of sovereign debt dynamics in a stochastic setting. A substantial literature exists documenting the low degree of inter-temporal substitution in developing countries, as well as the relatively high risk-aversion, although the evidence here is more mixed.

Households that face a subsistence consumption requirement will have an inter-temporal elasticity of substitution that is wealth dependent. As wealth increases, the inter-temporal substitution of elasticity converges to that under standard CES preferences. For poorer households, this implies a lower inter-temporal elasticity of substitution, and a savings rate that responds less to returns on capital (see Ravn et al. (2008)). At the aggregate level, in poor countries where budget shares of food are relatively high, the interest elasticity of saving is likely to be low, thus leading to a low inter-temporal elasticity (see Ogaki et al. (1996) for a discussion of these hypotheses). Atkeson & Ogaki (1996) use panel data in India to show that the inter-temporal elasticity amongst poor households is substantially lower that rich households.

On risk-aversion, the evidence here is informative but far from conclusive. A number of studies conclude that agents in developing countries are highly risk averse (Binswanger et al. (1980); Yesuf & Bluffstone (2009); Akay et al. (2012)). More recent work however, has argued that these earlier studies may be biased towards overestimating risk-aversion. Vieider et al. (2013) finds evidence of risk-neutrality amongst poor farmers in Vietnam. Gandelman & Hernandez-Murillo (2014) estimate a coefficient of relative risk aversion close to unity in developing countries, although risk aversion in African economies may be higher than in Asia.

The challenge here is that, up to now, the sovereign's behaviour has been modelled using deterministic preferences in which risk aversion and inter-temporal preferences are inversely related. It is not possible in the current set-up to alter the degree of risk aversion separately from the elasticity of inter-temporal substitution. One option is to use recursive preferences in which the risk-aversion and substitution parameters are independent.[21] To this end, I incorporate Epstein-Zin preferences into the model. Specifically, the value function for the sovereign under default in 4.10 becomes:

$$V^D(y_t, p_t) = \left[u(y_t^D + o^D)^{\frac{1-\sigma}{\eta}} + \beta\left(\mathbb{E}_t V^D(y_{t+1}, p_{t+1})^{\frac{1}{\eta}}\right)\right]^{\frac{\eta}{1-\sigma}} \tag{4.18}$$

whilst under no default the value function in 4.11 becomes:

$$V^N(d_t, y_t, p_t) = \max_{d_{t+1}} \left[u(c_t^N)^{\frac{1-\sigma}{\eta}} + \beta\left(\mathbb{E}_t V^N(y_{t+1}, p_{t+1})^{\frac{1}{\eta}}\right)\right]^{\frac{\eta}{1-\sigma}} \tag{4.19}$$

where $\sigma \geq 0$ is the coefficient of relative risk aversion, $\xi \geq 0$ is the inter-temporal elasticity of substitution (IES), and:

$$\eta = \frac{1-\sigma}{1-\frac{1}{\xi}} \tag{4.20}$$

Under deterministic consumption, we set the IES equal to the inverse of the risk aversion

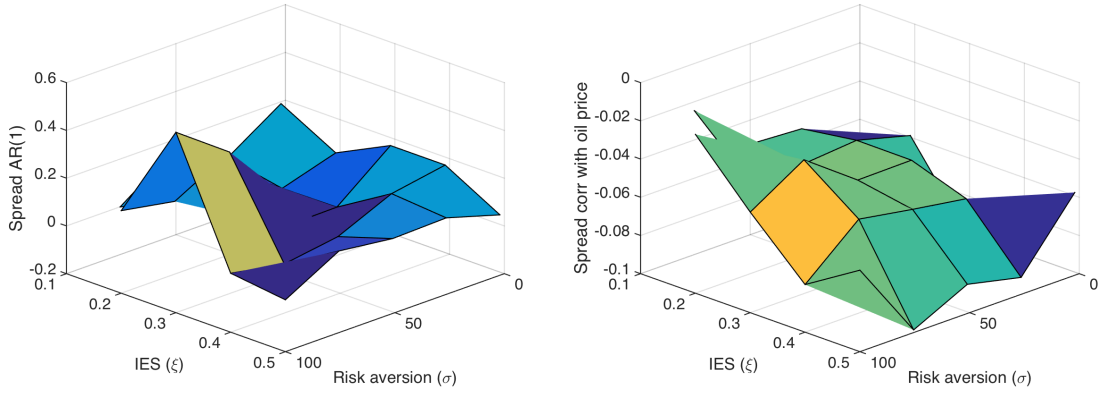
[21] Another option would be to use some form of hyperbolic preferences following work by Laibson (1997)

parameter, i.e. $\sigma = 1/\xi$ and $\eta = 1$, so the recursive preferences collapse to the standard time-separable expected discounted utility with discount factor β.

The use of recursive preferences allows us to move the IES independently from the risk-aversion parameter. Exactly what values to choose for these parameters however, is unclear. The micro and macroeconomic evidence suggests a range of possible values for these parameters in developing economy settings. I solve and simulate the model, as before, using a range of plausible values for the risk-aversion and inter-temporal substitution parameters, $\sigma \in [2, 100]$ and $\xi \in [0.1, 0.5]$.

Once again, figures 4.7a and 4.7b present the spread persistence and the spread's correlation with oil prices respectively for different combinations of risk aversion and inter-temporal substitutability. The persistence of the interest rate spread remains relatively unaffected by changes in either parameter. However, the same is not true for the correlation of the spread with the oil price. The easier it is for the sovereign to substitute consumption inter-temporally (that is a higher IES coefficient), the greater is the negative correlation between spreads and movements in the oil price. This is due to the trade-off between consumption and debt stability. The easier it is for the sovereign to substitute consumption between periods, the more volatile consumption will be. At the same time, the path for debt will be more stable. However, this immediately means that when the oil price rises again, the level of outstanding debt will be reduced relatively slowly, thus increasing the risk of default, ceteris paribus. International creditors know this, and so charge a higher risk premium when oil prices fall initially. A similar, albeit less pronounced, relationship exists with the risk aversion parameter.

Figure 4.7: Sensitivity to risk aversion, $\sigma \in [2, 100]$, and IES, $\xi \in [0.1, 0.5]$



(a) AR(1) parameter

(b) Correlations

4.6.5 Final model

The previous sections used comparative statics to investigate the behaviour of the model along key parameter dimensions. This section aims to present a sensibly parameterised model which best fits the observed data moments for the four resource rich developing countries in our sample. In particular we calibrate the model to reflect higher risk-aversion, more impatience, larger oil shares and a stronger correlation between domestic output and oil price shocks. Table 4.7 presents the full parameter set[22].

The model is then solved and simulated again over 1000 periods. Tables 4.8 and 4.9 compare the moments of the simulations data to actual data moments.

Overall, the final model calibration fits the data relatively well, and certainly better than any of the other models. The model is able to reproduce the correct ordering of autoregressive properties observed in the data. Consumption is the most persistent of the four key variables, followed by output, the oil price and then the interest rate spread. The model still under predicts the persistence and variability of the interest rate spread. However, this is substantially less than other models and is perhaps to be expected given the simplicity of the financial market structure in the model.

[22] ρ_{zh} is the off diagonal element of the coefficient matrix in equation 4.17 and σ_{hz} is the off diagonal element of the covariance matrix of the errors from the bivariate VAR.

Table 4.7: Model parameters - best fit

Parameter	Value	Description
ρ	.025	Discount rate
$\beta = \frac{1}{1+\rho}$	.980	Discount factor
σ	70	Coefficient of relative risk aversion
ξ	0.5	Inter-temporal elasticity of substitution
θ	.25	Probability of re-entry into capital market
ψ	.95	1 - output cost of default
$\bar{s}$	.212	Annual production of oil (oil share in GDP)
r^*	.004	risk-free interest rate
$\hat{r}$	.016	Risk premium (exog)
r	$r^* + \hat{r} = .020$	Default-risk-free rate
ρ_z	.841	Persistence of productivity
ρ_h	.725	Persistence of oil price
σ_z	.015	Std. dev of output shocks
σ_h	.136	Std. dev of oil price shock
ρ_{zh}	.05	Cross correlation
σ_{hz}	0	Cross std dev

Table 4.8: AR(1) model: autocorrelations and standard dev

	Data		Model	
	AR(1)	std	AR(1)	std
output (HP filt)	.841	.032	.893	.043
oil price (HP filt)	.725	.198	.657	.195
cons (HP filt)	.898	.026	.977	.107
spread	.645	.022	.304	.010

Table 4.9: AR(1) model: cross correlations with real interest rate spread

	Data	Model
output dev	-.138	-.148
oil price dev	-.279	-.160

The biggest success of the model is that it comes much closer to matching the correlation of interest rate spreads and the oil price. The model produces a spread-output correlation which is very close to that observed in the data. Crucially however, it is able to produce a spread-oil price correlation that is larger than the spread-output one, albeit lower in magnitude than that observed in the data. Countercyclical interest rate spreads and commodity prices are a key feature of the data in resource rich developing economies.

Finally, I also solve and simulating this model with oil revenues 'turned off'. The result is a substantial worsening in the model's predictive power. In particular the spread persistence and the countercyclical correlation of the spread and the oil price are dramatically reduced. See tables D.3 and D.4 in the appendix.

4.7 Conclusion

How do commodity price movements affect the pricing of sovereign debt in resource rich economies? And, what does this mean for the optimal management of volatile oil revenues in developing countries? With debt distress across developing economies on the rise, this is a pertinent question for policy makers and investors alike. This chapter sheds light on these two questions by investigating the link between commodity price movements and risk premiums in resource dependent developing economies.

I presented empirical evidence of a counter-cyclical relationship between oil prices and interest-rate spreads in these settings. I then developed a simple model that delivers this counter-cyclical relationship and showed that it comes relatively close to explaining some key macroeconomic co-movements in a sample of resource rich developing countries. The model generates an endogenous link between oil prices and interest rate risk premiums and is, to the best of my knowledge, new to the literature. This approach could be useful in a range of other applications involving optimal policy in resource rich settings going forward.

For developing countries, managing resource revenues can be challenging. Conventional wisdom suggests that part of these windfalls should be saved in a sovereign wealth fund to smooth consumption and avoid the Dutch disease. Another reason to save is to manage these inherently volatile revenue streams. Economists agree that macroeconomic volatility is bad for growth and a growing literature suggests that this is is particularly true in resource rich economies. Recent research highlights the importance of building up precautionary savings in the face of oil price volatility in a *liquidity fund*, the size of which will depend on the level of volatility, the size of the windfall and how prudent policy makers

are.

This chapter presented another reason that policy makers in developing countries may choose to use a volatility fund - to manage debt-dynamics. Large movements in the oil price can present an important source of default risk for resource dependent sovereigns. Managing this risk may be important for a country's growth trajectory, particularly in capital scarce settings. Higher default risk translates into higher interest rate premiums and so a lower rate of capital accumulation, which in turn is likely to affect the rate of growth. Policy makers should respond by building up liquid buffers to manage downward movements in the oil price. Generally speaking, these buffers should be larger the stronger the link between oil prices and the domestic economy is, the more impatient policy makers are and the more willing they are to substitute current for future consumption.

Finally, developing country governments that depend significantly on resource revenues are likely to face prolonged of substantial interest rate premiums. This can lead to sub-optimal investment and capital stocks, particularly if significant absorption constraints exist in the domestic economy. The importance of using resource revenues to set up an *investment* fund in developing countries has been emphasised in resent research. Future work could extend the model developed in this chapter to incorporate domestic physical capital. This would allow policy questions around the trade-offs between precautionary savings and domestic investment to be answered within a framework of endogenous capital accumulation and interest rate premiums.

Chapter 5

Concluding remarks

This thesis investigated the challenges and opportunities surrounding the macroeconomic management of natural resource revenues in developing economies. Natural resource wealth offers countries an opportunity to boost growth and foster sustainable economic development. However, for many the reality of resource dependence has been very different. The early work of Gelb (1988), Auty (1993) and Sachs & Warner (1995) posited the existence of a negative link between resource wealth and growth - the so called 'resource curse'. Since then, a rich literature on the economics of natural resources has developed, ranging from empirical studies of the local impacts of resource booms to theoretical investigations of optimal monetary policy in resource rich settings. Each line of enquiry has attempted to understand whether there are indeed any fundamental economic laws about resource dependence.

Whilst this has been a fruitful research agenda, there are a number of areas where progress is still being made. One of these is the identification of exogenous natural resource shocks, both at the local and the national level. In order to draw a causal interpretation of the evidence relating resource wealth to growth, one must be able to identify a change to resource wealth that is independent of underlying economic factors. Only then, can a meaningful counterfactual be constructed and any causal claims be made. Finding examples of an exogenous change in resource wealth has proven extremely challenging. Even where these have been identified, questions around the external validity of the results and

the wider applicability of policy conclusions have arisen.

Chapter 2 of this thesis contributes to this literature by developing an exogenous cross-country measure of resource wealth, and then uses it to investigate the dynamic macroeconomic response of resource rich, low income countries to global commodity price shocks. One important contribution is the development of a new index to measure natural resource wealth that is exogenous to demand and financial conditions at the global level. Future work could apply this panel-VAR approach to study the effect that natural resource wealth has on a number of different economic, social and political outcomes, both at the national and sub-national level.

Researchers are also beginning to think more carefully about the mechanisms through which natural resource wealth can impact aggregate economic variables. Given the difficulties involved when using cross-country evidence, particularly in controlling for confounding factors as well as drawing inference, some are beginning to focus on resource shocks at the local level. Local studies have two advantages. First, they allow the researcher to hold constant any aggregate economic or political factors that might confound empirical analysis at the cross-country level. Second, by focusing on local effects they abstract from general equilibrium processes that take place at the national level, thus allowing specific mechanisms to be tested. In addition, local studies often allow for a more rigorous defence of exogeneity than do cross-country studies.

Chapter 3 of this thesis investigated the economic impacts of high quality natural assets at the local level. Focusing on a specific local non-market asset - surf breaks - allowed us to exploit three different natural experiments to study the impact these assets have on growth. It also allowed us to abstract from the direct income effects that arise from typical market resources such as oil and gas. Instead we focused on the indirect effect that these local assets might have on growth at the local level. This work adds to an already vast literature on economic geography. By using a novel set of natural experiments, we presented evidence of the importance of location fundamentals, in this case local amenities, for economic develop-

ment. This work takes an innovative approach to studying the economic impact of natural amenities and could in principle be applied to a number of different economic questions.

Finally, the ultimate aim of much of the research on natural resources is to guide policy makers through the many complex decisions required in managing a nation's natural assets. This is particularly useful for developing and emerging economies for whom natural resource wealth has the potential to accelerate the development process. However, understanding the developing country context remains vital when thinking about optimal resource management in these settings. The permanent income hypothesis (PIH) often forms the basis for advice to resource rich economies. This advice traditionally includes the setting up of an *intergenerational* sovereign wealth fund, in which foreign assets are accumulated over time using resource revenues. A key assumption in the PIH is the ability for an economy to borrow or lend at a constant and exogenous world interest rate (Friedman (1957)). This assumption is likely to be inappropriate in a developing country context which can typically save more readily than it can borrow, often facing a substantial borrowing premium over world interest rates as indebtedness rises. Recent work by van der Ploeg & Venables (2017), for example, shows that when the interest rate is an increasing function of the debt stock, a larger share of resource revenues should be invested in domestic assets.

The aim of chapter 4 was to investigate one possible channel through which the link between debt levels and interest rate premiums may operate in resource rich economies. That channel is commodity prices movements. A model with stochastic resource revenues and sovereign default risk delivers a bond price that is decreasing in the level of indebtedness and negatively correlated with the oil price. Once the model is augmented to reflect some key developing country characteristics, it can explain a large proportion of business cycle fluctuations in these economies. This work adds to the existing literature by endogenising the interest rate premiums that can arise in resource rich settings. This is a new approach in the literature and could be useful in a range of other applications and future work. The model provides an additional rationale for a *volatility* fund in which liquidity buffers are accumulated to manage debt repayments. These buffers should be larger the stronger

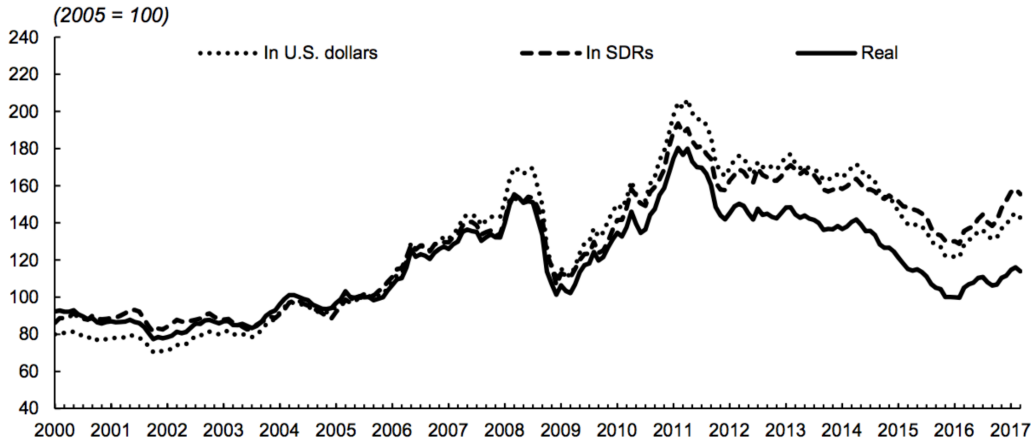
the link between oil prices and the domestic economy is, the more impatient policy makers are and the more willing they are to substitute between current and future consumption.

Appendix A

General

A.1 Figures

Figure A.1: Global commodity price indices (non-fuel)



Indices comprise 60 price series for 44 non-fuel primary commodities. Weights are based on the 2002-2004 average of world export earnings.
Real indices deflated by U.S. CPI.
Source: IMF primary commodity price index charts, April 20, 2017

Figure A.2: Oil reserves (bb)

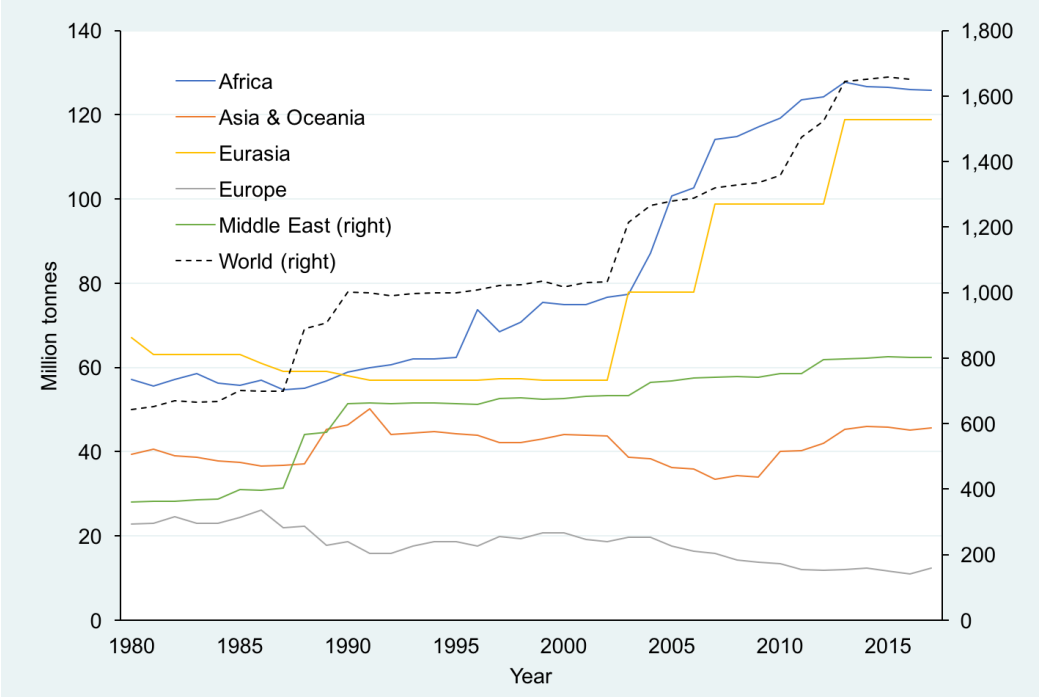


Figure A.3: Gas reserves (tcf)

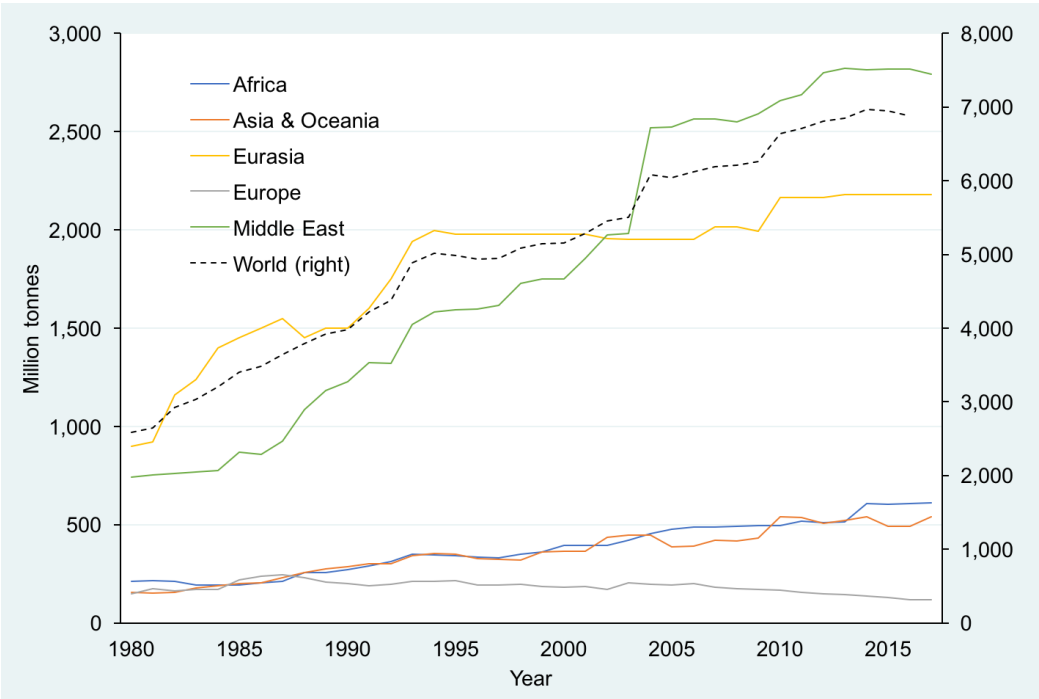


Figure A.4: Mineral reserves (Tril usd)

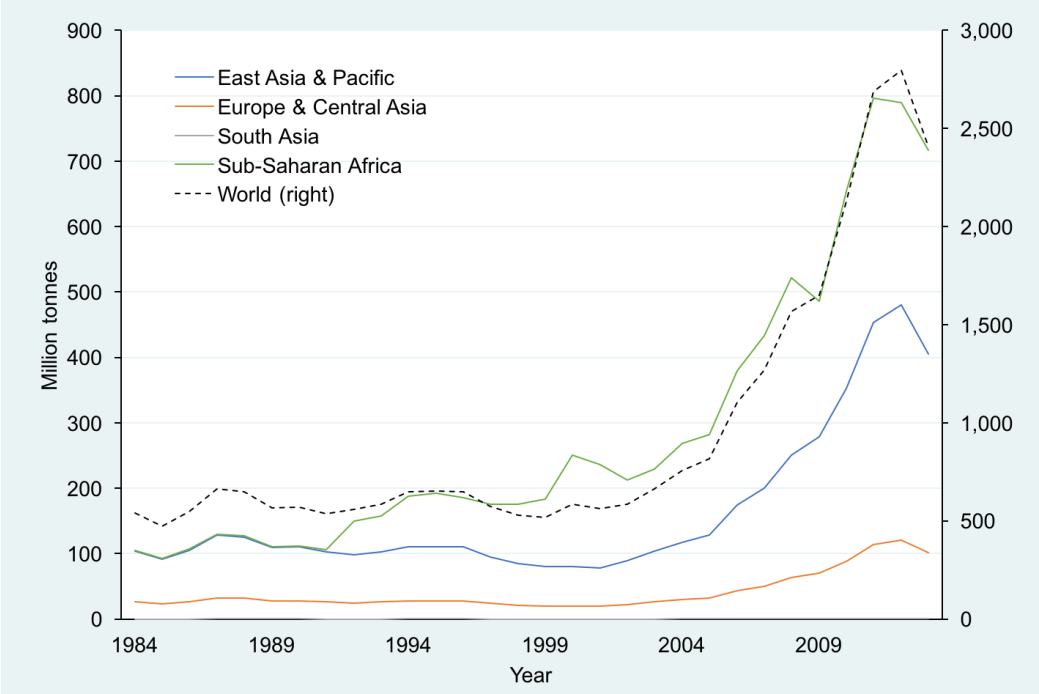
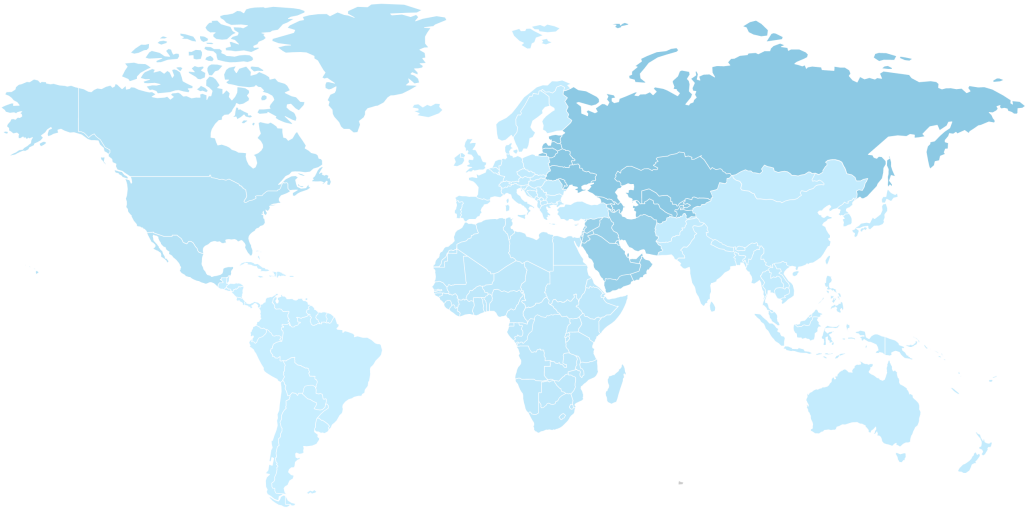
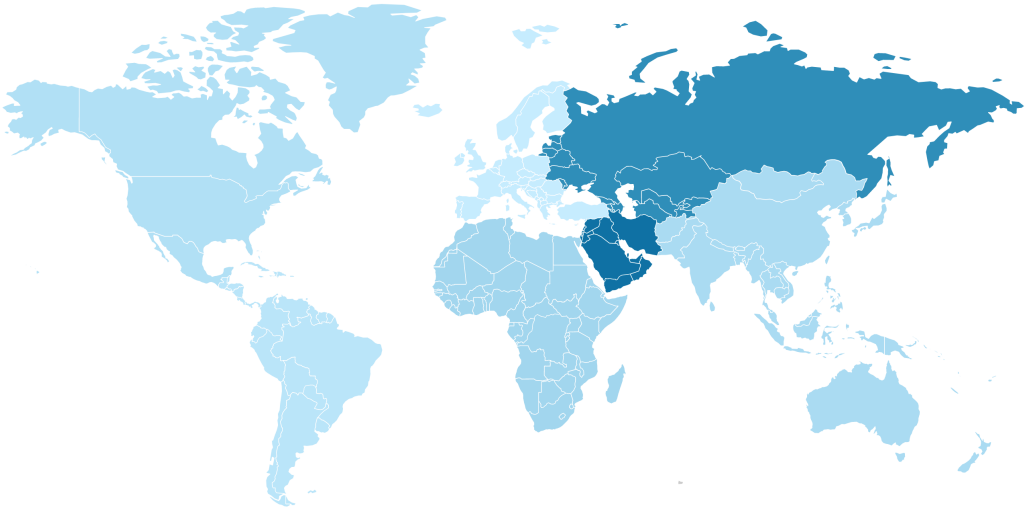


Figure A.5: Natural Gas proved reserves

(a) 1980



(b) 2016



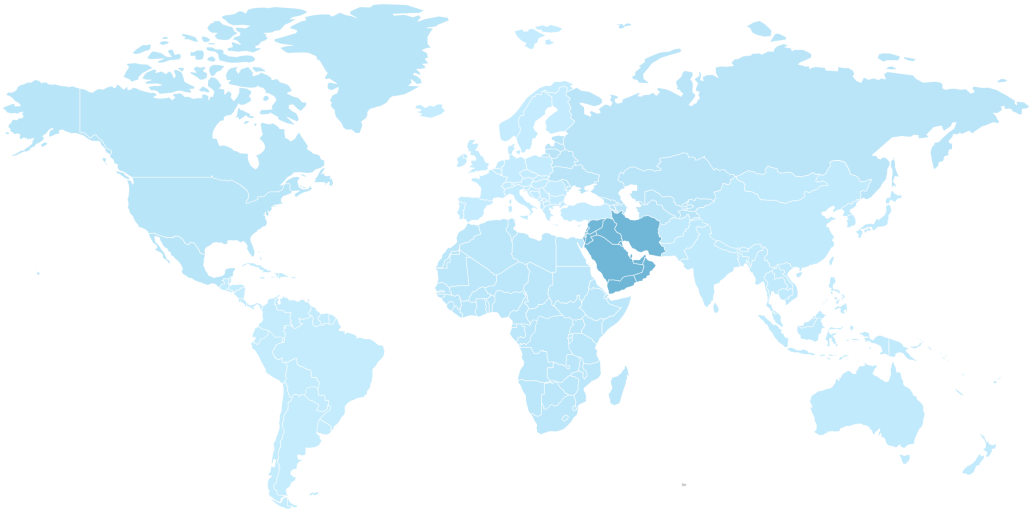
(c)



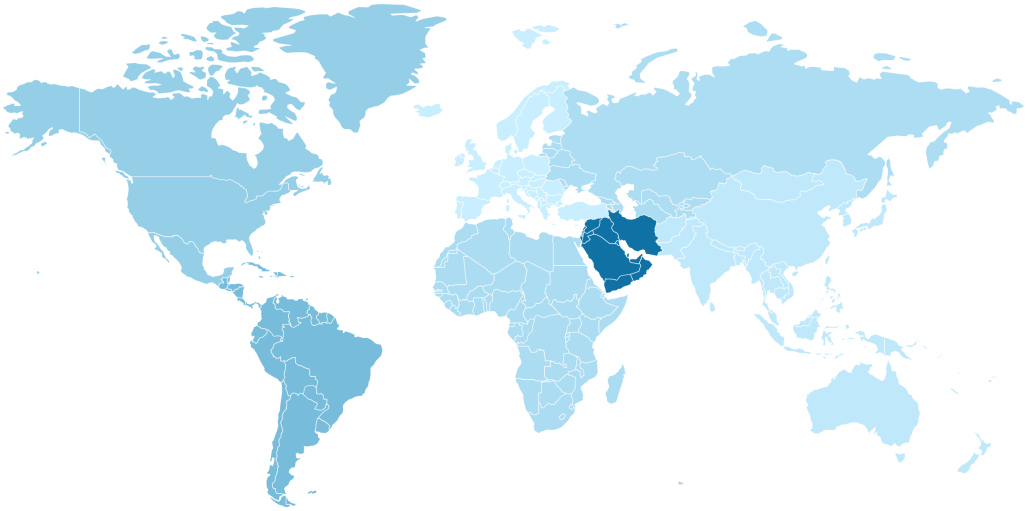
Trillion cubic feet per year
Source: US EIA International Energy Statistics, May 10, 2017

Figure A.6: Crude oil proved reserves

(a) 1980



(b) 2016



(c)



Billion barrels per year
Source: US EIA International Energy Statistics, May 10, 2017

A.2 Tables

Table A.1: Petroleum export share

Share of total exports (%)	2000	2015
0-0.1	68.2	61.7
0.1-0.2	9.5	16.3
0.2-0.3	5.4	7.1
0.3-0.4	0.7	2.8
0.4-0.5	2.0	2.1
0.5-0.6	2.7	2.1
0.6-0.7	2.7	2.1
0.7-0.8	1.4	0.7
0.8-0.9	4.7	2.1
0.9-1	2.7	2.8

Source: US EIA International Energy Statistics, May 10, 2017. Own calculations

Figure A.7: Wealth and per capita wealth by type of capital and income group, 1995 and 2005

	1995					2005				
Income Group	Total Wealth (US$ billions)	Per Capita Wealth (US$)	Intangible Capital (%)	Produced Capital (%)	Natural Capital (%)	Total Wealth (US$ billions)	Per Capita Wealth (US$)	Intangible Capital (%)	Produced Capital (%)	Natural Capital (%)
Low income	2,447	5,290	48	12	41	3,597	6,138	57	13	30
Lower middle income	33,950	11,330	45	21	34	58,023	16,903	51	24	25
Upper middle income	36,794	73,540	68	17	15	47,183	81,354	69	16	15
High income OECD	421,641	478,445	80	18	2	551,964	588,315	81	17	2
World	504,548	103,311	76	18	6	673,593	120,475	77	18	5

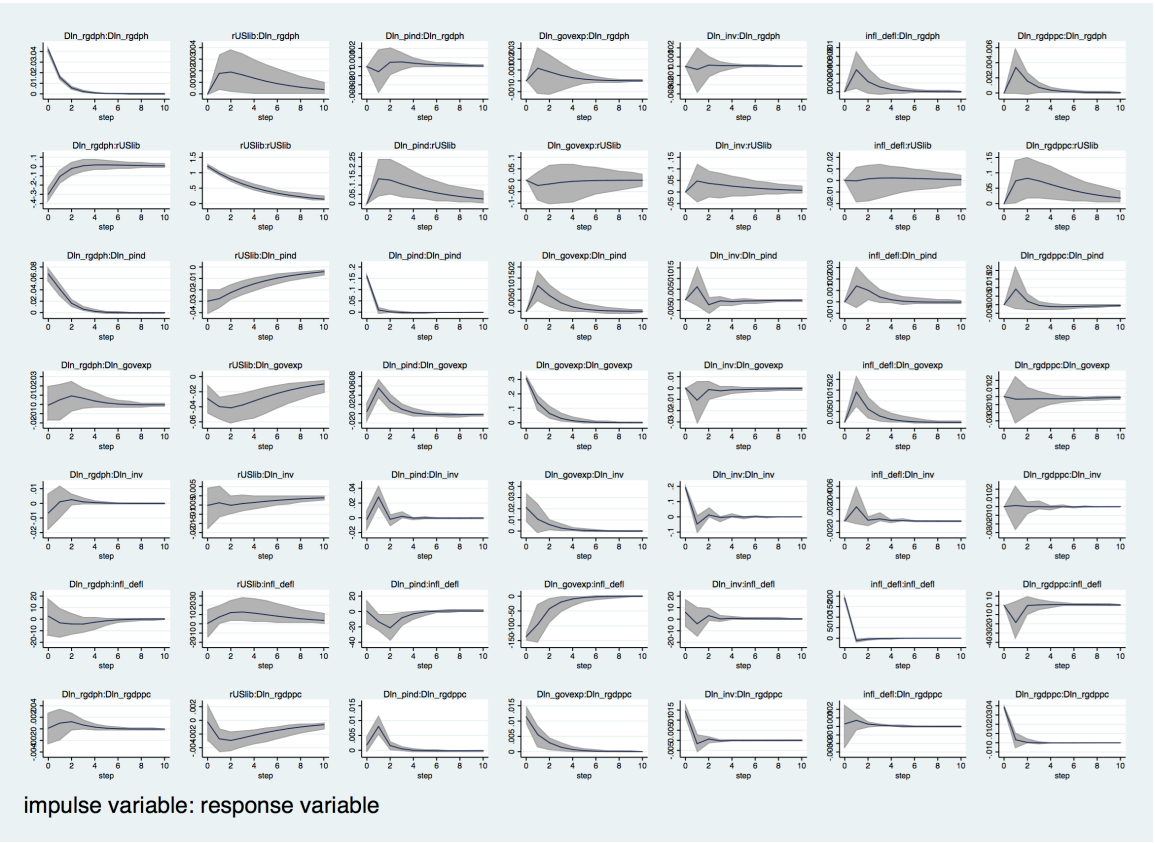
Source: World Bank (2011)
Figures are based on the set of countries for which wealth accounts are available from 1995 to 2005. Data in this table do not include high-income oil exporters.

Appendix B

Commodity price shocks

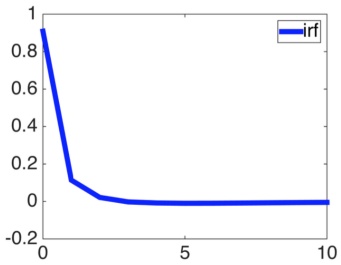
B.1 Figures

Figure B.1: Panel VAR IRFs - Cholesky - IRFs with 95% confidence bands

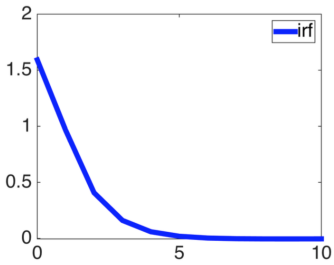


IRFs and confidence 95% intervals are presented before scaling.

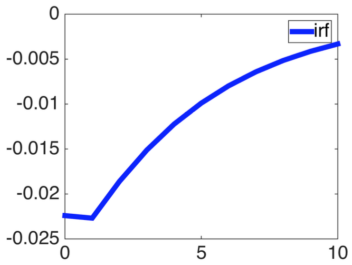
Figure B.2: Panel VAR IRFs - Cholesky - no OPEC members



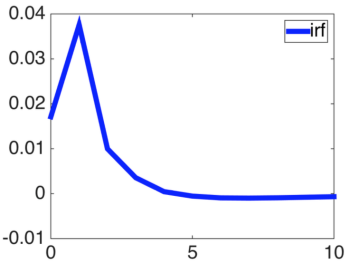
(a) imp: cpind - resp: cpind



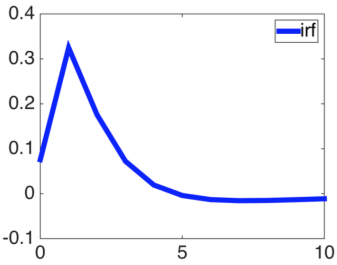
(b) imp: gdph - resp: cpind



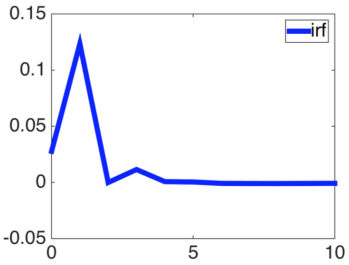
(c) imp: libor - resp: cpind



(d) imp: cpind - resp: rgdppc

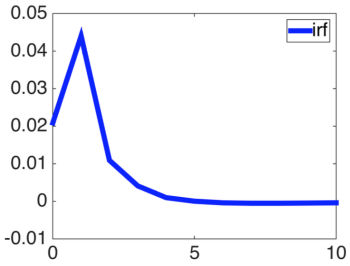


(e) imp: cpind - resp: g

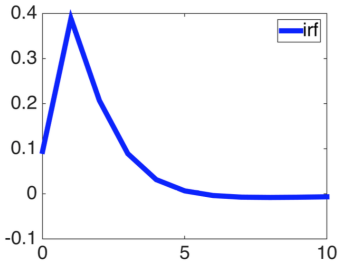


(f) imp: cpind - resp: i

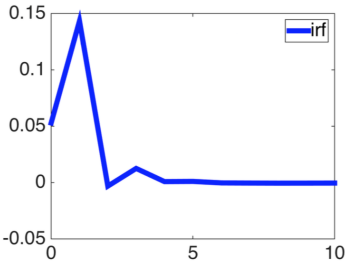
Figure B.3: Panel VAR IRFs - block recursive - no OPEC members



(a) imp: cpind - resp: rgdppc

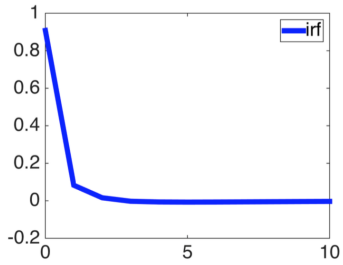


(b) imp: cpind - resp: g

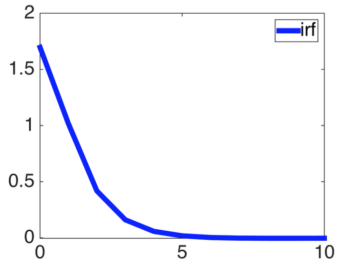


(c) imp: cpind - resp: i

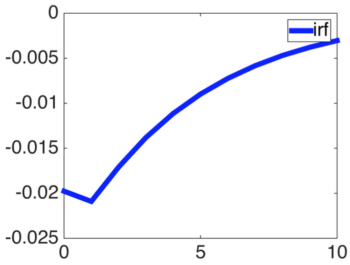
Figure B.4: Panel VAR IRFs - Cholesky - low and middle income countries



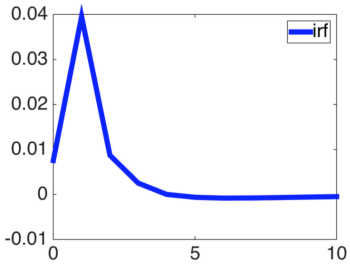
(a) imp: cpind - resp: cpind



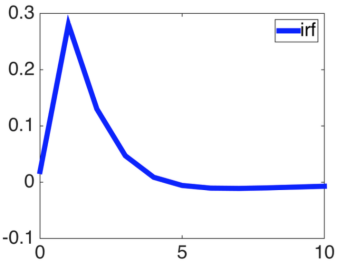
(b) imp: gdph - resp: cpind



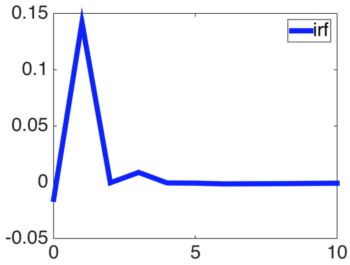
(c) imp: libor - resp: cpind



(d) imp: cpind - resp: rgdppc

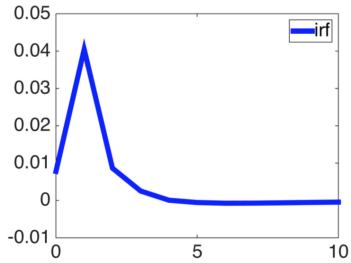


(e) imp: cpind - resp: g

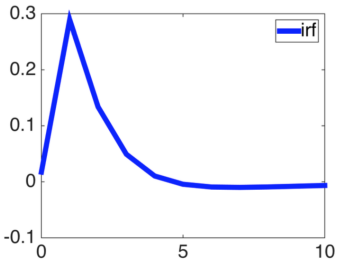


(f) imp: cpind - resp: i

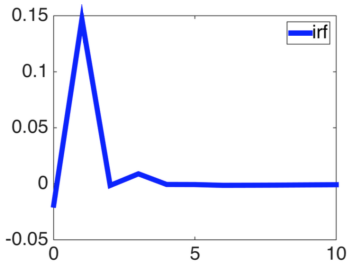
Figure B.5: Panel VAR IRFs - block recursive - low and middle income countries



(a) imp: cpind - resp: rgdppc

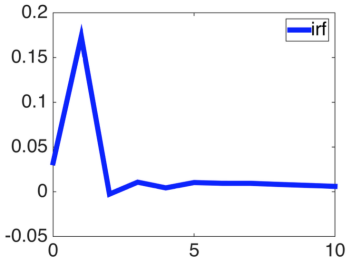


(b) imp: cpind - resp: g

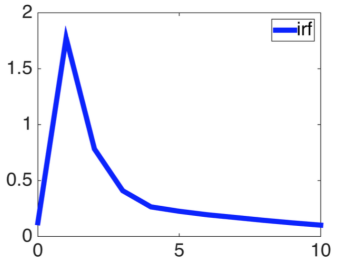


(c) imp: cpind - resp: i

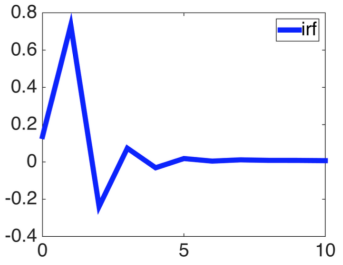
Figure B.6: Panel VAR IRFs - block recursive - fixed exchange rate



(a) imp: cpind - resp: rgdppc

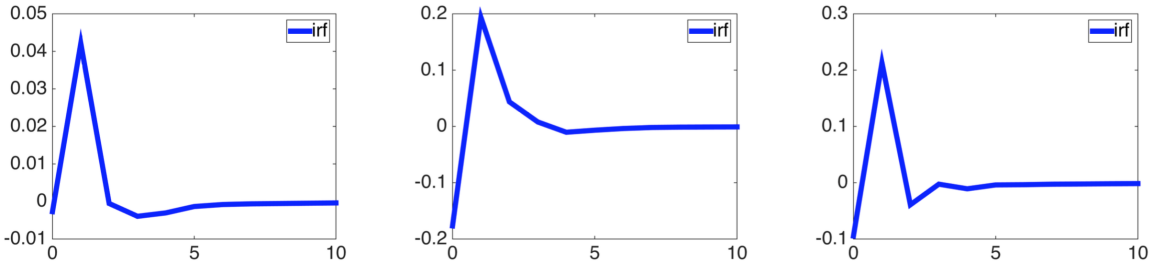


(b) imp: cpind - resp: g



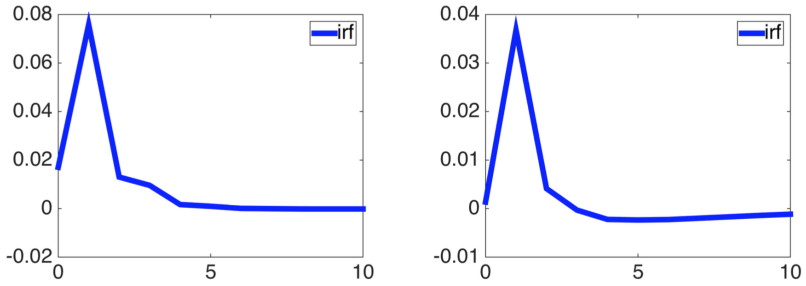
(c) imp: cpind - resp: i

Figure B.7: Panel VAR IRFs - block recursive - flexible exchange rates



(a) imp: cpind - resp: rgdppc (b) imp: cpind - resp: g (c) imp: cpind - resp: i

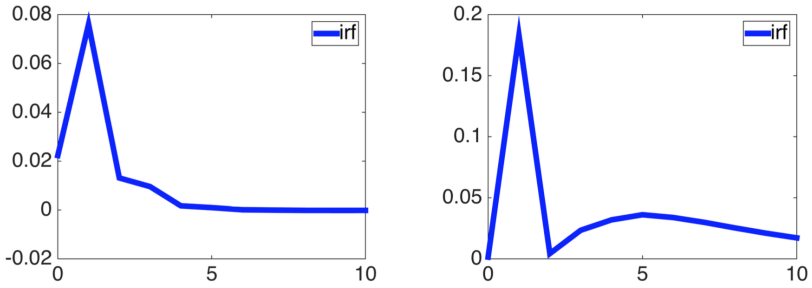
Figure B.8: Panel VAR IRFs - Cholesky - exposure levels



(a) imp: cpind - resp: rgdppc (b) imp: cpind - resp: rgdppc

IRFs for CPIND impulse on RGDPPC for different levels of exposure to commodity price shocks. Panel (a) shows countries with zero resource exports in 2013, panel (b) shows countries with positive resource exports but > USD 10 million.

Figure B.9: Panel VAR IRFs - block recursive - exposure levels



(a) imp: cpind - resp: rgdppc (b) imp: cpind - resp: rgdppc

IRFs for CPIND impulse on RGDPPC for different levels of exposure to commodity price shocks. Panel (a) shows countries with zero resource exports in 2013, panel (b) shows countries with positive resource exports but > USD 10 million.

B.2 Tables

Table B.1: Summary statistics

country	Region	2013 income group				Dln_rgdppc		Dln_pind		Dln_govexp		Dln_inv		infl_defl		rva_sharemanuf		rva_shareagric		Dln_rva_sharemanuf		Dln_rva_shareagric	
			N	min	max	mean	sd	mean	sd	mean	sd	mean	sd	mean	sd	mean	sd	mean	sd	mean	sd	mean	sd
Indonesia	EAP	LIC	28	1986	2013	0.036	0.041	0.053	0.128	-0.036	0.332	0.061	0.113	12.252	13.138	24.610	2.744	14.365	2.403	0.012	0.021	-0.022	0.035
Philippines	EAP	LMIC	28	1986	2013	0.020	0.023	0.053	0.128	0.025	0.114	0.051	0.081	6.937	4.461	24.365	1.210	13.315	1.640	-0.004	0.018	-0.018	0.027
Bolivia	LAC	LMIC	27	1986	2012	0.015	0.019	0.061	0.128	-0.058	0.317	0.057	0.121	17.209	42.845	13.719	0.301	13.828	1.094	0.003	0.020	-0.010	0.033
Chile	LAC	HIC	28	1986	2013	0.039	0.026	0.053	0.128	0.010	0.097	0.084	0.102	9.484	7.602	16.192	1.577	3.704	0.329	-0.008	0.022	0.007	0.040
Colombia	LAC	UMIC	13	2001	2013	0.027	0.017	0.045	0.350	0.050	0.106	0.095	0.052	5.352	1.772	14.632	0.973	7.876	0.763	-0.012	0.028	-0.017	0.020
Ecuador	LAC	UMIC	28	1986	2013	0.012	0.026	0.048	0.263	0.030	0.087	0.038	0.092	3.274	10.231	12.414	3.218	9.419	0.758	0.011	0.116	0.009	0.044
Guatemala	LAC	LMIC	28	1986	2013	0.011	0.013	0.053	0.128	-0.019	0.201	0.043	0.089	10.362	10.530	20.660	2.061	13.373	0.923	-0.011	0.010	-0.006	0.015
Mexico	LAC	UMIC	28	1986	2013	0.009	0.031	0.053	0.128	-0.105	0.282	0.033	0.096	23.189	32.909	17.011	0.921	3.513	0.338	0.001	0.023	-0.011	0.038
Peru	LAC	UMIC	28	1986	2013	0.020	0.057	0.053	0.128	-0.405	1.017	0.089	0.173	356.014	1163.234	17.167	0.750	6.927	0.632	0.001	0.034	0.004	0.054
Trinidad And Tobago	LAC	HIC	18	1991	2008	0.046	0.039	0.082	0.217	0.016	0.104	0.066	0.278	7.069	7.187	17.651	3.867	1.075	0.507	0.036	0.040	-0.089	0.123
Algeria	MENA	UMIC	24	1986	2009	0.004	0.029	0.046	0.117	-0.089	0.234	0.022	0.086	13.938	13.622	5.865	1.135	7.391	0.945	-0.016	0.047	0.022	0.094
Egypt	MENA	LMIC	28	1986	2013	0.024	0.016	0.053	0.128	-0.041	0.123	0.016	0.102	10.713	6.218	15.652	1.713	14.620	1.236	0.011	0.021	-0.010	0.014
Iran	MENA	UMIC	19	1986	2007	0.014	0.057	0.067	0.268	-0.158	0.204	0.032	0.112	20.460	7.972	8.320	2.099	9.898	0.731	0.049	0.045	0.004	0.057
Morocco	MENA	LMIC	27	1986	2012	0.023	0.045	0.061	0.123	0.050	0.060	0.047	0.053	3.245	3.207	16.363	1.413	16.483	3.039	-0.011	0.054	-0.010	0.220
India	SA	LIC	28	1986	2013	0.044	0.023	0.053	0.129	0.005	0.095	0.077	0.062	7.255	2.551	15.444	0.720	23.560	6.532	0.003	0.030	-0.032	0.036
Botswana	SSA	UMIC	27	1986	2012	0.037	0.044	0.061	0.123	-0.002	0.137	0.080	0.089	9.600	5.410	6.394	0.902	2.964	0.839	0.017	0.084	-0.030	0.115
Cameroon	SSA	LMIC	27	1986	2012	-0.011	0.040	0.050	0.267	0.027	0.227	0.015	0.133	3.200	4.414	18.265	0.856	19.852	1.866	0.001	0.033	0.016	0.041
Cote d'Ivoire	SSA	LMIC	20	1986	2005	-0.013	0.029	0.017	0.101	-0.020	0.269	-0.013	0.135	4.133	10.631	15.726	1.149	26.211	1.090	0.013	0.063	0.001	0.046
DRC	SSA	LIC	16	1994	2009	-0.017	0.046	0.062	0.123	-1.089	1.973	0.035	0.133	1962.317	6645.197	21.475	3.004	24.664	2.609	-0.019	0.094	-0.005	0.048
Gabon	SSA	UMIC	22	1986	2007	-0.009	0.063	0.044	0.263	-0.024	0.265	-0.010	0.279	5.275	14.703	3.494	1.000	5.599	0.634	0.012	0.097	0.001	0.093
Liberia	SSA	LIC	12	2001	2012	0.042	0.162	0.131	0.246	0.084	0.385	0.024	0.539	8.059	8.334	5.561	0.903	71.173	4.038	0.239	0.791	-0.002	0.039
Madagascar	SSA	LIC	24	1986	2009	-0.007	0.047	0.046	0.372	-0.082	0.301	0.070	0.223	14.788	9.989	13.963	0.520	28.981	1.348	-0.002	0.038	0.000	0.038
Mali	SSA	LIC	27	1986	2012	0.018	0.035	0.061	0.123	0.039	0.302	0.078	0.257	4.031	7.180	9.279	1.768	41.858	4.870	-0.013	0.123	0.003	0.077
Mauritania	SSA	LMIC	14	1999	2012	0.018	0.047	0.124	0.106	0.006	0.122	0.157	0.327	6.594	8.181	8.523	1.990	30.517	2.951	-0.008	0.161	-0.019	0.063
Mozambique	SSA	LIC	27	1986	2012	0.039	0.042	0.002	0.273	-0.175	0.465	0.087	0.149	28.382	36.820	12.545	2.995	29.055	3.203	0.000	0.096	-0.008	0.067
Nigeria	SSA	LMIC	8	2006	2013	0.034	0.014	0.083	0.246	0.062	0.237	0.130	0.255	19.638	34.567	7.191	1.007	25.080	0.779	0.051	0.056	-0.007	0.020
Rep Congo	SSA	LMIC	26	1986	2011	-0.004	0.039	0.052	0.272	0.002	0.287	0.019	0.240	6.719	17.836	3.796	0.651	4.612	0.252	0.005	0.083	0.004	0.052
Sierra Leone	SSA	LIC	24	1986	2011	0.001	0.073	0.077	0.179	-0.219	0.382	0.139	0.961	40.431	42.339	2.952	0.330	49.552	2.997	-0.017	0.081	0.005	0.018
Sudan	SSA	LMIC	23	1986	2008	0.028	0.040	0.044	0.120	-0.259	0.518	0.092	0.150	43.342	43.071	7.510	0.806	36.593	3.480	0.013	0.103	0.003	0.091
Zambia	SSA	LMIC	28	1986	2013	0.005	0.038	0.053	0.128	-0.243	0.440	0.103	0.172	42.095	42.328	10.631	0.679	22.065	3.367	-0.004	0.054	0.000	0.149
Total			705	1986	2013	0.017	0.047	0.055	0.187	-0.084	0.474	0.059	0.242	71.795	1039.664	13.342	6.210	18.018	14.180	0.007	0.123	-0.007	0.078

2013 income group	N	Dln_rgdppc		Dln_pind		Dln_govexp		Dln_inv		infl_defl		rva_sharemanuf		rva_shareagric		Dln_rva_sharemanuf		Dln_rva_shareagric	
		mean	sd	mean	sd	mean	sd	mean	sd	mean	sd	mean	sd	mean	sd	mean	sd	mean	sd
LIC	186	0.020	0.059	0.061	0.197	-0.184	0.529	0.071	0.305	259.689	845.694	13.229	1.623	35.401	3.500	0.025	0.159	-0.008	0.045
LMIC	284	0.013	0.030	0.059	0.156	-0.039	0.243	0.060	0.155	14.515	19.024	13.533	1.153	19.712	1.810	0.005	0.056	-0.005	0.064
UMIC	189	0.014	0.041	0.052	0.205	-0.088	0.292	0.048	0.122	54.638	156.232	10.662	1.375	6.698	0.705	0.008	0.059	-0.002	0.064
HIC	46	0.043	0.033	0.068	0.172	0.013	0.101	0.075	0.190	8.277	7.394	16.922	2.722	2.390	0.418	0.014	0.031	-0.041	0.082

Table B.2: SUR model - all sectors - no controls

VARIABLES	(1) agriculture	(2) construction	(3) manufactures	(4) minerals	(5) transport	(6) wholesale	(7) other
cshock	0.0786	-0.0191	-0.186**	-0.13	0.106**	0.0593	0.135
	(-0.779)	(-0.396)	(-2.107)	(-0.629)	(-2.476)	(-0.927)	(-1.081)
L.cshock	-0.0913	0.0188	-0.0955	-0.301	0.140***	0.0926	0.268**
	(-0.882)	(-0.44)	(-1.078)	(-1.459)	(-2.926)	(-1.458)	(-2.256)
L2.cshock	0.0325	-0.0346	-0.0745	-0.429**	0.127***	0.0918	0.272**
	(-0.337)	(-0.883)	(-0.943)	(-2.170)	(-2.945)	(-1.438)	(-2.284)
L3.cshock	0.0136	0.00336	0.0355	-0.392*	0.0391	0.0734	0.175
	(-0.135)	(-0.0884)	(-0.439)	(-1.830)	(-0.945)	(-1.171)	(-1.309)
initial rvashare_agric	1.153***						
	(-48.63)						
initial rvashare_constr		-0.193***					
		(-11.07)					
initial rvashare_manuf			0.278***				
			(-18.01)				
initial rvashare_min				0.831***			
				(-18.4)			
initial rvashare_trans					0.843***		
					(-43.68)		
initial rvashare_whole						-0.0151	
						(-0.304)	
initial rvashare_other							0.566***
							(-9.458)
Observations	577	577	577	577	577	577	577
Robust z-statistics in parentheses							
*** p¡0.01, ** p¡0.05, * p¡0.1							

Shock = cshock; Country and year fixed effects included; Low-income, resource-rich countries; standard errors clustered by country and year

Table B.3: SUR model - all sectors - national controls, no OPEC

VARIABLES	(1) agriculture	(2) construction	(3) manufactures	(4) minerals	(5) transport	(6) wholesale	(7) other
ln_govexp	-0.472	-0.156*	-0.711***	-0.121	0.342**	-0.0179	0.636***
	(-1.539)	(-1.841)	(-4.677)	(-0.343)	(-2.387)	(-0.0988)	(-3.352)
ln_inv	-0.318	1.883***	-0.875***	-0.692	-0.293	-0.315	0.27
	(-0.713)	(-7.88)	(-3.741)	(-1.139)	(-1.609)	(-0.853)	(-0.761)
infl_defl	-0.00029	-0.000186**	0.000326**	-0.000171	-0.000171	0.00018	0.000529***
	(-0.861)	(-2.520)	(-2.089)	(-0.741)	(-0.497)	(-0.978)	(-3.423)
ln_rgdppc	-6.762***	-1.318***	6.376***	6.435***	-1.127**	-0.0671	-2.837**
	(-4.958)	(-2.949)	(-7.767)	(-3.502)	(-2.006)	(-0.0751)	(-2.256)
cshock	-0.0176	0.0353	-0.221***	-0.0148	0.0941**	0.0285	0.13
	(-0.160)	(-0.855)	(-2.638)	(-0.0714)	(-1.97)	(-0.397)	(-1.287)
L.cshock	-0.0899	0.0393	-0.106	-0.223	0.135***	0.0439	0.244**
	(-0.876)	(-0.935)	(-1.346)	(-1.067)	(-2.586)	(-0.631)	(-2.472)
L2.cshock	0.0119	-0.0181	-0.125	-0.16	0.0988**	0.0616	0.186*
	(-0.13)	(-0.416)	(-1.641)	(-0.906)	(-1.963)	(-0.867)	(-1.886)
L3.cshock	0.0324	0.055	0.0143	-0.24	0.0195	0.0332	0.0594
	(-0.331)	(-1.432)	(-0.192)	(-1.387)	(-0.408)	(-0.503)	(-0.606)
initial rvashare_agric	0.580***						
	(-5.651)						
initial rvashare_constr		-0.0986**					
		(-2.301)					
initial rvashare_manuf			0.490***				
			(-17.94)				
initial rvashare_min				0.199*			
				(-1.674)			
initial rvashare_trans					0.643***		
					(-5.117)		
initial rvashare_whole						0.312***	
						(-3.989)	
initial rvashare_other							1.071***
							(-7.579)
Observations	577	577	577	577	577	577	577

Robust z-statistics in parentheses
*** p¡0.01, ** p¡0.05, * p¡0.1

Shock = cshock; Country and year fixed effects included; Low-income, resource-rich countries; standard errors clustered by country and year

Appendix C

Local impacts of resource booms

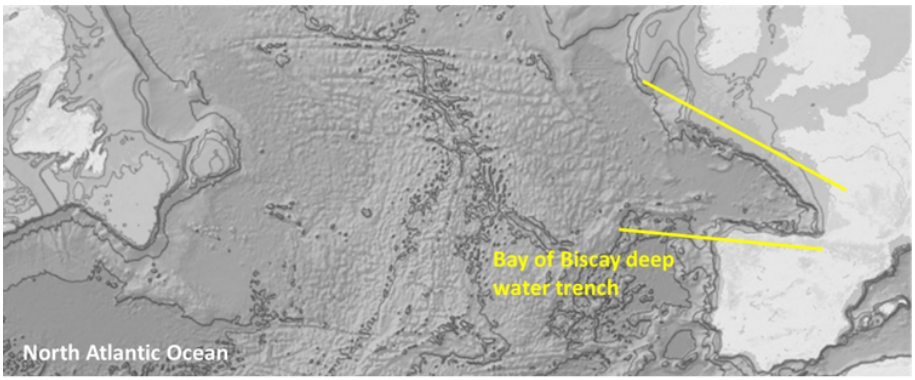
C.1 What makes a good wave?

Waves are created by wind acting on the surface of the ocean. The wind creates lines of swell which propagate up to 15,000km away from the originating low pressure system. During their travels the waves disperse, as longer wavelengths speed to the front; and group, as different wavelengths cancel and amplify one another. On their journey the waves warp and bend as they interact with the underwater topography (bathymetry), in a process called refraction. This can make the waves bigger, smaller, longer, shorter, faster, slower, fatter or hollower, depending on the bathymetry along the way. As waves approach the coast they break: the base slows as it interacts with the sea floor while the peak continues at speed. It is at the point of breaking that surfers draw off some of the waves' energy for their sport.[1] After the wave breaks it imparts energy to the shore, shaping sand bars which in turn affect breaking in a "self-organising system" (Butt & Russell (2004)).[2]

[1] Surfers ride waves by balancing the force of water pushing the surfboard up the wave's face against gravity pulling it down.

[2] Butt & Russell (2004) provide an introduction to oceanography and coastal engineering with a focus on surfing.

Figure C.1: Bay of Biscay trench



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

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

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

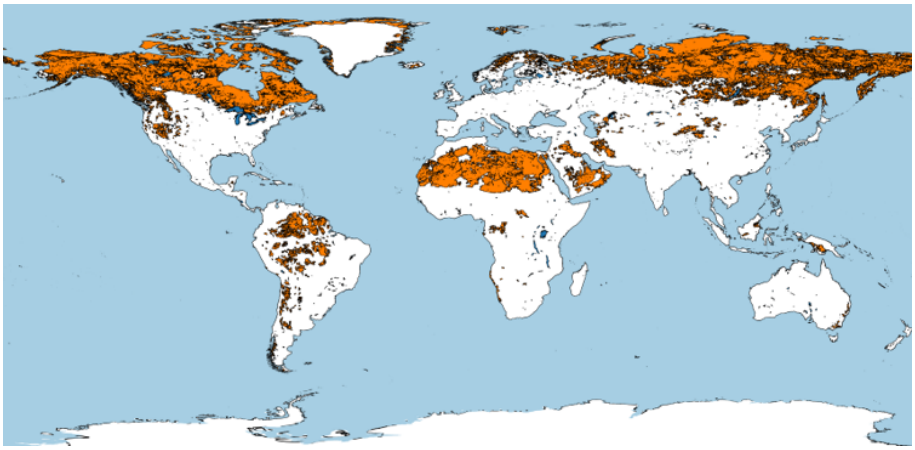
unlikely to be.

C.2 Overglow

Overglow is when light emitted in one pixel is also recorded in surrounding pixels. It has been documented as an issue when light is used to measure the level of economic activity (Doll (2008)). Here we test how overglow changes with emitted light in two ways: comparing light either side of a wasteland boundary (following Pinkovskiy (2013)), and light emitted from isolated pixels. Both tests find that overglow varies linearly with emitted light (as do Townsend & Bruce (2010)), so should not bias our study of light growth.

C.2.1 Test 1: Wasteland boundaries

Figure C.2: Global wastelands



1000 largest wasteland areas in the Anthropogenic Biomes v2 dataset.

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

$$\Delta Y^{in}_{j,2006} = \alpha + \beta \Delta Y^{out}_{j,2006} + \epsilon_t \tag{C.1}$$

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

Table C.1: Wasteland regressions

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

Results of regressing light growth 10km inside wasteland borders on that 10km outside.

C.2.2 Test 2: Isolated light sources

We also study overglow from a small number of isolated light sources in Algeria, using gas flare data from the NOAA Earth Observation Group, and wasteland data described above. We use gas flares that i) are in wastelands with no habitation or physical economic activity within 15km, ii) come from a $< 1\text{km}^2$ site as identified on Google Earth, and iii) have the brightest light in the central pixel (Figure C.3). This gives us an insight into how overglow is recorded by the DMSP-OLS satellites.

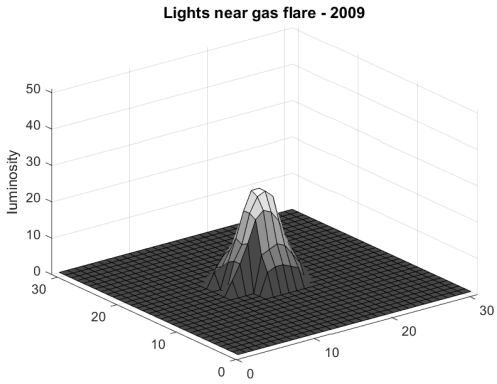
Figure C.3: Gas flares, Algeria



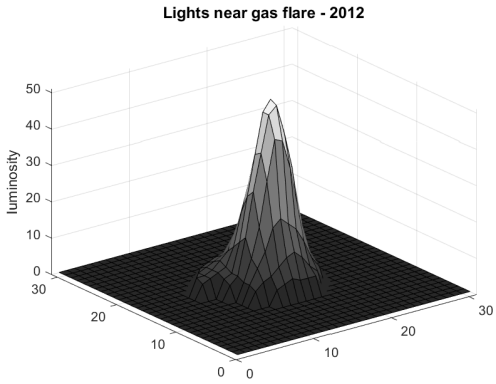
(a) Night-time lights



(b) Google earth (zoomed)



(c) 2009 light histogram

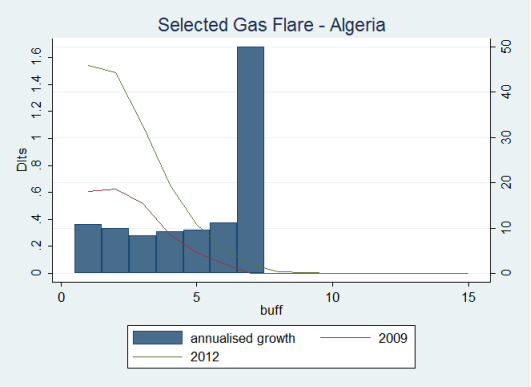


(d) 2012 light histogram

Example isolated gas flare site in Algerian wasteland

Figure C.4: Overglow comparison

(a) Same site, 2009 vs 2012



(b) Different sites, 2012

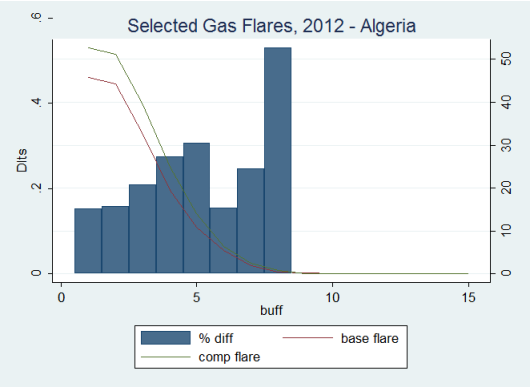


Figure C.5a illustrates the change in light between 2009 and 2012 at various distances from the flare. It shows that overglow grows proportionally to light in the central pixel, for non-trivial levels of light. This may be due to differences in the intensity of emitted light or the sensitivity of the satellite. Figure C.5b illustrates the difference in light between the original site, and another in the Algerian desert, both in 2012. Overglow differs from 15-30% but not systematically. We conclude that overglow varies proportionally to emitted light, and so will not bias our study of changes in light emissions.

Table C.2: Light growth regressions, 5-10km

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

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

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

Table C.3: Light growth regressions, 10-50km

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

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

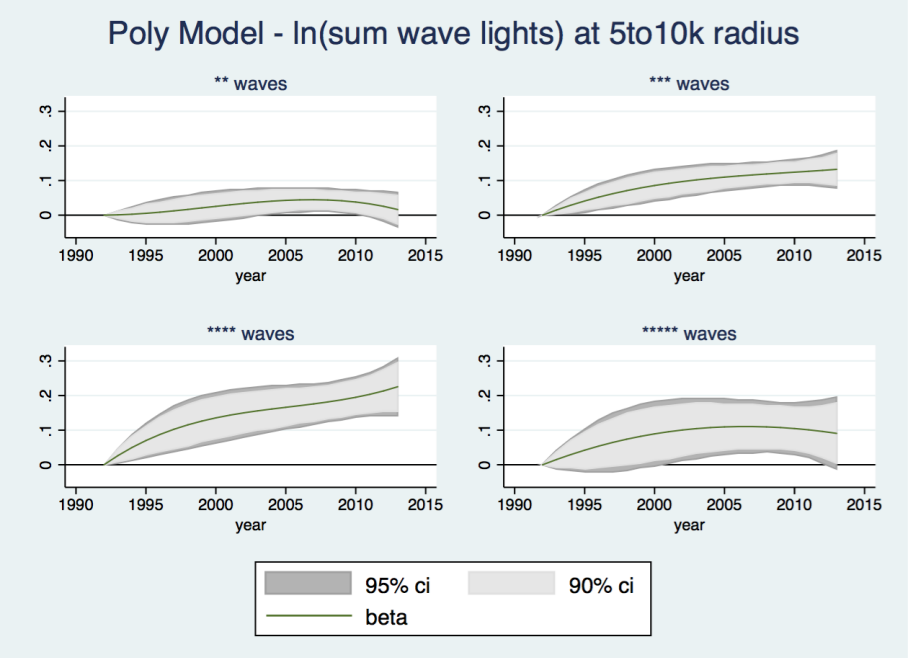
The effect of break quality on the change in ln(lights) from 10-50km of each break, from 1992-2013. Fixed effects are at country and zone level. Standard errors are clustered at the country and zone level, and allow for spatial correlation within 100km, and autocorrelation to 3 periods (Conley).

C.3 Experiment I: Additional results

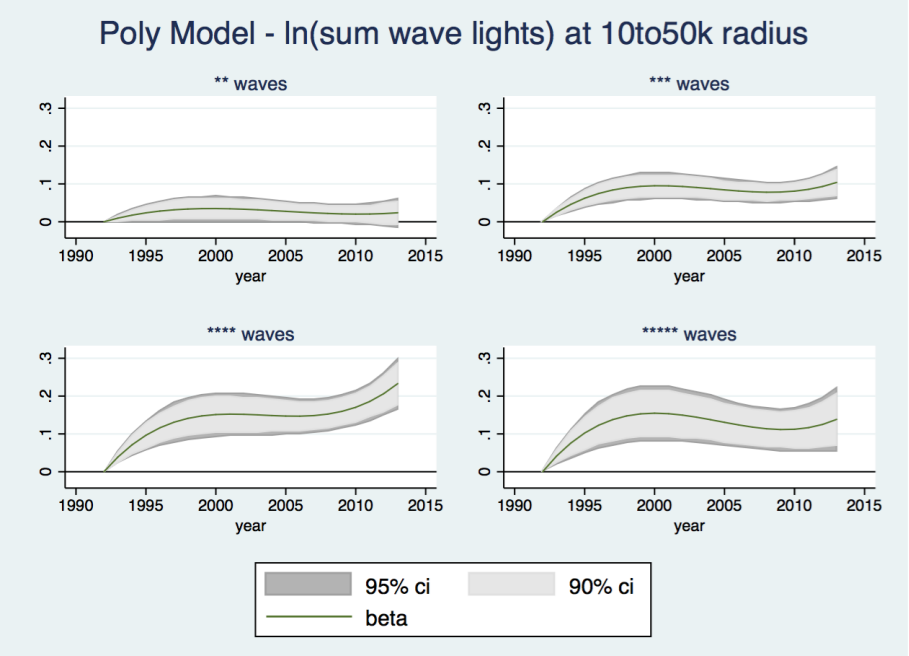
C.3.1 Surfing and economic activity

Figure C.5: Polynomial regressions

(a) 5 to 10km

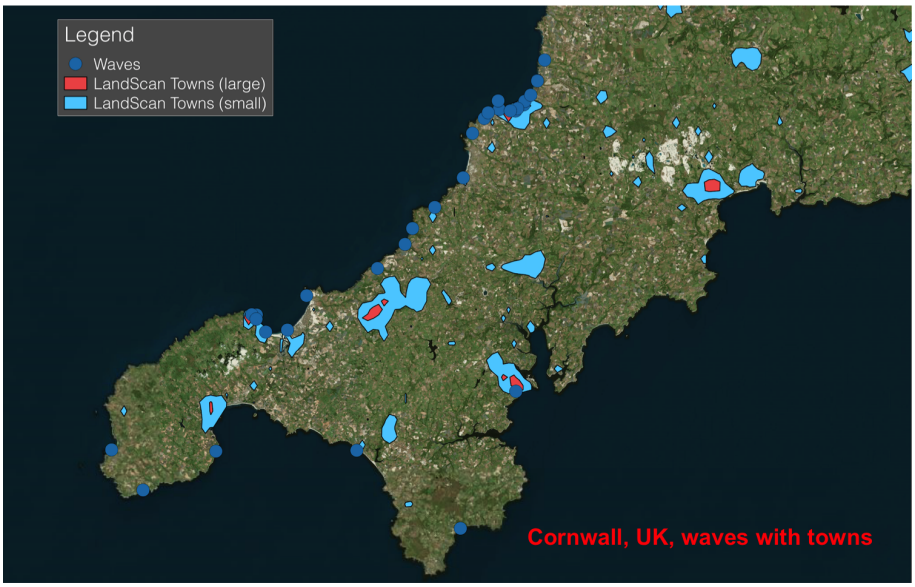


(b) 10 to 50km



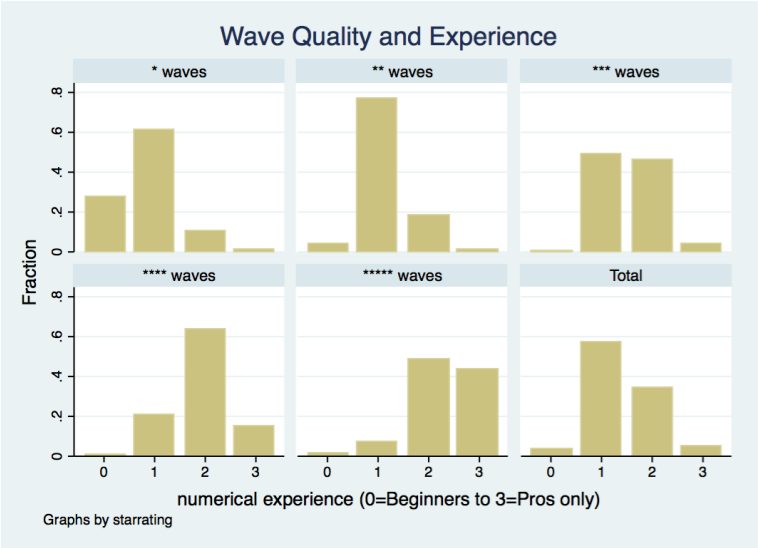
Polynomial model of the difference in difference of light levels from each break, using 1-star breaks as the baseline. Includes year and break fixed effects, and Conley standard errors.

Figure C.6: Endogenous town location - Cornwall, UK



Example of endogenously defining the location of towns by their population density, Cornwall, UK.

Figure C.7: Wave quality and experience



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

C.3.2 Ease of doing business and political stability categories

Table C.4: Surf breaks by World Bank Doing Business Index

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

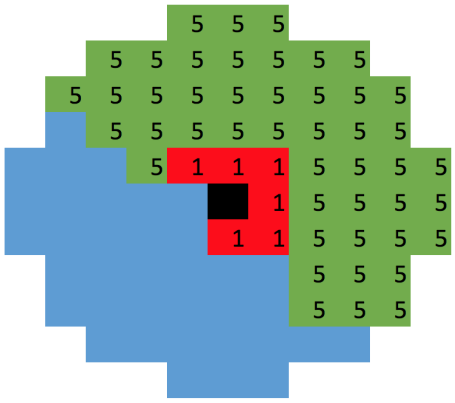
Table C.5: Surf breaks by World Bank Political Stability index

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

C.4 Experiment I: Robustness

C.4.1 Pixel-level analysis

Figure C.8: Pixel level, 1km and 5km



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

The analysis in the main text uses the area within 5km of a surf break as the unit of analysis. These areas may overlap. Here we use individual pixels as the unit of analysis, which are "treated" by surf breaks to give an estimate of the marginal contribution of each break quality to illumination. Every pixel within 5km of at least one break is included in our sample, excluding those wholly over water, which increases the sample size considerably (see Figure C.8). Each pixel is treated in two ways using the specification,

$$\ln(Y_{i,2013}) - \ln(Y_{i,1992}) = \alpha + \beta Q_i + W_i + Z_t + \epsilon_{i,t} \tag{C.2}$$

where $Y_{i.t}$ is light intensity in pixel i at time t, $Q_i \in [M_i, N_{i,k}]$, M_i is the maximum quality of breaks within 5km of pixel i (treated as both a continuous and a categorical variable), $N_{i,k}$ is the number of breaks of quality $k \in [1, 5]$ within 5km of pixel i, W_i is country fixed effects, and Z_t is year fixed effects. The results are given in Table C.6 and discussed in Section 3.3.3.3.

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

$$\ln(Y_{i,t}) = \alpha + \sum_{t=1992}^{2013} \beta_t T_t Q_i + F_i + T_t + \epsilon_{i,t} \tag{C.3}$$

where $Y_{i.t}$ is light intensity in pixel i at time t, T_t is a year dummy, $Q_i \in [M_i, N_{i,k}]$ is defined above, and F_i are pixel fixed effects. Standard errors are clustered at the break and country level. The control group is 1-star breaks. The results are given in Figure C.9 and discussed in Section 3.3.3.3.

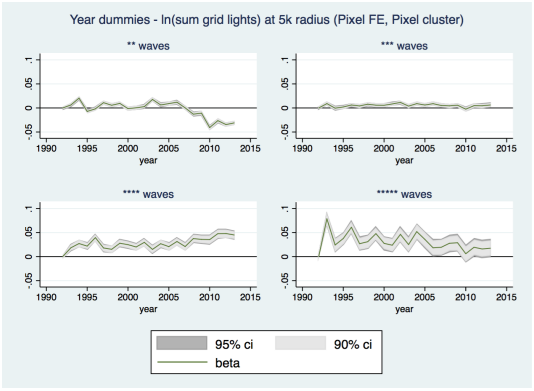
To control for possible effects of convergence between pixels with low and high we also control for the initial level of lights using the specification,

$$\Delta \ln(Y_{i,t}) = \alpha + \beta Q_i + \ln(Y_{i,1992}) + W_i + Z_t + \epsilon_{i,t} \tag{C.4}$$

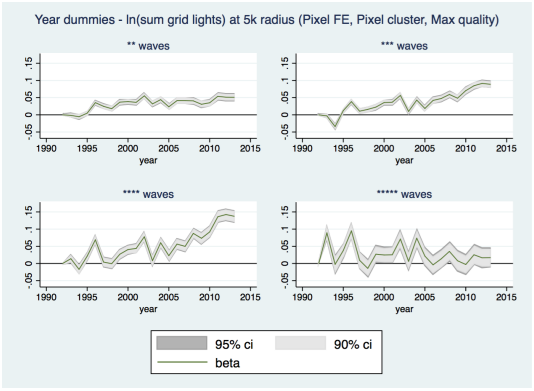
where $\Delta \ln(Y_{i,t})$ is the annual change in log lights (so $\ln(Y_{i,1992})$ doesn't appear on the RHS in all observations), and the other variables are defined above. These results are given in Table C.10 and discussed in Section 3.3.3.3.

Figure C.9: Pixel regressions

(a) Number of breaks of quality k within 5km



(b) Maximum break quality within 5km.



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

Table C.6: Pixel regressions, with controls

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

Robust p-values in parentheses

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

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

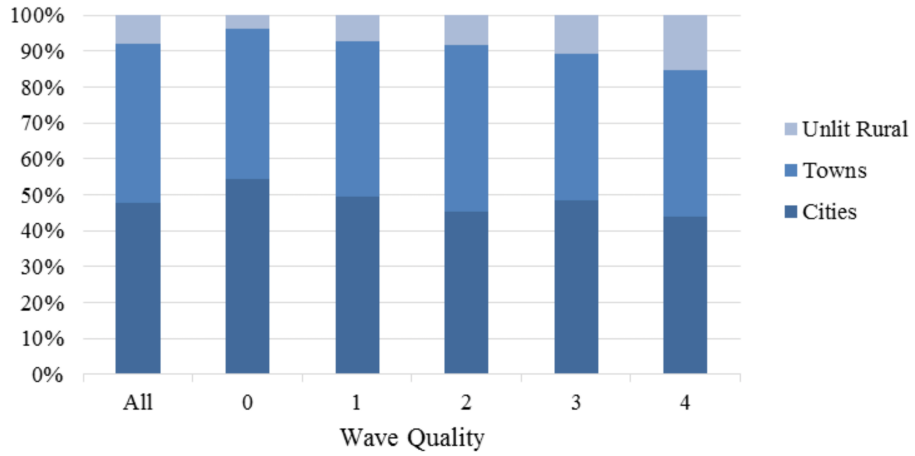
Table C.7: Breakdown of surf breaks by type

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

C.4.2 Omitted Geographic Characteristics

C.4.3 Selection

Figure C.10: Wave location breakdown



Breakdown of the 5km surrounding breaks of each quality

Table C.8: Growth regressions, without Australia or US

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

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

The effect of break quality on the change in ln(lights) within 5km of each break, from 1992-2013. Columns (1)-(4) exclude reefs, pointbreaks and rivermouths, and columns (5)-(8) exclude Australia and the USA.

Table C.9: Growth regressions, with 2-star base

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

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

The effect of break quality on the change in ln(lights) within 5km of each break, from 1992-2013, using 2-star breaks as the baseline.

Table C.10: Pixel level growth regressions, with initial light controls

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

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

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

Table C.11: Discovery regressions, diff-in-diff

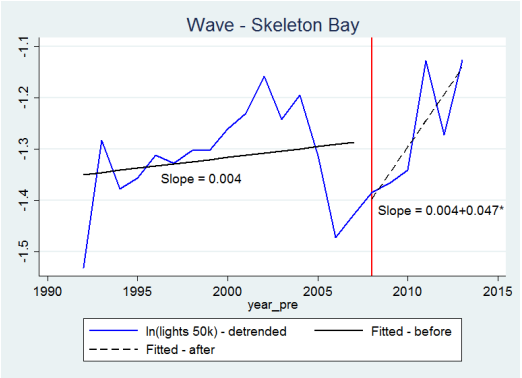
Detail: Discoveries and Disappearances		
	(1)	(2)
Year trend: post-event	0.0222***	-0.00495
	(0.003)	(0.655)
Constant: post-event, interacted with:		
Skeleton Bay	-0.0660***	-0.172***
	(0.001)	(4.74e-06)
Jardim do Mar	0.0504***	0.0839***
	(0.000)	(8.19e-05)
Mundaka	0.120***	0
	(1.05e-05)	
La Jolla	-0.143***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

Detail of Table 3.8. Difference-in-difference in average log lights around break discoveries and disappearances on (1) actual discovery dates, and (2) discovery dates drawn from a uniform U(1992,2013) distribution.

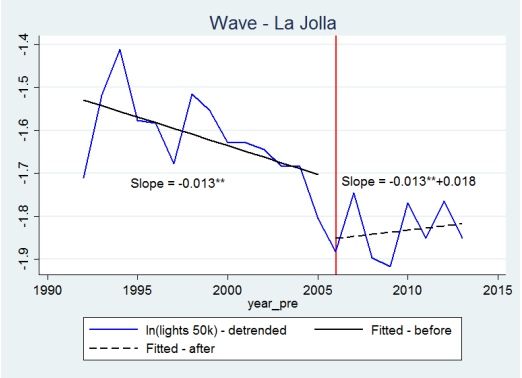
C.5 Experiment II: Additional Results

Figure C.11: Pixel regressions - discoveries

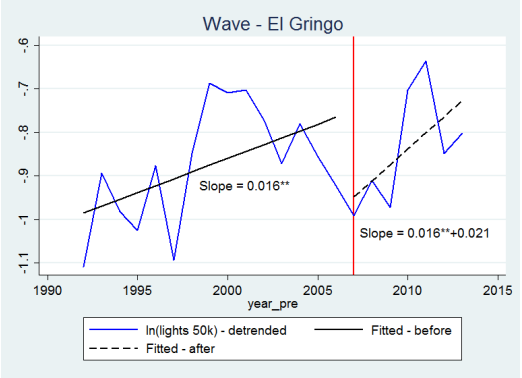
(a)



(b)



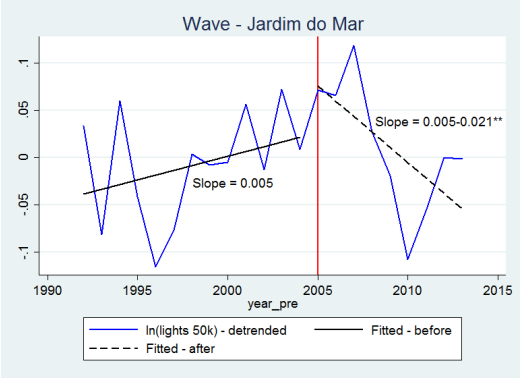
(c)



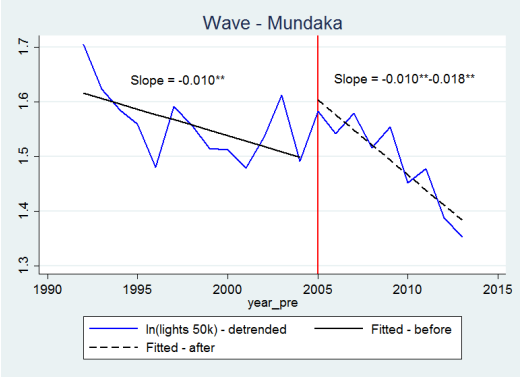
Three surf break discoveries (red) and their effect on the difference between log illumination in the 50km surrounding each break and the mean illumination in the 50km surrounding all surf breaks globally. Trend lines and their slopes are included before (solid) and after (dotted) discovery.

Figure C.12: Pixel regressions - disappearances

(a)



(b)



Two surf break disappearances (red) and their effect on the difference between log illumination in the 50km surrounding each break and the mean illumination in the 50km surrounding all surf breaks globally. Trend lines and their slopes are included before (solid) and after (dotted) disappearance.

Table C.12: Cold surf breaks

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

List of countries with surf breaks above +55 or below -55 degrees latitude.

Table C.13: Discovery of cold surf breaks

Alternative specification: Cold Water Breaks				
	(1)	(2)	(3)	(4)
	Main	Main	Robust	Robust
Year trend: post-event	0.0285***	0.0285**	0.00120	0.00120
	(2.95e-06)	(0.0139)	(0.931)	(0.952)
Constant: post-event	0.103***	0.103	0.00492	0.00492
	(0.000829)	(0.200)	(0.720)	(0.820)
Year trend: pre-event	-0.0160***	-0.0160**	-0.192***	-0.192*
	(1.92e-06)	(0.0283)	(0.000137)	(0.0695)
Constant: pre-event	31.51***	31.51**	-2.673	-2.673
	(2.48e-06)	(0.0301)	(0.924)	(0.946)
Observations	1,804	1,804	1,804	1,804
R-squared	0.032	0.032	0.038	0.038
Number of wid	83	83	83	83
Sample	Lat$\notin [-55, 55]$	Lat$\notin [-55, 55]$	Lat$\notin [-55, 55]$	Lat$\notin [-55, 55]$
FE	Break	Break	Break	Break
SE Cluster	Break	Zone	Break	Zone
Robust p-values in parentheses				
*** p<0.01, ** p<0.05, * p<0.1				

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

Appendix D

Sovereign debt dynamics

D.1 Technical appendix

D.1.1 Ramsey representation

D.1.1.1 Model setup

The model described in section 4.3 can be expressed as a Ramsey policy problem by setting the re-entry probability to 1. The agent maximises the following utility function:

$$\max \mathbb{E}_t \sum_{t=0}^{\infty} \beta^t u(c_t) \tag{D.1}$$

Domestic output is subject to an exogenous endowment shock, $y_t > 0$. In addition to this endowment, the country receives income from selling its annual production of oil, $\bar{s}$, abroad at a stochastic world oil price, $p_t > 0$. The total stochastic income stream in any given period is given by:

$$x_t = y_t + p_t \bar{s} \tag{D.2}$$

The government can save by accumulating safe (risk free) foreign assets, $f_t \geq 0$. Foreign bonds yield a risk-free return of r in each period, so the price of foreign assets is given by $\frac{1}{1+r}$.

The government can also issue bonds, $d_t \geq 0$, where a positive amount denotes positive

debt or negative assets. These are issued at time t and promise to pay d_t units of consumption in period $t+1$. The government can choose to default on debt in any given period subject to some fixed cost, $\lambda(y_t, p_t)$ which is a function of the output state, $y = \left\{y^1, ..., y^{ny}\right\}$, and oil price state, $p = \left\{p^1, ..., p^{np}\right\}$. We use $\delta = \{(y_t, p_t, d_{t-1}) \in [0,1]\mathbb{R}$ to denote an indicator for default where $\delta_{t-1} = 1$ signifies default and $\delta_{t-1} = 0$ signifies no default. The default decision is a function of the output and oil price states as different realisations of these states will result in different optimal default decisions. It is also a function of the level of inherited debt from the last period, d_{t-1}. This is important for two reasons. The higher is indebtedness the more incentive the government has to default and so avoid repayment, but this default cost is a function of output and the oil price. Finally, we impose that $\delta(\cdot, \cdot, d_{t-1}) = 0$ if $d_{t-1} \leq 0$. That is, the government cannot default if it has positive assets (or negative debt).

At the beginning of the period, the government's total wealth, ω_t is given by:

$$\omega_t = y_t + p_t\bar{s} + f_{t-1} - d_{t-1}(1 - \delta_{t-1}(y_t, p_t, d_{t-1})) - \lambda(y_t, p_t)\delta_{t-1}(y_t, p_t, d_{t-1}) \tag{D.3}$$

The resources available in the economy are on the right hand side of equation D.3 and are made up of income, y_t, oil revenues, $p_t\bar{s}$, foreign assets, f_t, net of the obligations on domestic debt repayment from last period, d_{t-1}.

The final two terms of the right hand side denote the decision by the government to default on existing debt through the choice of δ_{t-1} and the cost under default through λ respectively. When the government chooses to service the debt stock from the previous period (that is, no default) then $\delta_{t-1} = 0$, and these two terms collapse to $-d_{t-1}$. This means that wealth is simply given by: $\omega_t = y_t + p_t\bar{s} + f_{t-1} - d_{t-1}$. When the government chooses to default on the previous period's debt, $\delta_{t-1} = 1$, then these terms collapse to $-\lambda$, and wealth is given by: $\omega_t = y_t + p_t\bar{s} + f_{t-1} - \lambda$.

The economy's total budget constraint is given by:

$$\omega_t + \frac{d_t}{1 + R_t(y_t, p_t, \delta_t)} = c_t + \frac{f_t}{1 + r} \tag{D.4}$$

where c_t denotes consumption, f_t denotes new foreign asset accumulation and d_t denotes new debt issuance.

The interest rate on government debt is given by R_t and depends on the default profile, δ, chosen by the government and on the current output and oil price states, y_t and p_t, as these will affect the likelihood of entering different states in $t + 1$. The government takes into account in its optimisation problem the effect that future default decisions have on the bond price today. This is described in the function $R(\cdot, \cdot, \cdot)$ in the following bond pricing formula which, through a zero profits assumption, equates the return on domestic and foreign bonds:

$$q(y_t, p_t, \delta_t) = \frac{1}{1 + R_t(y_t, p_t, \delta_t)} = \frac{1}{1 + r} \sum_{i=1}^{ny} \sum_{j=1}^{np} \big(1 - \delta_t(y^{ny}, p^{np}, d^{nd})\big).\pi\big(y^i | y_t, p^j | p_t\big) \tag{D.5}$$

where r denotes the risk free world interest rate and $q(\cdot, \cdot, \cdot)$ denotes the prevailing market bond price. The indices ny and np denote the possible realisations of output and oil prices in the future respectively, whilst the nd index denotes the debt choice corresponding to these possible states.

The Ramsey problem characterising optimal policy is:

$$\max \mathbb{E}_t \sum_{t=0}^{\infty} \beta^t u(c_t) \tag{D.6}$$

s.t.:

$$c_t = \omega_t + \frac{d_t}{1 + R_t(y_t, p_t, \delta_t)} - \frac{f_t}{1 + r} \tag{D.7}$$

$$\omega_t = y_t + p_t \bar{s} + f_{t-1} - d_{t-1}(1 - \delta_{t-1}(y_t, p_t, d_{t-1})) - \lambda(y_t, p_t)\delta_{t-1}(y_t, p_t, d_{t-1}) \tag{D.8}$$

$$d_t \leq 0 \text{ and } f_t \geq 0$$

$$d_0, y_0, p_0 : \text{ given}$$

D.1.1.2 Solution

We formulate the Lagrangian L, letting ξ_t denote the multiplier on the budget constraint from combining D.7 and D.8 in period t:

$$\begin{aligned} L = \mathbb{E}_t \Bigg[& \sum_{t=0}^{\infty} \beta^t u(c_t) \\ & + \beta^t \xi_t \Bigg(y_t + p_t \bar{s} + \frac{d_t}{1 + R_t(y_t, p_t, \delta_t)} - \frac{f_t}{1+r} - c_t + f_{t-1} \\ & - d_{t-1}\big(1 - \delta_{t-1}(y_t, p_t, d_{t-1})\big) - \lambda(y_t, p_t)\delta_{t-1}(y_t, p_t, d_{t-1}) \Bigg) \Bigg] \end{aligned} \tag{D.9}$$

Differentiating the Lagrangian with respect to the choice variables $(c_t, f_t, d_t, \delta_t)$ gives:

$$\begin{aligned} c_t : \quad & u'(c_t) - \xi_t = 0 \\ f_t : \; & -\frac{\xi}{1+r} + \beta \xi_{t+1} = 0 \\ d_t : \; & \xi_t \Big[\frac{-d_t \delta_t' R_t + R_t + 1}{(1+R_t)^2} \Big] - \beta \xi_{t+1} \big(\delta_t'(\lambda - d_t) - \delta_t + 1 \big) = 0 \\ \delta_t : \; & -\xi_t \frac{d_t R_t'}{(1+R_t)^2} + \beta \xi_{t+1}(1 - \lambda) = 0 \end{aligned} \tag{D.10}$$

Using the FOC for consumption to substitute out for ξ_t and ξ_{t+1} in the remaining FOCs for assets, debt and default yields the following three conditions. The first is the standard Euler Equation on risk free assets given in equation D.11. It states that the growth in the marginal utility of consumption is determined by the risk-free interest rate, r, and the discount rate, β. When $\beta(1+r) = 1$, the ratio of marginal utilities is constant and so the agent expects pure consumption smoothing.

$$\begin{aligned} u'(c_t)\frac{1}{1+r} &= \beta u'(c_{t+1}) \\ \frac{u'(c_t)}{u'(c_{t+1})} &= \beta(1+r) \end{aligned} \tag{D.11}$$

The second is the Euler equation for debt which is the same as the standard Eurler equation except for the terms in square brackets on the left and right hand sides of equation D.12. The feedback of debt to interest rates in the model drives a wedge between the standard Euler equation and the equation for optimal debt which is governed by four things: the level of debt, d_t, the default decision, δ_t and its marginal effect, δ'_t, and the cost of default, λ.

$$\begin{aligned} u'(c_t)\left[\frac{-d_t\delta'_tR_t + R_t + 1}{(1+R_t)^2}\right] &= \beta u'(c_{t+1})\left[\delta'_t(\lambda - d_t) - \delta_t + 1\right] \\ \frac{u'(c_t)}{u'(c_{t+1})}\frac{1}{1+R_t}\left[\frac{-d_t\delta'_tR_t + R_t + 1}{1+R_t}\right] &= \beta\left[\delta'_t(\lambda - d_t) - \delta_t + 1\right] \\ \beta(1+r)\frac{1}{1+R_t}\left[\frac{-d_t\delta'_tR_t + R_t + 1}{1+R_t}\right] &= \beta\left[\delta'_t(\lambda - d_t) - \delta_t + 1\right] \\ \frac{1+r}{1+R_t}\left[\frac{-d_t\delta'_tR_t + R_t + 1}{1+R_t}\right] &= \left[\delta'_t(\lambda - d_t) - \delta_t + 1\right] \end{aligned} \tag{D.12}$$

When $\left[\frac{-d_t\delta'_tR_t+R_t+1}{1+R_t}\right] = [\delta'_t(\lambda - d_t) - \delta_t + 1]$ then $\frac{1+r}{1+R_t} = 1$, and the interest rate on sovereign debt is equal to the risk-free interest rate, $R_t = r$. In this case, the default model converges to an RBC model without default. One case where this occurs is when there is no option of default for the government, such that $\delta = \delta' = 0$ and $\lambda = 0$ due to no default cost. In this case the terms in square brackets collapse to unity, and equations D.11 and D.12 are identical. This gives rise to the first proof:

Proof 1 *The default model described in the system of equations D.10 collapses to a standard Ramsey model with no financial market frictions when* $\delta = 0$ *and* $\lambda = 0$.

We can also solve the above condition for d_t as follows:

$$\begin{aligned} \frac{-d_t\delta_t' R_t + R_t + 1}{1 + R_t} &= \delta_t'(\lambda - d_t) - \delta_t + 1 \\ d_t &= (\lambda - 1)(R_t + 1) \end{aligned} \tag{D.13}$$

Given that the fixed default cost, λ, and the risk-free interest rate, r, are exogenous parameters, there is no reason for the expression in equation D.13 to hold. When it does however, we define this as the level of debt that yields a perfect Ramsey model solution to the model. This gives rise to the second proof:

Proof 2 *There is a single parameterisation of the fixed default cost, λ, and risk-free interest rate, R, such that $d_t = (\lambda - 1)(R_t + 1)$ holds and the optimal debt choice, $d_t = d_t^{Ramsey}$, is such that the default model described in the system of equations D.10 approximates a standard Ramsey model with no financial market frictions.*

The third condition, which is obtained from substituting out for ξ_t and ξ_{t+1} in the final FOC of D.10, is the Euler equation for default given in equation D.14. Again, this is identical to the standard Euler equation but for the terms in square brackets. A wedge is driven between the standard Euler equation and the equation for optimal default, which changes with respect to the level of debt, d_t, the marginal effect on the interest rate of debt, R_t', and the cost of default, λ.

$$\begin{aligned} u'(c_t)\left[\frac{d_t R_t'}{(1+R_t)^2}\right] &= \beta u'(c_{t+1})\left[1-\lambda\right] \\ \frac{u'(c_t)}{u'(c_{t+1})}\frac{1}{1+R_t}\left[\frac{d_t R_t'}{1+R_t}\right] &= \beta\left[1-\lambda\right] \\ \beta(1+r)\frac{1}{1+R_t}\left[\frac{d_t R_t'}{1+R_t}\right] &= \beta\left[1-\lambda\right] \\ \frac{1+r}{1+R_t}\left[\frac{d_t R_t'}{1+R_t}\right] &= \left[1-\lambda\right] \end{aligned} \tag{D.14}$$

When $\left[\frac{d_t R_t'}{1+R_t}\right] = [1-\lambda]$, then again the default model converges to an RBC model without default. When default is not an option for the sovereign $\lambda = 0$ and this expression collapses to $d_t = \frac{1+R_t}{R_t'}$, then the expression in D.14 collapses to the simple Ramsey Euler

equation with risk free assets. Note that the function $R'_t \geq 0$ is the first derivative of the interest rate charged on sovereign debt with respect to the level of debt and is assumed to be positive - higher debt leads to a higher interest rate. This expression shows that when default is not available, the sovereign chooses to borrow less, the more responsive it's interest rate is to the level of debt. This is the first lemma:

Lemma 1 *Without default as an option for the sovereign, the level of debt is decreasing in the elasticity of the sovereign's interest rate with respect to indebtedness according to the following expression:* $d_t = \frac{1+R_t}{R'_t} \mid R'_t \geq 0$.

We can also use the expression in D.14 to derive an expression for the optimal level of debt when default is available to the sovereign. By simply re-arranging we get:

$$\begin{aligned} u'(c_t)\frac{1}{1+R_t}\left[\frac{d_t R'_t}{1+R_t}\right] &= \beta u'(c_{t+1})\left[1-\lambda\right] \\ \frac{u'(c_t)}{u'(c_{t+1})}\left[\frac{d_t R'_t}{(1+R_t)(1+\lambda)}\right] &= \beta(1+R_t) \end{aligned} \tag{D.15}$$

When $\beta(1+r) = 1$ we get pure consumption smoothing only when the term in square brackets, $\left[\frac{d_t R'_t}{(1+R_t)(1+\lambda)}\right]$, is equal to unity. If this term is less than unity then the ratio of marginal utilities is required to be greater than unity. This means that the marginal utility of consumption is higher today that it is tomorrow and thus consumption is lower today that is is tomorrow. This requires that $d_t \leq \frac{(1+R_t)(1+\lambda)}{R'_t}$ or, in other words, the elasticity of the sovereign's interest rate is sufficiently low relative to the prevailing marked rate and the cost of default. This makes sense as when the sovereign's interest rate does not responds strongly to the level of debt (low debt elasticity of the interest rate) then the sovereign will increase indebtedness to bring forward consumption. This gives rise to the second lemma:

Lemma 2 *When sovereign default is possible the degree of consumption tilting depends on the level of impatience,* β*, the prevailing market interest rate,* R_t*, and its elasticity with respect to the level of debt,* R'_t*, relative to the cost of default,* λ*. Under the usual consumption smoothing condition of* $\beta = (1+r)$*, consumption will be brought forward, via debt issuance, when the condition* $d_t \leq \frac{(1+R_t)(1+\lambda)}{R'_t}$ *holds.*

D.1.1.3 Steady state

We now solve for the steady-state of the model. In the steady state, consumption, debt and assets are all constant, and output and the oil price are untty. We can represent the full system of equations which include the three Euler equations, D.11, D.12, D.14, and the combined budget constraint from D.7 and D.8 as follows, where we have substituted out for $1 = \beta(1 + r)$ in equations D.17 and D.18:

$$\begin{aligned} \frac{u'(c_t)}{u'(c_{t+1})} &= \beta(1+r) \\ \frac{u'(c)}{u'(c)} &= \beta(1+r) \\ 1 &= \beta(1+r) \end{aligned} \tag{D.16}$$

$$\begin{aligned} u'(c_t)\frac{1}{1+R_t}\left[\frac{-d_t\delta_t' R_t + R_t + 1}{1+r}\right] &= \beta u'(c_{t+1})\left[\delta_t'(\lambda - d_t) - \delta_t + 1\right] \\ u'(c)\frac{1}{1+R}\left[\frac{-d\delta' R + R + 1}{1+r}\right] &= \beta u'(c)\left[\delta'(\lambda - d) - \delta + 1\right] \\ \left[\frac{-d\delta' R + R + 1}{1+R}\right] &= \left[\delta'(\lambda - d) - \delta + 1\right] \end{aligned} \tag{D.17}$$

$$\begin{aligned} u'(c_t)\frac{1}{1+R_t}\left[\frac{d_t R_t'}{1+R_t}\right] &= \beta u'(c_{t+1})\left[1 - \lambda\right] \\ u'(c)\frac{1}{1+R}\left[\frac{dR'}{1+R}\right] &= \beta u'(c)\left[1 - \lambda\right] \\ \left[\frac{dR'}{1+R}\right] &= \left[1 - \lambda\right] \end{aligned} \tag{D.18}$$

$$\begin{aligned} c_t &= y_t + p_t\bar{s} + f_{t-1} - d_{t-1}\left(1 - \delta_{t-1}\right) - \lambda\delta_{t-1} + \frac{d_t}{1+R_t} - \frac{f_t}{1+r} \\ c &= \theta + f - d\left(1 - \delta\right) - \lambda\delta + \frac{d}{1+R} - \frac{f}{1+r} \end{aligned} \tag{D.19}$$

where variables without time subscripts denote steady-state values and $\theta = y + p\bar{s}$ denotes the steady state level of income which depends on the share of oil in the domestic

economy, $\bar{s}$.

D.1.1.4 Optimal debt choice

We can now solve for the optimal choice of debt in the steady-state by combining the three Euler equations in D.16, D.17 and D.18 with the budget constraint in D.19.

$$\begin{aligned} 1 &= \beta(1+r) \\ \left[\frac{-d\delta' R + R + 1}{1+R}\right] &= \left[\delta'(\lambda - d) - \delta + 1\right] \\ \left[\frac{dR'}{1+R}\right] &= [1-\lambda] \\ c &= \theta - d\,(1-\delta) - \lambda\delta + \frac{d}{1+R} \end{aligned} \tag{D.20}$$

Substituting out for $1/1+R$ from the 2nd equation of D.20 into the 4th equation, and solving for optimal debt in the steady-state, d^*, yields:

$$d^* = \frac{(R+1)(c+\theta+\delta\lambda)}{(\delta-1)R} \tag{D.21}$$

Noting that:

$$\begin{aligned} \frac{\partial d^*}{\partial c} &= \frac{R+1}{(\delta-1)R} \\ \frac{\partial d^*}{\partial R} &= \frac{c+\delta\lambda+\theta}{(\delta-1)R^2} \end{aligned} \tag{D.22}$$

where $(R+1)/R > 0$ and $\delta - 1 < 0$, leads to the third lemma

Lemma 3 *The optimal level of debt for a resource rich sovereign is a non-linear function of the demands on revenue, the prevailing interest rate and the government's likelihood of default. Optimal debt is increasing in the the steady state level of consumption, c, but at a rate that is decreasing in the default decision, δ. It is also decreasing in the interest rate, R, but at an increasing rate*

D.1.1.5 Pricing sovereign debt

Dividing equation D.17 by D.18 gives:

$$\frac{-d\delta' R + R + 1}{dR'} = \frac{\delta'(\lambda - d) - \delta + 1}{1 - \lambda} \tag{D.23}$$

Now substituting in for the bond price, $q = \frac{1}{1+R}$, from D.5 gives:

$$q = \frac{1 - \lambda}{R\left(\delta'(\lambda - d) - \delta + 1\right)} + \frac{1}{d\delta' R} \tag{D.24}$$

When default is not possible for the sovereign, $\delta = \delta' = 0$ and $\lambda = 0$, so the bond pricing expression collapses to $q = \frac{1}{R}$. This simple expression states that the steady state bond price of sovereign debt is inversely related to the level of debt, d and the interest rate facing the sovereign, R. When default is possible, this inverse relationship is augmented by the cost of default, λ, and the marginal effect of the level of debt on the optimal default decision by the sovereign, δ'. The more sensitive the default decision of the sovereign is with respect the level of debt (that is, the larger is δ'), the lower the price investors are willing to pay to hold this debt (and the higher the effective interest rate). This is the fourth lemma:

Lemma 4 *Without default, the the steady state bond price of sovereign debt, q, is inversely related to the level of debt, d, and the interest rate facing the sovereign, R, according to the expression $q = \frac{1}{R}$. When default is possible, this inverse relationship is augmented by the sensitivity of the sovereign's default decision with respect to the level of debt, δ'; the more sensitive this choice is, the less investors are willing to pay to hold the sovereign's debt. The bond pricing expression in the steady state is given by $q = \frac{1-\lambda}{R(\delta'(\lambda-d)-\delta+1)} + \frac{1}{d\delta' R}$.*

D.1.2 Bivariate Markov shocks

I follow Tauchen (1986) in the exposition of a discretised state-space representation of a Markov shock process. Instead of using a single variable process however, I assume a

bivariate auto-regressive shock process consisting of two correlated variables, x_t and y_t. Consider the following shock process in matrix form:

$$\begin{bmatrix} x_{t+1} \\ y_{t+1} \end{bmatrix} = \begin{bmatrix} \rho_x & \rho_{xy} \\ \rho_{yx} & \rho_y \end{bmatrix} \begin{bmatrix} x_t \\ y_t \end{bmatrix} + \begin{bmatrix} \epsilon_{x,t+1} \\ \epsilon_{y,t+1} \end{bmatrix} \tag{D.25}$$

$$\boldsymbol{z}_{t+1} = \boldsymbol{\rho} z_t + \boldsymbol{\epsilon}_{t+1}$$

where $\boldsymbol{\epsilon}_{t+1} \sim N(\mu, \Sigma)$, $\boldsymbol{\mu} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$ and $\boldsymbol{\Sigma} = \mathbb{E}(\boldsymbol{\epsilon}\boldsymbol{\epsilon}') = \begin{bmatrix} \sigma_x & \sigma_{xy} \\ \sigma_{yx} & \sigma_y \end{bmatrix}$. Then we can express the expected value of $\boldsymbol{z}_{t+1}$ and it's variance as:

$$\mathbb{E}(\boldsymbol{z}_{t+1}) = 0 \tag{D.26}$$

$$\text{Var}(\boldsymbol{z}_{t+1}) = \boldsymbol{\Omega} = \begin{bmatrix} \omega_x & \omega_{xy} \\ \omega_{yx} & \omega_y \end{bmatrix} = \boldsymbol{\Sigma}\left[\mathrm{I} - \boldsymbol{\rho}\right]^{-1} \tag{D.27}$$

We define a state space grid for the correlated variables as:

$$\begin{bmatrix} x_1 \\ y_1 \end{bmatrix} < \begin{bmatrix} x_2 \\ y_2 \end{bmatrix} < ... < \begin{bmatrix} x_m \\ y_n \end{bmatrix}$$

where

$$\begin{aligned} \boldsymbol{z}_1 &= \begin{bmatrix} x_1 \\ y_1 \end{bmatrix} = \begin{bmatrix} -r\omega_x \\ -r\omega_y \end{bmatrix}, \\ \boldsymbol{z}_{mn} &= \begin{bmatrix} x_m \\ y_n \end{bmatrix} = \begin{bmatrix} r\omega_x \\ r\omega_y \end{bmatrix} \text{ and} \\ \boldsymbol{d} &= \begin{bmatrix} d_x \\ d_y \end{bmatrix} = \begin{bmatrix} x_i - x_{i-1} \\ y_k - y_{k-1} \end{bmatrix} \end{aligned} \tag{D.28}$$

where r is the state-space factor, which we set to 3, and m and n denote the respective

number of state spaces. Since $\Pr(\epsilon_{t+1} \leq \boldsymbol{u}) = F\left([\boldsymbol{u} - 0]\,\boldsymbol{\Sigma}^{-1}\right)$, we can define the transition probability matrix $\boldsymbol{\Pi}_{ij,kl}$ as:

$$\begin{aligned}\boldsymbol{\Pi}_{ij,kl} &= \Pr\{\boldsymbol{z}_{jl} - \boldsymbol{d} \setminus 2 \leq \boldsymbol{\rho}\boldsymbol{z}_{ik} + \boldsymbol{\epsilon}_{t+1} \leq \boldsymbol{z}_{jl} + \boldsymbol{d} \setminus 2\} \\ &= F\left([\boldsymbol{z}_{jl} - \boldsymbol{\rho}\boldsymbol{z}_{ik} + \boldsymbol{d} \setminus 2]\,\boldsymbol{\Sigma}^{-1}\right) - F\left([\boldsymbol{z}_{jl} - \boldsymbol{\rho}\boldsymbol{z}_{ik} - \boldsymbol{d} \setminus 2]\,\boldsymbol{\Sigma}^{-1}\right)\end{aligned} \tag{D.29}$$

where $\boldsymbol{\Pi}$ is a matrix of dimensions $(m \times n \times m \times n)$, and $F(\cdot)$ denotes the CDF of the bivariate normal distribution. Subscript ij denotes the probability of moving from the ith position to the jth position of variable x, and kl denotes the probability of moving from the kth position to the lth position of variable y. The boundary transition probabilities are given by:

$$\boldsymbol{\Pi}_{i1,k1} = F\left([\boldsymbol{z}_{11} - \boldsymbol{\rho}\boldsymbol{z}_{ik} + \boldsymbol{d} \setminus 2]\,\boldsymbol{\Sigma}^{-1}\right) \tag{D.30}$$

and

$$\boldsymbol{\Pi}_{i\mathrm{m},k\mathrm{n}} = F\left([\boldsymbol{z}_{mn} - \boldsymbol{\rho}\boldsymbol{z}_{ik} - \boldsymbol{d} \setminus 2]\,\boldsymbol{\Sigma}^{-1}\right) \tag{D.31}$$

I have developed Matlab code that estimates the transition matrix for the bivariate VAR process described here, and another that simulates a Markov chain using this bivariate transition matrix. These codes allow the user to set the deep parameters of the bivariate process as well as the state space for the model.[1]

[1] Please visit `https://thomasmjmcgregor.wordpress.com/data/`

D.2 Appendix

D.2.1 Solution algorithm

The numerical solution method involves the following steps:

1. Choose parameter values for: $(\rho, \theta, \sigma, r, \bar{s}, \psi, \rho_z, \rho_h, \sigma_z^2, \sigma_h^2)$.

2. Discretise the state spaces for y, d, and p. I use 11 output and oil price states, and 125 debt states.

3. Start with a guess of the bond price schedule $q^0(d, y, p) = 1/(1 + r)$ and use this to solve the sovereign's problem using value function iteration, obtaining the optimal policy functions for consumption, asset holdings and default choice.

4. Given these policy functions, compute the probability of default, $\delta(d', y, p)$.

5. Update the bond pricing schedule given this probability of default and use this updated schedule to repeat steps 3-4 until the the convergence criterion has been reached, namely $q^0(d, y, p) - q^1(d, y, p) < \Delta$, where Δ is some very small number.

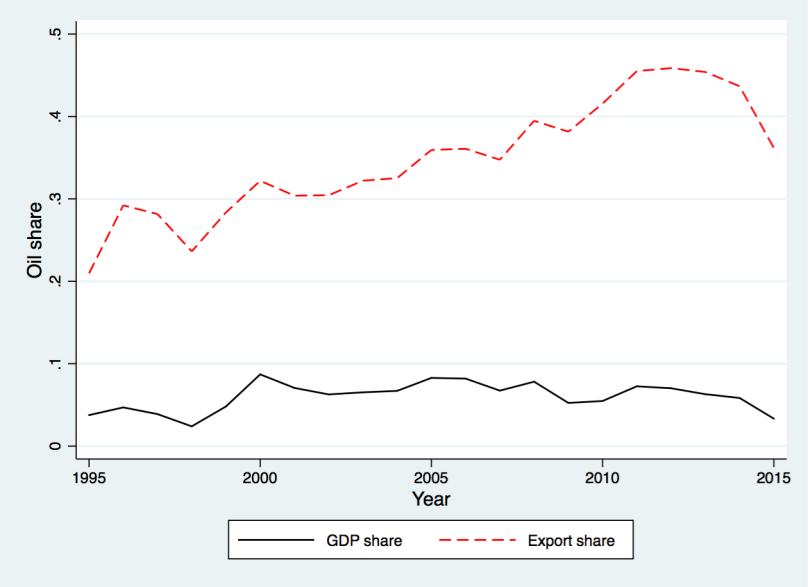
D.2.2 Oil export share calibration

This appendix discusses the performance of the model with a substantially higher oil share parameter, $\bar{s}$. It turns out that the model is fairly sensitive to the calibration of this parameter. In general a higher oil share increases the spread persistence as well as the strength of the negative correlation between interest rate spreads and the oil price. Increasing the oil share parameter does, however, present challenges for the numerical solution algorithm. A higher share, means that the sovereign also has more resources at it's disposal on average. Relying on the same solution algorithm, with a pre-defined debt state space for example, often results in 'non-convergence' problems.

As discussed in section 4.4.2, an alternative way to calibrate the oil share parameter would be to use data on oil export shares. Figure D.1 plots the average share of oil in total GDP and the oil export share in total exports for our four economies, between 1995

and 2015. The export share has been broadly rising, whilst the output share has remained relatively constant.

Figure D.1: Oil and gas shares over time



Here I present model simulation results setting with $\bar{s} = .35$. Table D.1 presents the mean AR(1) coefficients and standard deviations for output, the oil price, consumption and the interest rate spread from simulating the model over 1000 periods. The key difference between these simulations and those in section 4.6.1 is that the spread persistence is increased substantially. Table D.2 presents the mean correlation between spreads and deviations of output and the oil price in the data respectively. The negative correlations are stronger in both cases.

Table D.1: AR(1) model: autocorrelations and standard dev ($\bar{s} = .35$)

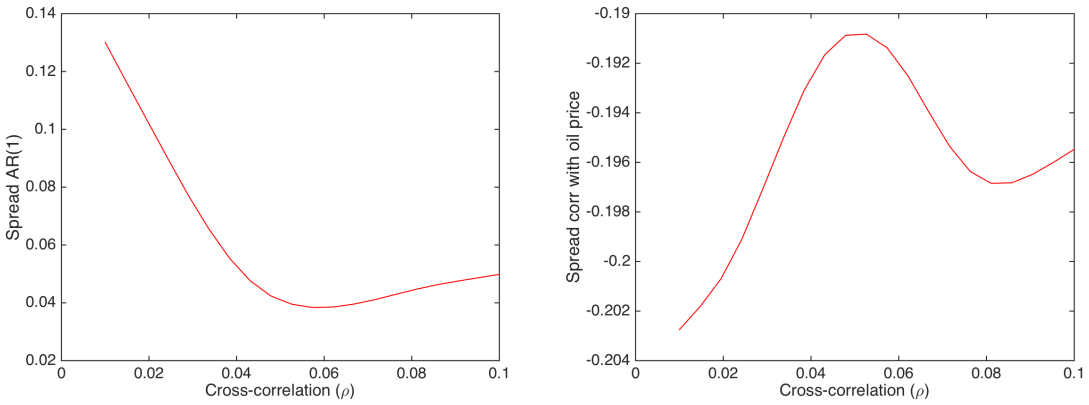
	Data		Model	
	AR(1)	std	AR(1)	std
output (HP filt)	.841	.032	.893	.043
oil price (HP filt)	.725	.198	.657	.195
cons (HP filt)	.898	.026	.977	.107
spread	.645	.022	.304	.010

Table D.2: AR(1) model: cross correlations with real interest rate spread ($\bar{s} = .35$)

	Data corr	Model corr
output dev	-.138	-.148
oil price dev	-.279	-.160

Figures D.2a and D.2b present the AR(1) parameter for the spread and the spread's correlation with oil prices respectively, for different levels of the correlation between output and oil price shocks. The shape of the spread persistence is similar to the model with $\bar{s} = .06$ for strongly correlated shocks, but very different for less correlated shocks. The shape of the cross correlation between spreads and oil prices is very similar to the model with $\bar{s} = .06$, but with a more negative correlation on average.

Figure D.2: Sensitivity to correlation - $\rho_{zh} \in [.01, .10]$ ($\bar{s} = .35$)



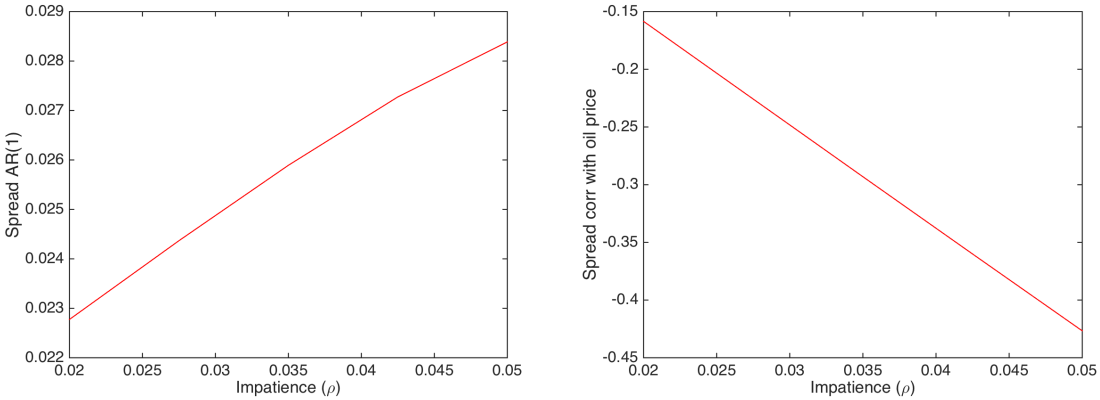
(a) AR(1) parameter

(b) Correlations

Figures D.3a and D.3b present the same figures, but for different levels of impatience. I was forced to limit the solution algorithm to solve the model for only 3 intermediate points on the impatience range due to non-convergence problems. In general the spread persistence responds in a similar way to impatience as in the model with $\bar{s} = .06$, that is, it is increasing with impatience. The main difference between the two models is the increase in persistence on average. The correlation between spreads and the oil price also responds in a similarly

negative way to impatience, albeit with the negative correlation substantially increased on average.

Figure D.3: Sensitivity to impatience, $\rho \in [.01, .05]$ ($\bar{s} = .35$)



(a) AR(1) parameter

(b) Correlations

D.2.3 Model sensitivity

Figure D.4: Output truncation under default

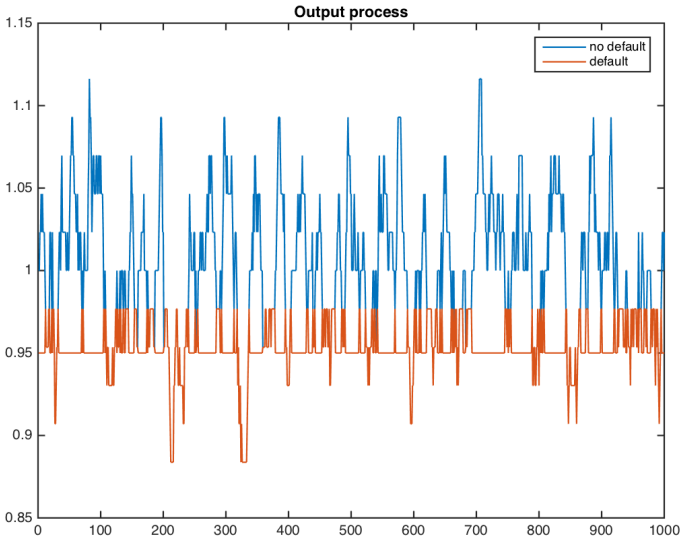


Figure D.5: Bond pricing schedule (no oil)

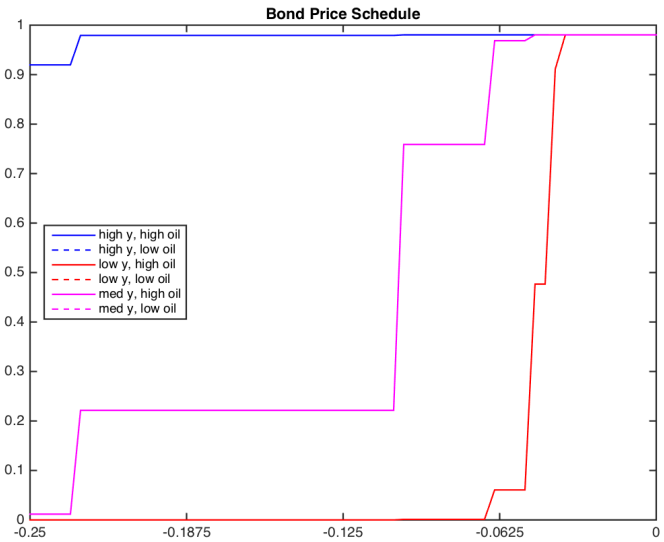


Table D.3: AR(1) model (no oil): autocorrelations and standard dev

	Data		Model	
	AR(1)	std	AR(1)	std
output (HP filt)	.841	.032	.716	.024
oil price (HP filt)	.725	.198	.650	.195
cons (HP filt)	.898	.026	.839	.016
spread	.645	.022	.001	.005

Table D.4: AR(1) model (no oil): cross correlations with real interest rate spread

	Data	Model
	corr	corr
output dev	-.138	-.116
oil price dev	-.279	-.129

Bibliography

Acemoglu, D., Johnson, S. & Robinson, J., 2002. Reversal of Fortune: Geography and Institutions in the Making of the Modern World Income Distribution. *The Quarterly Journal of Economics*, **117**(4), 1231–1294.

Acemoglu, D., Johnson, S., Robinson, J. & Thaicharoen, Y., 2003. Institutional causes, macroeconomic symptoms: volatility, crises and growth. *Journal of monetary economics*, **50**(1), 49–123.

Acemoglu, D. & Robinson, J., 2006. Economic backwardness in political perspective. *American Political Science Review*, **100**(1), 115–131.

Adam, C. S. & Bevan, D. L., 2006. Aid and the supply side: Public investment, export performance, and Dutch disease in low-income countries. *The World Bank Economic Review*, **20**(2), 261–290.

Adam, C. S. & O'Connell, S. A., 2004. Aid versus Trade Revisited: Donor and Recipient Policies in the Presence of Learning-by-Doing. *The Economic Journal*, **114**(492), 150–173.

Adam, K. & Grill, M., 2017. Optimal Sovereign Default. *American Economics Journal: Macroeconomics*, **9**(1), 128–164.

Aghion, P., Bacchetta, P., Ranciere, R. & Rogoff, K., 2009. Exchange rate volatility and productivity growth: The role of financial development. *Journal of monetary economics*, **56**(4), 494–513.

Aguiar, M. & Gopinath, G., 2006. Defaultable debt, interest rates and the current account. *Journal of international Economics*, **69**(1), 64–83.

Akay, A., Martinsson, P., Medhin, H. & Trautmann, S. T., 2012. Attitudes toward uncertainty among the poor: an experiment in rural Ethiopia. *Theory and Decision*, **73**(3), 453–464.

Alesina, A., Michalopoulos, S. & Papaioannou, E., 2016. Ethnic inequality. *Journal of Political Economy*, **124**(2), 428–488.

Allcott, H. & Keniston, D., 2014. Dutch disease or agglomeration? The local economic effects of natural resource booms in modern America. Tech. rep., National Bureau of Economic Research.

Aragón, F. M. & Rud, J. P., 2013. Natural resources and local communities: evidence from a Peruvian gold mine. *American Economic Journal: Economic Policy*, **5**(2), 1–25.

Arellano, C., 2008. Default risk and income fluctuations in emerging economies. *The American Economic Review*, **98**(3), 690–712.

Arezki, M. R., Piotrowski, J. M. & Cherif, R., 2009. Tourism Specialization and Economic Development: Evidence from the UNESCO World Heritage List. Tech. rep.

Arezki, R. & Brückner, M., 2011. Oil rents, corruption, and state stability: Evidence from panel data regressions. *European Economic Review*, **55**(7), 955–963.

Arezki, R. & Brückner, M., 2012. Commodity windfalls, democracy and external debt. *The Economic Journal*, **122**(561), 848–866.

Atkeson, A. & Ogaki, M., 1996. Wealth-varying intertemporal elasticities of substitution: evidence from panel and aggregate data. *Journal of Monetary Economics*, **38**(3), 507–534.

Auty, R. M., 1993. *Sustaining Development in Mineral Economies: The Resource Curse Thesis*. Routledge, London.

Bannon, I. & Collier, P., 2003. *Natural resources and violent conflict: Options and actions*. World Bank publications.

Barro, R. J., 1979. On the determination of the public debt. *Journal of political Economy*, **87**(5, Part 1), 940–971.

Barsky, R. B. & Kilian, L., 2004. Oil and the Macroeconomy since the 1970s. *The Journal of Economic Perspectives*, **18**(4), 115–134.

Bems, R. & de Carvalho Filho, I., 2011. The current account and precautionary savings for exporters of exhaustible resources. *Journal of International Economics*, **84**(1), 48–64.

Bhattacharyya, S. & Hodler, R., 2010. Natural resources, democracy and corruption. *European Economic Review*, **54**(4), 608–621.

Binswanger, H. P., Barah, B. et al., 1980. Yield risk, risk aversion, and genotype selection: conceptual issues and approaches.

Bleakley, H. & Lin, J., 2012. Portage and path dependence. *The quarterly journal of economics*, **127**(2), 587.

Bleakley, H. & Lin, J., 2015. History and the Sizes of Cities. *The American Economic Review*, **105**(5), 558–563.

Boone, P., 1996. Politics and the effectiveness of foreign aid. *European economic review*, **40**(2), 289–329.

Braeu, R., 2010. Consumption tilting and the current account: Evidence from Canada. *International Review of Economics & Finance*, **19**(2), 304–312.

Brokensha, C., 2012. Skeleton Bay: The Death of a Wave Foretold. *Swellnet.com*.

Brunner, A. D., 2002. El Nino and world primary commodity prices: warm water or hot air? *Review of Economics and statistics*, **84**(1), 176–183.

Brunnschweiler, C. N. & Bulte, E. H., 2008. The resource curse revisited and revised: A tale of paradoxes and red herrings. *Journal of environmental economics and management*, **55**(3), 248–264.

Brunnschweiler, C. N. & Bulte, E. H., 2009. Natural resources and violent conflict: resource abundance, dependence, and the onset of civil wars. *Oxford economic papers*, **61**(4), 651–674.

Burnside, C. & Dollar, D., 2004. Aid, policies, and growth: reply. *The American Economic Review*, **94**(3), 781–784.

Butt, T., 2009. *The Surfer's Guide to Waves, Coasts and Climates*. Alison Hodge Publishers.

Butt, T. & Russell, P., 2004. *Surf science: An introduction to waves for surfing*. University of Hawaii Press.

Caldara, D., Fernandez-Villaverde, J., Rubio-Ramirez, J. F. & Yao, W., 2012. Computing DSGE models with recursive preferences and stochastic volatility. *Review of Economic Dynamics*, **15**(2), 188–206.

Calvo, G. A., 1998. Capital flows and capital-market crises: the simple economics of sudden stops. *The Journal of Economic Perspectives*.

Calvo, G. A. & Mendoza, E. G., 2000. Capital-markets crises and economic collapse in emerging markets: An informational-frictions approach. *The American Economic Review*, **90**(2), 59–64.

Carroll, C. & Kimball, M., 2008. Precautionary Saving and Precautionary Wealth. in The New Palgrave Dictionary of Economics, SN Durlauf and LE Blume, eds.

Carter, P., Postel-Vinay, F. & Temple, J., 2015. Dynamic aid allocation. *Journal of International Economics*, **95**(2), 291–304.

Caselli, F. & Michaels, G., 2013. Do oil windfalls improve living standards? Evidence from Brazil. *American Economic Journal: Applied Economics*, **5**(1), 208–238.

Cashin, P., Haque, N. U. & Olekalns, N., 2003. Tax smoothing, tax tilting and fiscal sustainability in Pakistan. *Economic Modelling*, **20**(1), 47–67.

Cashin, P. & McDermott, C. J., 2002. Intertemporal consumption smoothing and capital mobility: evidence from Australia. *Australian Economic Papers*, **41**(1), 82–98.

Cashin, P., Mohaddes, K. & Raissi, M., 2017. Fair weather or foul? The macroeconomic effects of El Niño. *Journal of International Economics*, **106**, 37–54.

Chatterjee, S. & Eyigungor, B., 2012. Maturity, indebtedness, and default risk. *The American Economic Review*, **102**(6), 2674–2699.

Chay, K. Y. & Greenstone, M., 2005. Does Air Quality Matter? Evidence from the Housing Market. *Journal of Political Economy*, **113**(2), 376–424.

Chen, X. & Nordhaus, W. D., 2011. Using luminosity data as a proxy for economic statistics. *Proceedings of the National Academy of Sciences*, **108**(21), 8589–8594.

Cheshire, P. C. & Magrini, S., 2006. Population growth in European cities: weather matters–but only nationally. *Regional studies*, **40**(1), 23–37.

Christiano, L. J., Eichenbaum, M. & Evans, C. L., 1999. Monetary policy shocks: What have we learned and to what end? *Handbook of macroeconomics*, **1**, 65–148.

Clarida, R., Gali, J. & Gertler, M., 2000. Monetary policy rules and macroeconomic stability: evidence and some theory. *The Quarterly journal of economics*, **115**(1), 147–180.

Coffman, M. & Burnett, K., 2009. The value of a wave: An analysis of the mavericks region, half moon bay, California. *Study commission for Save The Waves Coalition*.

Coibion, O. & Gorodnichenko, Y., 2011. Monetary policy, trend inflation, and the great moderation: An alternative interpretation. *The American Economic Review*, **101**(1), 341–370.

Collier, P. & Dehn, J., 2001. *Aid, shocks, and growth*, vol. 2688. World Bank Publications.

Collier, P. & Goderis, B., 2012. Commodity prices and growth: An empirical investigation. *European Economic Review*, **56**(6), 1241–1260.

Collier, P. & Gunning, J. W., 1999. Why has Africa grown slowly? *The Journal of Economic Perspectives*, **13**(3), 3–22.

Collier, P., van Der Ploeg, R., Spence, M. & Venables, A. J., 2010. Managing resource revenues in developing economies. *IMF Staff papers*, 84–118.

Conley, T. G., 1999. GMM estimation with cross sectional dependence. *Journal of econometrics*, **92**(1), 1–45.

Conley, T. G., 2010. Spatial econometrics. In *Microeconometrics*, 303–313. Springer.

Corden, W. M. & Neary, J. P., 1982. Booming sector and de-industrialisation in a small open economy. *The economic journal*, **92**(368), 825–848.

Cullen, J. B. & Levitt, S. D., 1999. Crime, urban flight, and the consequences for cities. *Review of Economics Studies*, **81**, 159–169.

Cust, J. & Poelhekke, S., 2015. The local economic impacts of natural resource extraction. *Annual Review of Resource Economics*, **7**(1), 251–268.

Dean, J., 2016. Kelly Slater built the perfect wave. Can he sell it to the world? *Bloomberg Businessweek*, **May 24**.

Deaton, A. & Miller, R., 1995. *International commodity prices, macroeconomic performance, and politics in Sub-Saharan Africa*. International Finance Section, Department of Economics, Princeton University.

Dehn, J., 2000. *The effects on growth of commodity price uncertainty and shocks*, vol. 2455. World Bank Publications.

Dell, M., 2010. The persistent effects of Peru's mining mita. *Econometrica*, **78**(6), 1863–1903.

Deller, S. C., Tsai, T.-H., Marcouiller, D. W. & English, D. B. K., 2001. The Role of Amenities and Quality of Life in Rural Economic Growth. *American Journal of Agricultural Economics*, **83**(2), 352–365.

Deverajan, S., Lewis, J. D. & Robinson, S., 1990. Policy lessons from trade-focused, two-sector models. *Journal of policy modeling*, **12**(4), 625–657.

Diamond, J., 1999. *Guns, germs, and steel: The fates of human societies*. WW Norton & Company.

Diamond, R., 2016. The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000. *The American Economic Review*, **106**(3), 479–524.

Doll, C. N., 2008. CIESIN thematic guide to night-time light remote sensing and its applications. *Center for International Earth Science Information Network of Columbia University, Palisades, NY*.

Donaldson, D. & Storeygard, A., 2016. Big Grids: Applications of Remote Sensing in Economics. *Journal of Economic Perspectives*, **30**(4), 171–198.

Dorélien, A., Balk, D. & Todd, M., 2013. What Is urban? Comparing a satellite view with the Demographic and Health Surveys. *Population and Development Review*, **39**(3), 413–439.

Durlauf, S. N., Johnson, P. A. & Temple, J. R., 2005. Growth econometrics. *Handbook of economic growth*, **1**, 555–677.

Easterly, W., 2003. Can foreign aid buy growth? *The journal of economic perspectives*, **17**(3), 23–48.

Easterly, W., Kremer, M., Pritchett, L. & Summers, L. H., 1993. Good policy or good luck? *Journal of Monetary Economics*, **32**(3), 459–483.

Eaton, J. & Gersovitz, M., 1981. Debt with potential repudiation: Theoretical and empirical analysis. *Review of Economic Studies*, **48**(2), 289–309.

Eilat, Y. & Einav, L., 2004. Determinants of international tourism: a three-dimensional panel data analysis. *Applied Economics*, **36**(12), 1315–1327.

Ellison, G. & Glaeser, E. L., 1997. Geographic Concentration in US Manufacturing Industries: A Dartboard Approach. *Journal of Political Economy*, **105**(5).

Ellison, G. & Glaeser, E. L., 1999. The geographic concentration of industry: does natural advantage explain agglomeration? *The American Economic Review*, **89**(2), 311–316.

Ellison, G., Glaeser, E. L. & Kerr, W. R., 2010. What causes industry agglomeration? Evidence from coagglomeration patterns. *The American Economic Review*, **100**(3), 1195–1213.

Erten, B. & Ocampo, J. A., 2013. Super cycles of commodity prices since the mid-nineteenth century. *World Development*, **44**, 14–30.

Faber, B. & Gaubert, C., 2015. Tourism and Economic Development: Evidence from Mexicos Coastline. *UC Berkeley Economics Department working paper*.

Ferrero, A. & Seneca, M., 2015. Notes on the underground: monetary policy in resource-rich economies. *OxCarre Working Paper*, (158).

Fetzer, T., 2014. Can workfare programs moderate violence? Evidence from India.

Fink, F. & Scholl, A., 2016. A quantitative model of sovereign debt, bailouts and conditionality. *Journal of International Economics*, **98**, 176–190.

Frankel, J. A., 2010. The natural resource curse: a survey. Tech. rep., National Bureau of Economic Research.

Friedman, M., 1957. *A theory of the consumption function: A study by the National Bureau of Economic Research*. Princeton University Press.

Fujita, M., Krugman, P. R. & Venables, A., 2001. *The spatial economy: Cities, regions, and international trade*. MIT press.

Gallup, J. L. & Sachs, J. D., 2000. Agriculture, climate, and technology: why are the tropics falling behind? *American Journal of Agricultural Economics*, **82**(3), 731–737.

Gandelman, N. & Hernandez-Murillo, R., 2014. Risk Aversion at the Country Level. Tech. rep., Working Paper 5, Federal Reserve Bank of St Louis.

Garcia-Ponce, O. & Wantchekon, L., 2014. Critical Junctures: Independence Movements and Democracy in Africa. Tech. rep., Working paper, Princeton University.

Gelb, A. H., 1988. *Oil windfalls: Blessing or curse?*. Oxford University Press.

Glaeser, E. L. & Gottlieb, J. D., 2009. The wealth of cities: Agglomeration economies and spatial equilibrium in the United States. *Journal of Economic Literature*, **47**(4), 983–1028.

Glaeser, E. L., La Porta, R., Lopez-de Silanes, F. & Shleifer, A., 2004. Do institutions cause growth? *Journal of economic Growth*, **9**(3), 271–303.

Grossman, H. & Van Huyck, J., 1988. Sovereign Debt as a Contingent Claim: Excusable Default, Repudiation, and Reputation. *The American Economic Review*, **78**(5), 1088–1097.

Gylfason, T., Herbertsson, T. T. & Zoega, G., 1999. A mixed blessing. *Macroeconomic dynamics*, **3**(02), 204–225.

Hamilton, J. D., 2003. What is an oil shock? *Journal of econometrics*, **113**(2), 363–398.

Handler, P. & Handler, E., 1983. Climatic anomalies in the tropical Pacific Ocean and corn yields in the United States. *Science*, **220**(4602), 1155–1156.

Hansen, H. & Tarp, F., 2001. Aid and growth regressions. *Journal of development Economics*, **64**(2), 547–570.

Harding, T. & Venables, A. J., 2010. Foreign exchange windfalls, imports and exports. *OxCarre Working Paper*.

Harrison, M., 2010. Valuing the Future: the social discount rate in cost-benefit analysis.

Heap, A., 2005. Riding the Super Cycle. *Metals and Mining Global Equity Research, Citigroup*, **31**.

Hemer, M., Jack Katzfey, J. & Hotan, C., 2011. The wind-wave climate of the Pacific Ocean. *Report for the Pacific Adaptation Strategy Assistance Program Department of Climate Change and Energy Efficiency*.

Henderson, J. V., Storeygard, A. & Weil, D. N., 2011. A bright idea for measuring economic growth. *The American economic review*, **101**(3), 194.

Henderson, J. V., Storeygard, A. & Weil, D. N., 2012. Measuring economic growth from outer space. *The American Economic Review*, **102**(2), 994–1028.

Hodler, R. & Raschky, P. A., 2014. Regional favoritism. *The Quarterly Journal of Economics*, **129**(2), 995–1033.

Housman, J., 2016. El Nio's Relentless Perfection. *Surfer Magazine*.

Hsiang, S. M., 2010. Temperatures and cyclones strongly associated with economic production in the Caribbean and Central America. *Proceedings of the National Academy of sciences*, **107**(35), 15367–15372.

Hsiang, S. M. & Meng, K. C., 2015. Tropical economics. *The American Economic Review*, **105**(5), 257–261.

Hsiang, S. M., Meng, K. C. & Cane, M. A., 2011. Civil conflicts are associated with the global climate. *Nature*, **476**(7361), 438–441.

Iizumi, T., Luo, J.-J., Challinor, A. J., Sakurai, G., Yokozawa, M., Sakuma, H., Brown, M. E. & Yamagata, T., 2014. Impacts of El Niño Southern Oscillation on the global yields of major crops. *Nature communications*, **5**.

IMF, 2016. World Economic Outlook: Subdued Demand: Symptoms and Remedies. *International Monetary Fund*, **October**.

Isham, J., Woolcock, M., Pritchett, L. & Busby, G., 2005. The Varieties of Resource Experience: Natural Resource Export Structures and the Political Economy of Economic Growth. *World Bank Economic Review*, **19**(2), 141–174.

Ismail, K., 2010. *The Structural Manifestation of theDutch Disease: The Case of Oil Exporting Countries*. No. 10-103. International Monetary Fund.

James, A., 2015. The resource curse: A statistical mirage? *Journal of Development Economics*, **114**, 55–63.

Jedwab, R., Kerby, E. & Moradi, A., 2017. History, path dependence and development: Evidence from colonial railroads, settlers and cities in Kenya. *The Economic Journal*, **127**, 1467–1494.

Jedwab, R. & Moradi, A., 2016. The permanent effects of transportation revolutions in poor countries: evidence from Africa. *Review of Economics and Statistics*, **98**(2), 268–284.

Judson, R. A. & Owen, A. L., 1999. Estimating dynamic panel data models: a guide for macroeconomists. *Economics letters*, **65**(1), 9–15.

Kilian, L., 2008. Exogenous oil supply shocks: how big are they and how much do they matter for the US economy? *The Review of Economics and Statistics*, **90**(2), 216–240.

Kletzer, K. M. & Wright, B. D., 2000. Sovereign debt as intertemporal barter. *The American Economic Review*, **90**(3), 621–639.

Laibson, D., 1997. Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, **112**(2), 443–478.

Lazarow, N., 2009. Using observed market expenditure to estimate the value of recreational surfing to the Gold Coast, Australia. *Journal of Coastal Research*, 1130–1134.

Lee, S. & Lin, J., 20157. Natural Amenities, Neighborhood Dynamics, and Persistence in the Spatial Distribution of Income. *Review of Economic Studies*, **0**, 1–32.

Lei, Y.-H. & Michaels, G., 2014. Do giant oilfield discoveries fuel internal armed conflicts? *Journal of Development Economics*, **110**, 139–157.

Lizaro, S., 2013. Default Risk and Risk Averse International Investors. *Journal of International Economoics*, **89**.

Longman, J., 2016. Surfing under Northern Lights. *The New York Times*, **18 October**.

Love, I. & Zicchino, L., 2006. Financial development and dynamic investment behavior: Evidence from panel VAR. *The Quarterly Review of Economics and Finance*, **46**(2), 190–210.

Malik, A. & Temple, J. R., 2009. The geography of output volatility. *Journal of Development Economics*, **90**(2), 163–178.

Manzano, O. & Rigobon, R., 2001. Resource curse or debt overhang? Tech. rep., National bureau of economic research.

Marshall, A., 1890. Principles of economics. *London: Macmillan*.

McGregor, T., 2011. Fiscal Options for Absorbing a Windfall of Natural Resource Revenues A CGE Model of Oil Discovery in Uganda. *OxCarre Working Paper*, (186).

McGregor, T., 2017. Commodity price shocks, growth and structural transformation in low-income countries. *The Quarterly Review of Economics and Finance*, **65**, 285–303.

McGregor, T., forthcoming. Pricing sovereign debt in resource-rich economies. *OxCarre Working Paper*.

McGregor, T. & Wills, S., 2017. Surfing a wave of economic growth. *OxCarre Working Paper*, (170).

Mehlum, H., Moene, K. & Torvik, R., 2006a. Cursed by resources or institutions? *The World Economy*, **29**(8), 1117–1131.

Mehlum, H., Moene, K. & Torvik, R., 2006b. Institutions and the resource curse. *The economic journal*, **116**(508), 1–20.

Mendoza, E. G., 2006. Lessons from the debt-deflation theory of sudden stops. Tech. rep., National Bureau of Economic Research.

Mendoza, E. G. & Yue, V. Z., 2012. A general equilibrium model of sovereign default and business cycles. *The Quarterly Journal of Economics*, **127**(2), 889–946.

Michaels, G. & Rauch, F., 2016. Resetting the Urban Network: 117-2012. *The Economic Journal*.

Michalopoulos, S. & Papaioannou, E., 2013. Pre-Colonial Ethnic Institutions and Contemporary African Development. *Econometrica*, **81**(1), 113–152.

Michalopoulos, S. & Papaioannou, E., 2014. National Institutions and Subnational Development in Africa. *The Quarterly Journal of Economics*, **129**(1), 151–213.

Murphy, K. M., Shleifer, A. & Vishny, R. W., 1989. Industrialization and the big push. *Journal of political economy*, **97**(5), 1003–1026.

Neumeyer, P. A. & Perri, F., 2005. Business cycles in emerging economies: the role of interest rates. *Journal of monetary Economics*, **52**(2), 345–380.

Nickell, S., 1981. Biases in dynamic models with fixed effects. *Econometrica: Journal of the Econometric Society*, **49**, 1417–1426.

Noble, F., 2017. The search is on! Mick Fanning surfs incredible 5km mystery wave... *The Daily Mail and AAP*, **15 February**.

North, D. C., 1989. Institutions and economic growth: An historical introduction. *World development*, **17**(9), 1319–1332.

North, D. C., 1990. *Institutions, institutional change and economic performance*. Cambridge university press.

Obstfeld, M., 1999. Foreign resource inflows, saving, and growth. Tech. rep., University of California at Berkeley.

Obstfeld, M. & Rogoff, K., 1995. The intertemporal approach to the current account. *Handbook of international economics*, **3**, 1731–1799.

Ogaki, M., Ostry, J. D. & Reinhart, C. M., 1996. Saving behavior in low-and middle-income developing countries: a comparison. *IMF Staff Papers*, **43**(1), 38–71.

Orphanides, A., 2002. Monetary-policy rules and the great inflation. *The American Economic Review*, **92**(2), 115–120.

Orphanides, A., 2003. Historical monetary policy analysis and the Taylor rule. *Journal of monetary economics*, **50**(5), 983–1022.

Pagan, A., 1984. Econometric issues in the analysis of regressions with generated regressors. *International Economic Review*, **25**, 221–247.

Park, J., 2017. Sovereign Default and Capital Accumulation. *Journal of International Economics*, **106**, 119–133.

Pastén, R. & Cover, J. P., 2015. Tax tilting and politics: Some theory and evidence for Latin America. *Journal of Macroeconomics*, **44**, 208–218.

Pinkovskiy, M. L., 2013. Economic Discontinuities at Borders: Evidence from Satellite Data on Lights at Night. Tech. rep., New York Federal Reserve.

Pouzo, D. & Presno, I., 2016. Sovereign default risk and uncertainty premia. *American Economic Journal: Macroeconomics*, **8**(3), 230–266.

Raddatz, C., 2007. Are external shocks responsible for the instability of output in low-income countries? *Journal of Development Economics*, **84**(1), 155–187.

Rajan, R. G. & Subramanian, A., 2011. Aid, Dutch disease, and manufacturing growth. *Journal of Development Economics*, **94**(1), 106–118.

Ramey, G. & Ramey, V., 1995. Cross-country evidence on the link between volatility and growth. *The American Economic Review*, **85**, 1128–1151.

Ravn, M. O., Schmitt-Grohé, S. & Uribe, M., 2008. Macroeconomics of subsistence points. *Macroeconomic Dynamics*, **12**(S1), 136–147.

Redding, S. J., 2010. The empirics of new economic geography. *Journal of Regional Science*, **50**(1), 297–311.

Redding, S. J. & Rossi-Hansberg, E., 2016. Quantitative Spatial Economics. Tech. rep., National Bureau of Economic Research.

Richmond, C. & Dias, D. A., 2009. Duration of capital market exclusion: An empirical investigation. Tech. rep.

Roback, J., 1982. Wages, rents, and the quality of life. *Journal of political Economy*, **90**(6), 1257–1278.

Roback, J., 1988. Wages, rents, and amenities: differences among workers and regions. *Economic Inquiry*, **26**(1), 23–41.

Robinson, J. A., Torvik, R. & Verdier, T., 2006. Political foundations of the resource curse. *Journal of development Economics*, **79**(2), 447–468.

Rodrik, D., 1996. Coordination failures and government policy: A model with applications to East Asia and Eastern Europe. *Journal of international economics*, **40**(1), 1–22.

Rodrik, D., 2004. Industrial policy for the twenty-first century. *CEPR Discussion paper*, (4767).

Rodrik, D., 2008. The real exchange rate and economic growth. *Brookings papers on economic activity*, **2008**(2), 365–412.

Rodrik, D., Subramanian, A. & Trebbi, F., 2004. Institutions rule: the primacy of institutions over geography and integration in economic development. *Journal of economic growth*, **9**(2), 131–165.

Rogers, J., 2004. *Hot commodities: how anyone can invest profitably in the world's best market*. New York, Random House Incorporated.

Romer, C. D. & Romer, D. H., 2004. A new measure of monetary shocks: Derivation and implications. *The American Economic Review*, **94**(4), 1055–1084.

Rosen, S., 1974. Hedonic prices and implicit markets: product differentiation in pure competition. *Journal of political economy*, **82**(1), 34–55.

Ross, M. L., 1999. The political economy of the resource curse. *World politics*, **51**(02), 297–322.

Sachs, J., 1982. The current account in the macroeconomic adjustment process. *The Scandinavian Journal of Economics*, **84**, 147–159.

Sachs, J., 2000. Notes on a new sociology of economic development. *Culture matters: How values shape human progress*, 29–43.

Sachs, J. D., 2001. Tropical underdevelopment. Tech. rep., National Bureau of Economic Research.

Sachs, J. D. & Warner, A. M., 1995. Natural resource abundance and economic growth. Tech. rep., National Bureau of Economic Research.

Sachs, J. D. & Warner, A. M., 1997b. Sources of slow growth in African economies. *Journal of African economies*, **6**(3), 335–376.

Sachs, J. D. & Warner, A. M., 1999. The big push, natural resource booms and growth. *Journal of development economics*, **59**(1), 43–76.

Sala-i Martin, X. & Subramanian, A., 2013. Addressing the natural resource curse: An illustration from Nigeria. *Journal of African Economies*, **22**(4), 570–615.

Scorse, J., Reynolds, I., Sackett, A. et al., 2015. Impact of surf breaks on home prices in Santa Cruz, CA. *Tourism Economics*, **21**(2), 409–418.

Sequeira, T. N. & Maçãs Nunes, P., 2008. Does tourism influence economic growth? A dynamic panel data approach. *Applied Economics*, **40**(18), 2431–2441.

Smith, B. & Wills, S., 2016. Left in the Dark? Oil and Rural Poverty. Tech. rep., Oxford Centre for the Analysis of Resource Rich Economies, University of Oxford.

Surfer Magazine, 2006. The Best Surfers in the World Land Somewhere in Mexico for the Rip Curl Pro. *Surfer Magazine*.

Tauchen, G., 1986. Finite state markov-chain approximations to univariate and vector autoregressions. *Economics letters*, **20**(2), 177–181.

The BBC, 2016. Surfonomics: las razones de dos economistas de Oxford para estudiar el valor de las olas. *The BBC Mundo*, **Aug 15**.

The Economist, 2012. Surf Economics: Beach Rush. *The Economist*, **March 17**.

The Economist, 2016. The Leisure Economy: Surfing to Success. *The Economist*, **Aug 13**.

The Financial Times, 2017. Australia seeks to ride the inland surfing boom. *The Financial Times*, **Jan 7**.

Torvik, R., 2002. Natural resources, rent seeking and welfare. *Journal of development economics*, **67**(2), 455–470.

Townsend, A. C. & Bruce, D. A., 2010. The use of night-time lights satellite imagery as a measure of Australia's regional electricity consumption and population distribution. *International Journal of Remote Sensing*, **31**(16), 4459–4480.

Ubilava, D., 2012. El Nino, La Nina, and world coffee price dynamics. *Agricultural Economics*, **43**(1), 17–26.

UNWTO, 2016. Tourism Highlights, 2016 Edition. Tech. rep., United Nations World Tourism Organisation.

Uribe, M. & Yue, V. Z., 2006. Country spreads and emerging countries: Who drives whom? *Journal of international Economics*, **69**(1), 6–36.

Van den Bremer, T. S. & van Der Ploeg, F., 2013. Managing and harnessing volatile oil windfalls. *IMF Economic Review*, **61**(1), 130–167.

van der Ploeg, F., 2011. Natural resources: Curse or blessing? *Journal of Economic Literature*, **49**(2), 366–420.

van der Ploeg, F., 2012. Bottlenecks in ramping up public investment. *International Tax and Public Finance*, **19**(4), 509–538.

van der Ploeg, F. & Poelhekke, S., 2009. Volatility and the natural resource curse. *Oxford economic papers*, **61**(4), 727–760.

van der Ploeg, F. & Poelhekke, S., 2010. The pungent smell of red herrings: Subsoil assets, rents, volatility and the resource curse. *Journal of Environmental Economics and Management*, **60**(1), 44–55.

van der Ploeg, F. & Poelhekke, S., 2017. The impact of natural resources: Survey of recent quantitative evidence. *The Journal of Development Studies*, **53**(2), 205–216.

van der Ploeg, F. & Venables, A. J., 2011. Harnessing Windfall Revenues: Optimal Policies for Resource-Rich Developing Economies. *The Economic Journal*, **121**(551), 1–30.

van der Ploeg, F. & Venables, A. J., 2012. Natural resource wealth: The challenge of managing a windfall. *Annual Review of Economics*, **4**, 315–337.

van der Ploeg, F. & Venables, A. J., 2017. Extractive revenues and government spending: Short-versus long-term considerations. Tech. rep., WIDER Working Paper.

Van Wijnbergen, S., 1984. The Dutch Disease: a disease after all? *The Economic Journal*, **94**(373), 41–55.

Venables, A. J., 2016. Using natural resources for development: why has it proven so difficult? *The Journal of Economic Perspectives*, **30**(1), 161–183.

Venables, A. J. & Wills, S. E., 2016. Resource Funds: Stabilising, Parking, and Inter-generational Transfer. *Journal of African Economies*, **25**(suppl 2), ii20–ii40.

Vicente, P. C., 2010. Does oil corrupt? Evidence from a natural experiment in West Africa. *Journal of development Economics*, **92**(1), 28–38.

Vieider, F. M., Truong, N., Martinsson, P. & Khanh, N. P., 2013. Risk preferences and development revisited. Tech. rep., Working paper.

Waltert, F. & Schläpfer, F., 2010. Landscape amenities and local development: A review of migration, regional economic and hedonic pricing studies. *Ecological Economics*, **70**(2), 141–152.

Williamson, J., 1990. Latin American Adjustment: How Much Has Happened? Tech. rep., Institute for International Economics.

Wills, S., 2015. Seven principles for managing resource wealth. *OxCarre Working Paper*.

World Bank, 2011. The Changing Wealth of Nations: Measuring Sustainable Development in the New Millennium. Tech. rep., World Bank, Washington DC.

Wu, J. & Gopinath, M., 2008. What causes spatial variations in economic development in the United States? *American Journal of Agricultural Economics*, **90**(2), 392–408.

Yesuf, M. & Bluffstone, R. A., 2009. Poverty, risk aversion, and path dependence in low-income countries: Experimental evidence from Ethiopia. *American Journal of Agricultural Economics*, **91**(4), 1022–1037.

Zhuang, J., Liang, Z., Lin, T. & De Guzman, F., 2007. Theory and practice in the choice of social discount rate for cost-benefit analysis: a survey. *ERD Working Paper*, **94**.