The London School of Economics and Political Science

Of Naxalites, Pirates and Microfinance Borrowers: Three Essays in Applied Microeconomics

Thiemo Rene Fetzer

October 2014

A thesis submitted for the degree of Doctor of Philosophy in Economics at the London School of Economics

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made. This thesis may not be reproduced without my prior written consent.

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of approximately 55,000 words.

Statement of Conjoint Work

I certify that chapters 2 and 3 of this thesis were co-authored. Chapter 2 was co-authored with Timothy Besley and Hannes Mueller, while chapter 3 was co-authored with Jon de Quidt and Maitreesh Ghatak. I contributed 40% of the work in chapter 2, and 33% of the work in chapter 3.

Abstract

This thesis consists of three chapters that fall under the broad banner of applied microeconomics, with a particular focus on the study of conflict and its social cost.

The importance of social security systems, providing a safety net for individuals to cope with shocks is well understood in developed countries. Less developed countries struggle with implementing social security schemes, due to a lack of state capacity, which often render these schemes less effective than designed. However, social insurance schemes may have far greater benefits in societies, where there is latent conflict. In the first chapter of this thesis, I study the Indian employment guarantee act and its effect on rural labour markets, mitigating adverse shocks by providing safe outside options. I show that this scheme has a significant effect on the dynamics of intra-state conflict: it moderates the cyclical nature of conflict triggered by adverse shocks and thus, helps to contribute to substantially lower levels of overall violence. The importance of technologies to smooth adverse shocks, in particular, shocks due to bad weather, will become increasingly important due to climate change.

The second chapter analyses a dimension along which conflict is costly. We estimate the impact of Somali piracy on the costs of trade. In spite of general agreement that establishing the rule of law is central to properly functioning economies, little is known about the cost of law and order breakdowns. We study shipping routes whose shortest path exposes them to the risk of piracy and find that the increase in attacks in 2008 lead to an 8 to 12 percent increase in shipping costs. We estimate the welfare loss due to piracy based on these estimates and arrive at a fairly conservative estimate: generating around 120 USD million of revenue for Somali pirates led to a welfare loss in excess of 630 USD million, highlighting that the welfare losses from trade disruptions are substantial and piracy capture only a small share relative to the loss in welfare.

The third chapter is reflecting my earlier research interest in the economics of micro finance. The chapter provides a theoretical model contrasting individual liability lending with and without groups to joint liability lending. This research is motivated by an apparent shift away from joint liability lending, while still retaining the group structure. We show under what conditions individual liability can deliver welfare improvements over joint liability, conditions that depend on the joint income distribution and social capital. We then show that lower transaction costs that mechanically favour group lending may also encourage the creation of social capital. In the last section, we draw on estimated parameters to simulate the model and to

quantify our welfare conclusions.

Acknowledgements

This thesis benefited from discussions, help and support from many people. I am extremely thankful to Gerard Padró i Miquel, who has helped me develop my thinking about my paper on the dynamics of conflict in India. Furthermore, I am thankful to my co-authors and advisors for their guidance. My work has greatly benefited from interactions with Tim Besley, Gharad Bryan, Mirko Draca, Greg Fischer, Christian Fons-Rosen, Maitreesh Ghatak, Vernon Henderson and Daniel Sturm .

I want to thank the Konrad Adenauer Foundation and the "Promotionskolleg Soziale Marktwirtschaft" for supporting me financially, academically and personally throughout the course of my undergraduate and postgraduate studies. Further, I would like to thank STICERD and CEP for providing an environment that fostered my research.

I would like to thank, especially, my colleagues with whom I ventured on this academic journey; especially, Michael Best, Konrad Burchardi, Jonathan Colmer, Laura Derksen, Erika Deserranno, Jon de Quidt, Jason Garred, Benjamin Guin, Sam Marden, Luis Martinez, Oliver Pardo, Amar Shanghavi, Miriam Sinn, Pedro Souza, and Dimitri Szerman.

I would like to thank my parents, Jakob and Andrea, and my sisters, Jana and Lara. Their loving support enabled me to carry on with my studies, despite the multiple setbacks I experienced along the way. Lastly, I would like to thank Claudia for her continued support, love, patience, care and understanding.

Contents

List of Tables

List of Figures

Preface

This thesis consists of three chapters that fall under the broad banner of applied microeconomics, with a particular focus on the study of conflict and its social cost.

Conflict is socially very costly but remains a prevalent feature in the 21st century. The first chapter studies the relationship between social insurance and the impact of local Monsoon weather shocks on conflict in India. Considering the vast literature that has studied the relationship between weather shocks and human conflict (see Hsiang et al. (2013)), the concern is that climate change could induce more human conflict. Hence, it becomes increasingly important to understand whether and how social insurance schemes could moderate the weather and conflict relationship by providing insurance. A central economic mechanism that drives conflict is the opportunity cost channel. A productivity shock puts downward pressure on workers' outside options, which renders joining or supporting insurgency movements incentive compatible. Insurgents draw from this increased support base and are able to affect more violence. The core of this argument implies that any intervention that smoothes away negative shocks should contribute to weaken the link between economic shocks and conflict, through its stabilising effect on workers outside options. The first chapter is a contribution that tries to tackle the question on whether public interventions can achieve this end.

The testing ground for this study is India. India has suffered from many low-intensity intra-state conflicts throughout its history; while these are endemic, they are still not too intense for the state to not function altogether. This allows me to study how a workfare scheme, introduced under the National Rural Employment Guarantee Act (NREGA) in 2006, could moderate the relationship between income and conflict.

The chapter focuses not on the levels of violence, but rather on the elasticity of violence with respect to income and how this relationship changes after the introduction of the workfare program. Before the introduction of NREGA, agricultural production, wages and violence in India were strongly rainfall dependent to the present day. This is a surprising finding, since the dependence of Monsoon rainfall should have been weakened through decades worth of investment in physical infrastructure such as damns, irrigation canals, or railroads and roads. Nevertheless, the elasticity between Monsoon rainfall and agricultural GDP estimated in this paper is actually higher than the one presented in the existing literature derived from historical data. A one percent increase in Monsoon rain, increases agricultural GDP per capita by

0.36%. This relationship between rainfall and agricultural incomes appears to be the driving force behind the strong reduced form relationship between Monsoon rain and conflict in India before the introduction of NREGA.

Following the introduction of NREGA, I highlight in the second step that NREGA appears to have completely removed the relationship between Monsoon rain and conflict. A similar pattern emerges when studying agricultural wages. The introduction of NREGA insulates agricultural wages from shocks, while agricultural output is still very much dependent on Monsoon rainfall. This suggests two things: first, NREGA serves as an effective tool to stabilise agricultural wages and thus incomes; however, it is not able to affect the underlying agricultural production function, at least in the time-period under study.

In the third step, I explore the underlying mechanisms that explain the reduced form findings. I show that NREGA does function as a stabiliser with take-up - both on the extensive, and the intensive margin strongly responding to contemporaneous and lagged rainfall. An 1% lower Monsoon rainfall realisation, increases NREGA participation by 0.2%. I show that NREGA expenditures offset around 1/3 of the income losses that can be attributed to Monsoon shocks.

My findings do not imply that India has become a more *peaceful* place since the results only suggest that a particular driver of conflict has lost its bite. Nevertheless, despite identification concerns, I provide some tentative evidence that suggests that overall levels of violence, following the introduction of NREGA, have gone down. I highlight that at least 1/3 of this decrease is driven by NREGA shutting down the opportunity cost channel, highlighting the importance of this mechanism driving conflict.

The second chapter analyses a dimension along which conflict is costly. We estimate the impact of Somali piracy on the costs of trade. For centuries, piracy has posed a threat to ocean-going trade. In essence, it is organized private predation which thrives in locations in which law and order is weak, either because particular states provide a safe haven or due to poor international cooperation. And it has repercussions for worldwide trade. However, despite the long-standing importance of piracy, little is known about its economic costs. The issue has been brought into sharp relief by the upsurge of piracy in the Gulf of Aden which poses a threat to one of the world's busiest shipping routes. Frequently attributed to the collapse of effective authority in Somalia, it has provoked an international response.

We match data on piracy attacks in the maritime area around Somalia to data on around 24,000 shipping contracts by constructing the closest navigable sea distance between each origin and destination port for which a ship has been chartered. This allows us to exploit the monthly time-series variation in the frequency of piracy attacks in the main areas affected by Somali piracy to estimate the impact of piracy on shipping costs. We then use these estimates to calibrate a model of the welfare cost of Somali piracy.

Our regression results show that shipping costs for dry bulk goods rose by between 8 and 12 percent when pirate activity increased in Somalia. We also show that these larger shifts mask significant variation across months. Charter rates fluctuate by 18 percent between the most and least dangerous months. This seasonal pattern in shipping prices is absent prior to the upsurge in pirate activity in the region during 2008. Accounting for this seasonal variation highlights that the average shipping costs through the Somali area did not increase during the months in which weather conditions inhibit pirates from operating.

The extra shipping costs that we uncover are mostly due to higher insurance costs and the increased security measures that are needed to repel pirate attacks. These constitute a welfare cost to the extent that labor and resources are allocated from productive tasks towards protection. Our model compares the extraction of resources through pirate attacks to a tax on shipping which finances an equivalent transfer. This allows us to calculate the welfare loss caused by piracy. Our central estimate suggests that the resource costs incurred in transferring around 120 million USD annually to Somali pirates is well in excess of 630 million USD.

The third chapter reflects the early part of my academic journey, in which I focused on studying micro finance. The chapter contrasts individual liability lending with and without groups to joint liability lending. The motivation for this research is the apparent and documented shift away from lending methods that use explicit joint liability in giving out loans. We show under what conditions individual liability can deliver welfare improvements over joint liability, conditions that depend on the joint income distribution and social capital. We then show that lower transaction costs that mechanically favour group lending may also encourage the creation of social capital. Finally, we simulate the model to quantify our welfare conclusions.

Chapter 1

Can Workfare Programs Moderate Violence? Evidence from India

Can public interventions persistently reduce conflict? This question is more important than ever. The last decades have seen dramatic episodes of social unrest, some of which turned violent leading to civil war and state failure. This has affected the lives of millions of people.[1] Between 1946 - 2005 it is estimated that civil wars claimed 10.1 million lives and currently, more than one third of developing countries are affected by internal conflict. In an effort to contain spreading conflicts, billions of dollars are spent on military interventions. This often takes the form of providing arms and training for different fighting groups. The open question is whether such money could have been spent to prevent conflict in the first place. The academic literature can help guide policy making as it has has put a lot of emphasis on identifying drivers of conflict. Two interlinked empirical regularities stand out. Low incomes provide a breeding ground for civil conflict (Collier and Hoeffler, 1998, 2004; Hegre and Sambanis, 2006) and adverse shocks to incomes induce new conflicts to break out or lead to an intensification of existing ones (Bazzi and Blattman, 2014; Dube and Vargas, 2013; Besley and Persson, 2008; Miguel et al., 2004; Fearon and Laitin, 2003). This robust empirical relationship provides a blue print for a policy: social insurance. Any public intervention that helps households smooth income following adverse shocks has the potential to break the link between income shocks and conflict.

The scope for public interventions to protect households from income risks is huge in developing countries. The 35 poorest countries with real GDP per capita less than USD 1000 have experienced 2.8 times more volatile growth in consumption per capita compared to the 35 richest countries.[2] The recent World Development Report 2014 arrives at similar figures suggesting that household consumption risk in

[1] A large literature in economics has tried to assess the true social and economic cost of conflict and the many channels through which it operates, such as by deterring human capital investment (Blattman and Annan, 2010; Leon, 2009; Akresh and Walque, 2008), affecting time preferences (Voors et al., 2012), affecting capital investments (Singh, 2013), diverting foreign direct investment (Abadie and Gardeazabal, 2008) or increasing trade costs (Besley et al., 2014).

[2] Computed as simple average using the World Bank Development Indicators studying growth in consumption per capita between 1995-2011.

developing countries is of several orders of magnitude larger compared to developed countries. Yet, the share of public resources devoted to social insurance in developing countries is dismal. Data from the International Labor Organisation for the same set of countries suggests that the share of GDP devoted to social protection for the poorest countries with real GDP per capita less than USD 1000 lies at 4.4 % compared to an average of 20.5 % for the set of richest countries.[3] At the same time countries with volatile consumption per capita experience a lot more social unrest and crime. Homicide rates in the poorest countries with most volatile consumption growth are six times larger compared to developed countries.[4] Since income risks are so pronounced in developing countries, an effective social insurance could have profound effects on the dynamics of conflict: it could break the link between income shocks and conflict. This chapter shows that social insurance can achieve just that.

The challenge for researchers is to find a context in which the interplay between social insurance and conflict can be studied. This is not easy to come by. First, it is difficult to find a developing country context in which an effective social insurance has been introduced. This is not aided by the fact that developing countries spend little public resources on social protection despite the pronounced consumption risk. Even if a country spends significant resources on social protection it is not clear whether this truly reflects social insurance: providing a state-contingent pay out to individuals that are adversely affected by shocks. Given a set of policies that developing countries classify as social protection, this needs to be refined to only include policies that have the potential to function as insurance. Last but not least, if such a policy has been identified in a country, it is important to bear in mind that delivery of social insurance may be particularly difficult in countries that already experience conflict. India checks all three boxes providing a unique testing ground to study the relationship between social insurance and conflict. First, the country has suffered from many low-intensity intra-state conflicts throughout its history. These conflicts are endemic but have a relatively low intensity so that the state still functions on many dimensions. Secondly, India has put forth many development schemes. Most recently from 2006 onwards, India has introduced a public employment program through the National Rural Employment Guarantee Act (NREGA). This programme has the potential to function as social insurance by providing public employment on local infrastructure projects at minimum wages when households demand it. Third, NREGA is large and due to its scale may have an impact on the dynamics of conflict. It is the biggest public employment scheme in mankind's history, currently reaching up to 47.9 million rural households annually, generating 210 million person-days of employment. On a typical day, 7.7 million workers are expected to show up to work on one of nearly 294 thousand work sites.[5]

This chapter makes three contributions. First, the chapter highlights that income

[3] Computed as simple average of using data on Total public social expenditure as a percentage of GDP collected by the International Labor Organization for the most recent year available for each country. If one only looks at non-pension spending, the shares are 3 % and 13.7% respectively.

[4] Computed using data from the most recent United Nations Homicide Statistics.

[5] See `http://nrega.nic.in/netnrega/mpr_ht/nregampr.aspx`, accessed on 14.06.2014.

variation is an important driver of conflict in India. The chapter then studies the relationship between insurgency conflict and social insurance in India. The focus is not on levels of conflict but rather on the relationship linking adverse shocks to incomes with conflict and how this relationship changes once NREGA has been introduced. This is a natural extension to the vast literature studying the effect of income shocks on conflict and crime. The chapter highlights that public employment delivers effective insurance as it is designed to induce self-targeting. This helps overcome problems of adverse selection inherent to transfer schemes. The NREGA employment requires households to exert effort by working on public infrastructure projects at minimum wages. This ensures that households with ample economic opportunities have no incentive to work under NREGA even when they could credibly signal that they have been adversely affected by a shock. The self-selection of households ensures that NREGA resources are sent to households that are most vulnerable and have no access to a better paying job opportunity (see Besley and Coate, 1992; Nichols and Zeckhauser, 1982). The requirement for infrastructure construction through the public employment can further reduce moral hazard problems, since the public employment should produce relatively easily verifiable output. The chapter also studies the potential indirect insurance benefits that public employment can deliver: first, agricultural labor markets could become more resilient to adverse shocks and second, agricultural production could become more resilient against local productivity shocks.

The chapter has three main findings. First, I show that before the introduction of NREGA there is a strong relationship between local Monsoon shocks, proxys of agricultural income and conflict. This complements and reinforces the findings of the existing literature on conflict in India. I then show that since the introduction of NREGA the relationship between Monsoon shocks and agricultural wages has statistically disappeared. Local Monsoon shocks cease to have an effect on agricultural wages, while they continue to strongly predict agricultural output. This suggests that agricultural productivity and agricultural wages have decoupled since NREGA was introduced. More importantly, I show that the relationship linking Monsoon rainfall shocks and conflict has disappeared. Studying violent crimes and rioting suggests similar results. The introduction of NREGA induced an inward rotation of the relationship linking Monsoon shocks and conflict or crime. The finding is robust to an array of checks and is most pronounced when studying conflict events where the target of violence are civilians. This suggests that NREGA helps bring civilians out of the line of fire. The third findings highlight that NREGA provides insurance. I show that public employment under NREGA is utilized as a tool helping households to smooth consumption following adverse shocks. Participation in the program strongly responds to adverse Monsoon shocks along the extensive and intensive margin. In addition I provide a simple quantification exercise suggesting that 30% of the district level income losses due to an adverse Monsoon shock are offset by direct NREGA expenditure flowing into a district. This does not capture the indirect benefits gained by households from stabilized agricultural wages.

This chapter also makes some advances methodologically. This is the first chapter to use a novel conflict dataset that covers the whole of South Asia and has been constructed using scalable Natural Language Processing Tools (presented in Fetzer, 2013). The semi-automated coding procedure makes the process of coding data highly transparent and can be used to complement human coding of conflict data. This highlights the possibility to use semi-automated machine-learning routines for data cleaning and preparation in a field of economics research, where data availability and coding routines have been identified as an important constraint (Blattman and Miguel, 2010).

This chapter contributes to the nascent literature that evaluates the extent to which public interventions can affect the dynamics of conflict. Moderating the relationship between conflict and productivity shocks requires insulating personal incomes from these shocks. Technologies that can break the link between productivity shocks and incomes can be classified into three categories: (1) physical infrastructure, (2) new production technologies or (3) politically created institutions. Most of the empirical literature has focused on evaluating whether these technologies increase levels of income, rather than moderating income volatility.[6] Only recently, some studies have emerged that take the results form these chapters to study whether they help break the link between productivity and conflict. In the first category falls Sarsons (2011)'s paper, which builds on work by Duflo and Pande (2007) suggesting that the construction of dams moderated wage volatility, but appear not to have moderated Hindu-Muslim riots. This is not too surprising as physical infrastructure, while increasing levels of income, may not prove to be effective in providing insurance. Hornbeck and Keskin (2014) finds that farmers adjust their production technologies to take advantage of irrigation, which leads to higher production levels but not necessarily lower volatility. The insurance that is provided by access to irrigation may induce households to take more risks ex-ante in the crop choice. In the second category falls the work by Jia (2014), who studies the moderating effect of the drought resistant sweet potato as a new technology on the incidence of riots in historical China. She finds that the sweet potato persistently reduced the impact of droughts on rioting. Production technologies that help households cope with shocks may not be available in certain contexts. This chapter is the first to fall into the third category, evaluating whether a politically created institution such as India's National Rural Employment Guarantee achieves the goal to insulate personal incomes from negative shocks and through that, remove the income dependence of conflict.

This chapter also relates to the wider literature on the economics of conflict and labor markets. Shapiro et al. (2011) study how levels of unemployment affect levels

[6] Duflo and Pande (2007) evaluate the construction of dams and its impact on agricultural production in India. Aggarwal (2014) evaluates the impact of road construction, while Donaldson (2010) studies the impact of railroad construction in colonial India. Burgess and Donaldson (2010) build on that work to study how trade integration may have cushioned the effect of adverse productivity shocks on famine mortality. Another vast literature tries to understand and design effective rainfall or weather insurance schemes (see e.g. Lilleor and Giné (2005) or Cole et al. (2008)).

of insurgency violence in Afghanistan, Iraq and the Philippines, finding no support for an opportunity cost channel at work. Iyengar et al. (2011) on the other hand, highlight that increased construction spending seems to cause lower levels of labour intensive violence. Annan and Blattman (2014) present results from a randomised control trial in Liberia, indicating that interventions providing training and capital can greatly increase the opportunity cost of becoming a mercenary and thus, contribute to weaken the relationship between shocks and conflict.[7] A smaller literature studies conflict in India, in particular studying the Maoist movement and the driving forces behind this conflict (Gomes, 2012). Vanden Eynde (2011) and Gawande et al. (2012) established that the Naxalite conflict varies systematically with incomes or proxies thereof, suggesting an opportunity cost channel at work. This chapter builds on to their work studying conflict and crime across the whole of India and how the NREGA workfare scheme weakened the link between income shocks and conflict.

The chapter also relates to recent research efforts that seek to estimate level effects of the introduction of NREGA on conflict levels in India. This chapter does not focus on the level effect of NREGA for two reasons. First, there are identification concerns as the roll-out of NREGA was clearly targeted towards poor and vulnerable districts: districts that experienced conflict received NREGA earlier, making districts that received NREGA later a poor counterfactual. It is conceptually not clear why one should expect to see an effect of NREGA on conflict levels, bearing in mind that NREGA provides social insurance. This chapter highlights that NREGA functions as social insurance following adverse shocks. The insurance effect of NREGA on conflict levels should manifest itself through the interaction with adverse shocks. The effect of NREGA on conflict levels should manifest itself indirectly over time. With this in mind, I estimate level effects in the appendix and find similar results to Dasgupta et al. (2014). They use a difference-in-difference estimator to estimate the level effect of NREGA. The identifying variation this relies on is coming from the sequential roll out of NREGA, coming solely from two years of data. Khanna and Zimmermann (2013) address the endogeneity of the roll out directly and use a fuzzy regression discontinuity design. They reverse engineer the NREGA roll out algorithm to identify districts that were near the cutoff of being assigned into an earlier or later phase. Districts close to the cutoff serve as counterfactual. Their results indicate that the introduction of NREGA lead to an increase in levels of conflict in the short run. The concern with this research design is that it lives off very few observations of districts on either side of the cutoff that experience conflict variation. However, some districts included may have never experienced any conflict event throughout the whole sample period and thus, may not be a good counterfactual.

The chapter is organised as follows. The second section provides some background on the context of conflict in India and the NREGA workfare program. Section 3 presents the data used, discussing the novel conflict dataset created for this chapter. Section 4 presents the empirical strategy used in this chapter. Section 5 discusses

[7] This contrasts with Blattman et al. (2014), who find that a Ugandan employment program, despite large income gains, is not correlated with lower levels of aggression or protests.

the main results and provides robustness checks. In Section 6 I show that NREGA functions as insurance and quantify to what extent it offsets risks. The last section concludes.

1.1 Context: Conflict and Insurance in India

India is affected by multiple insurgency movements. The most prolific insurgencies are the Maoist insurgency that stretches across East India and into several districts in India's Northeast (also known as the 7 Sister States). In the North East, various insurgency movements strive for political independence from the Indian government. Most prolific are the movements in the states of Manipur and Assam. It is difficult to separate these conflicts from one another due to the geographic proximity and existing interlinkages. The Maoist groups have documented ties with insurgencies in the North East, in particular with the Manipur based People's Liberation Army (PLA) and the Assam based United Liberation Front (ULFA).[8] This makes insurgency a regional phenomena covering most of East India and the North East. A particular focus of the academic literature has been the Maoist conflict.[9] The movement started out as a peasant revolt against extortive labor relationships with landlords in West Bengal. In May 1968 the "All India Coordination Committee of Communist Revolutionaries" (AICCCR) was formed. This organization became the root for the armed struggle of multiple organizations, including the Communist Party of India-Maoist (CPI-M). The CPI-M, in its present state, is the result of mergers of various groups beginning in the late 1990s and the early 2000s. CPI-M as an organization consists of a political wing and an armed wing and is considered a terrorist organization. It is estimated that the military wing the People's Liberation Guerilla Army consists of at least 10,000 combatants.[10] Originating in West Bengal, the movement has spread to less developed areas of rural southern and eastern India. The Naxalites are especially prolific in states of Chhattisgarh, Jharkhand, Bihar, Orissa and Andhra Pradesh; but they are also present in some states in the North East, in particular in Assam, Arunachal Pradesh and Tripura. In 2006, around 1/3 of India's roughly 600 districts were considered under the influence of left-wing extremism or subject to violence, forming a "red corridor" that stretches across India (see Figure 1.1).[11] The aim of the Maoists is to overthrow the existing government and establish a communist state. The Union government under Manmohan Singh has announced that Naxalism poses India's largest internal security threat.[12]

The Maoist movement is most prolific in India's rural districts. Districts affected

[8] See Lok Shaba Unstarred Questions 3138 (13.12.2011) and No 1964 (17.12.2013). Ramana (2007) and `http://goo.gl/bB8wLp`, accessed 20.01.2013.

[9] See Gawande et al. (2012); Gomes (2012); Hoelscher et al. (2012); Vanden Eynde (2011); Khanna and Zimmermann (2013); Morrison (2012); Morgan and Reiter (2013).

[10] Henceforth, I will refer to the CPI-M and its military wing as the Naxalites, the Naxalite movement or the Naxals.

[11] I will work with district definitions as per the 2001 census.

[12] See `http://www.economist.com/node/15579946`, accessed 25.02.2011.

by left-wing extremism are marked by underdevelopment: districts under left wing extremist influence are characterized by lower rates of urbanization, higher degrees of illiteracy and limited access to infrastructure (such as paved roads, electricity, primary education or health care facilities, see Table 1.11.) The Naxalites are also very prolific in the jungle districts of Chhattisgarh, Bihar and Jharkhand, where they control swaths of land and are said to draw support from the high share of tribal people living there.

The economic livelihood in left-wing extremist affected districts is dominated by subsistence farming, sharecropping or wage employment in the agricultural sector. The states of Jarkhand and Chhattisgarh host most of India's coal and bauxite reserves which attracts a lot of investment for resource exploitation. The Naxalties are said to extort taxes from mining companies and intimidate firms. In forested areas the main source of income is the production of forest produce such as the collection of *tendu* leaves, which are used to produce cigarettes. According to data from the National Sample Survey 2001, 64.9% of households directly rely on agriculture as primary means of income in India. In states with significant Naxalite presence, this share is significantly higher; for example, up to 90% of Chhattisgarh's population is employed in agriculture. For the tribal population, Gawande et al. (2012) highlight the relative importance of income from forest produce. Any shock to local incomes has dramatic consequences to the rural livelihoods. This results in very poor developmental indicators with tribal households displaying significantly higher levels of food insecurity. Nearly 71.6% of tribal households being food deficit for 2-3 months in a year and 79% of tribal children being anemic (Radhakrishna and Ray, 2006).

The Naxalites are highly organized. Across regions, the political wing has a Central Committee that makes key strategic decisions, while Regional- and State Bureaus are responsible for organization of coordinated activities, such as strikes. Local Squad Area Committees have a high degree of autonomy on individual operations. The military wing is called the People's Liberation Guerilla Army (PLGA) and has similar structure. At the village level a civilian militia operates. They act as informants and provide direct support and shelter for armed squads. New recruits are typically sent for training into training camps. There they receive a basic military training, lasting between 6-12 months, which equips them with the necessary knowledge of guerilla warfare, including handling of rifles and minor explosives (such as hand grenades, land-mines and improvised explosive devices), before joining an active fighting squad in their home district. The Naxalites have ties with insurgency movements in the North East. They cooperate in the procurement of arms and the training of new recruits.[13]

The governments response to the various insurgency movements is hindered by the federal structure of the Indian Union. Law and order rests in the domain of state governments. The Maoist insurgency is flexible across borders. The Operation Green Hunt, that took off in early 2010 was the first integrated response against Naxalism

[13] See Lok Shaba Unstarred Questions 3138 (13.12.2011) and No 1964 (17.12.2013).

by the central government conjoint with the states. Estimates suggest that around 100,000 Central Reserve Police Force personnel operate along the side of State Police Forces, mainly in Chhattisgarh, Bihar, Jharkhand and Orissa.[14] The deployment of paramilitary is the latest military response of the Indian state. However, some development policies were also put forward to tackle chronic underdevelopment. One policy is the Integrated Action Plan (IAP) which in December 2010 released a block grant of INR 25 crore (roughly USD 4 million) as additional funding for infrastructure development. The scheme was first conceptualized to provide additional funds to 35 districts that were severely affected by left wing extremism, but was later expanded and now covers around 80 backward districts. Projects funded through the IAP are decided upon by district-level committees with limited involvement of local stakeholders.

Naxalite's do want to be seen as the advocates of local interests to gain legitimacy and through that, foster their popular support and recruit active fighters. A lot of these grievances are brought into relief by economic shocks. Some examples of issues aggravated by economic shocks are relationships with moneylenders who forcefully demand repayment, grievances surrounding sharecropping arrangement, which leave farmers with little of the produce, or low wages being paid. In these environments Naxalites are said to step in and protect the interests.[15] This could involve launching of "famine raids" (Dash, 2006), calls for *bandhs* (strikes) to push for higher wages (Ranjan and Prasad, 2012) or targeted violence against civilians that may turn to becoming police informers (Vanden Eynde, 2011).

It is in environments of deprivation and hunger when they are able to actively recruit new fighters from the local population without having to rely on coercive means. The anecdotal accounts on recruiting of insurgents are concentrated around the Maoist insurgency. Verma (2011) argues that environments of deprivation are ideal for "Maoists to step in, by paying a handsome amount of around Rs. 3000 to the young and promising parents that their kid will have food and money." There are accounts suggesting that Naxalites use the extortion revenues and money from Narcotics trade to pay monthly stipends of around Rs 1,500 (see Ramana, 2007). This figure is significant compared to average agricultural wages ranging between Rs 50 - 70 per day in India's poorest districts (see Table 1.1). Higher levels of recruitment will ultimately lead to more violence, once the new recruits have been trained and sent back to their home districts.

While these are distinct mechanisms, they have in common that they are brought into relief by adverse shocks. This generates the widely observed correlation between contemporaneous violence and lagged Monsoon season rainfall that has been

[14]Some states have responded by setting up local militias recruited as Special Police Officers (SPOs). In Chhattisgarh, the state government has commissioned and armed local militia known as "Salwa Juddum" (peace hunt), which is said to operate with complete impunity and lead to an escalation of violence.

[15]See discussion in Pandita (2011) relating to money lenders, Deshpande and Shah (2010) for a discussion about the relationship between droughts and farmer suicides in relation to money lending, Prasad (1987) for a discussion of tenancy arrangements and conflict in Bihar.

observed in the data. This chapter studies how the relationship between local Monsoon shocks and conflict (or general violent crimes) changes with the introduction of the Mahatma Gandhi National Rural Employment Guarantee Act.

The Mahatma Gandhi National Rural Employment Guarantee Act (NREGA) was passed in 2005 and establishes a "right to work" by providing a legal entitlement to 100 days of (minimum) wage employment per household and financial year to rural households. The Ministry of Rural Development considers the program to be "the largest and most ambitious social security and public works program in the world" (Ministry of Rural Development, 2012). The goal of NREGA is to develop a strong social safety net for vulnerable groups by providing employment sources when other employment alternatives are scarce (Ministry of Rural Development, 2008). It is also envisioned as a "growth engine for sustainable development of an agricultural economy, through the process of providing employment on works that address causes of chronic poverty such as drought, deforestation and soil erosion. The Act seeks to strengthen the natural resources base of rural livelihood and create durable assets which have potential to generate additional employment in the years to come in rural areas."[16]

The employment scheme that is defined under the act was rolled out in three phases between 2006 and 2008 and covers nowadays all Indian districts.[17] In the first phase, 200 districts received the program from the first quarter of 2006 onwards. In 2007, 130 further districts were added, while, in early 2008 the last phase the remaining rural districts received the scheme. The order of the roll out was far from being random. Table 1.1 provides summary statistics for the districts falling in different phases. Districts that received NREGA in earlier rounds had significantly lower agricultural output per capita and wages. They were more likely to be considered under left-wing extremist influence and more likely to be experiencing any conflict event. Access to infrastructure, such as roads, health-care and postal services are also a lot worse. The endogeneity of the roll-out makes it difficult for any empirical study that aims to identify a level effect of the program.

The administration of NREGA is decentralized and aims to empower local governance structures down to the Panchayat level, the lowest level of governance in India. A panchayat typically comprises a few villages or hamlets. In case people want to work under the scheme they approach their local panchayat representative and express their interest to work. The Gram Panchayat will issue a Job Card used to identify a household. Each household can request one Job Card and all adult members of a household that are willing to work under NREGA will be registered on that households job card. There is no cost for creating a job card borne by households. The panchayat then has to provide work on a public project within a two week period at the given state-level minimum wage.[18] As NREGA employment requires house-

[16]See `http://www.nrega.nic.in`, accessed 12.02.2014.

[17]Excluded are a few purely urban districts. I refer to the scheme that is defined under the Act as NREGA from hereon as well.

[18]Minimum Wage Laws existed before the inception of NREGA. These had limited impact as Mini-

holds to exert effort by working on public infrastructure projects at minimum wages, this ensures that households with ample economic opportunities have no incentive to work under NREGA even when they could credibly signal that they have been adversely affected by a shock. Such targeting is not achieved by any other government schemes that could be classified as providing social protection; in particular, the Public Distribution System or the system of Minimum Support Prices for agricultural produce.[19] If the panchayat fails to provide work, a daily unemployment allowance (which is below minimum wage), financed by the state government, is to be paid. The projects on which workers are employed have to be in close proximity to the home of the worker (at most 5 km distance) and there is additional remuneration for transportation costs or living expenses, while on the work site. The NREGA act further requires that 60 percent of the budget for a project be allocated to wages. Also, the use of machines or contractors is prohibited. Another requirement is that at least 1/3 of the workers need to be female.

The design of NREGA included a major push for financial inclusion by requiring that all wage payments be made through the banking sector or through bank accounts held with postal offices. NREGA income is to be paid weekly by wire transfer to local post offices or bank accounts. This is not fully implemented across India to date which is a big concern due to corruption in the scheme (see Niehaus and Sukhtankar, 2013b).

The types of works are decided at the local level. Districts prepare a shelf of projects which need to be agreed on with local panchayats. In particular, the act seeks to empower panchayats by giving them the right to assign priorities to infrastructure projects that meet local needs and preferences. The type of infrastructure projects range from drought proofing of land, to micro-irrigation works, rural-sanitation and rural connectivity. As NREGA provides a legal entitlement that becomes available when households demand employment it can be used by households as insurance against adverse shocks. This allows NREGA to function as a form of social insurance, providing protection against idiosyncratic shocks. Low Monsoon rainfall is robustly correlated with low agricultural output and wages. Employment on public works may be an attractive outside option in such situations. On the other hand, wages paid under NREGA are not too high as they are fixed at the respective minimum wages. This ensures that NREGA work becomes unattractive in times of ample economic opportunities elsewhere.

In addition to the direct employment generation there are also indirect effects that can contribute to the program functioning as insurance of local incomes. These indirect effects may be attributable to the types of infrastructure being constructed under the scheme. Micro-irrigation infrastructure, constructed through the scheme, may persistently moderate the rainfall income relationship.

---

mum Wage Laws were either not enforced (Planning Commission, 2008) or simply did not apply as in the case of self-employment which is the predominant form of agricultural employment.

[19] Despite a reform in 1997, turning the Public Distribution System into a Targeted Public Distribution System targeting is still extremely poor (Gadenne, 2014).

The scheme under the NREGA Act is the largest known workfare program, generating 2.84 billion person-days of employment during the financial year 2009-2010 for 53 million rural households and thus, benefiting 291 million individuals. For that year, each participating household worked an average of 54 days under the scheme. The fiscal expenditure for that financial year amounted to INR 37,900 crore, or USD 6.3 billion. Out of this, wages amounted to INR 25,500 crore, or USD 4.2 billion, implying additional labor income of USD 79 per household and financial year. This stands significant in contrast to an agricultural output per capita of INR 13,500 or USD 226.[20] NREGA flows into districts that are vulnerable to conflict. The 222 districts out of 543 districts that experience conflict variation in my sample accounted for more than 50% of the expenditures under NREGA throughout. The mere size of the scheme in relation to any other development scheme generate the possibility for NREGA to have a profound impact on the relationship between weather shocks and agricultural sector income. In the next section, I discuss the main data sources used in this study, before proceeding to the empirical strategy of the chapter.

1.2 Data

This chapter combines data from many different sources to provide an overarching picture on how NREGA affected the dynamics of conflict. There are three main data sources to be highlighted. First, the data covering conflict across India. This dataset is developed using an approach that is novel to the conflict literature by relying on language processing algorithms to code conflict events based on newspaper reports. I complement this dataset with official crime reporting. The second major effort is to develop agricultural income and wages data exploiting a multitude of different data sources. Lastly I use some novel remote sensed rainfall data and various other sources of weather data.

District Level Conflict data Empirical research on the economics of conflict almost always suffers from severe data limitations. This lies in the nature of the subject of study. Typically places that exhibit conflict are only weakly institutionalized with little official reporting and little press and media coverage. Blattman and Miguel (2010)'s review cites that the correlation across different civil war datasets ranges from 0.42 to 0.96, which may be the reason why empirical results are often not reproducible using similar identification strategies, but different datasets or variable definitions (Ciccone, 2011).

For civil war datasets differences can easily be reconciled. However, the conflict literature is moving increasingly to study more micro-datasets at finer spatial and temporal resolutions. Researchers are often left with a set of primary data sources, such as newspaper reports or news-feeds from wire services that need to be translated into a workable dataset for the econometrician, providing conflict event or inci-

[20] Agriculture accounts for 18% of GDP per capita, but employs 51% of the labor force. See World Development Indicators, `http://data.worldbank.org/indicator/SL.AGR.EMPL.ZS`, accessed 22.08.2014.

dence counts at a certain spatial- and temporal resolution. As these research efforts are decentralized this could result in many different datasets being coded up from the same raw data sources. The datasets need not coincide as researchers apply different coding practices. This renders the datasets not easily comparable, subject to subjectivity bias which also makes them difficult to expand. This chapter uses a novel approach to code violence data for the whole of India stemming from 28,638 newspaper clippings collected by the South Asia Terrorism Portal (SATP). The SATP newspaper clippings represents the most extensive and systematic collection of raw sources covering conflict in India.[21] The primary sources have been used in the context of studying conflict in India by many different authors.[22]

This chapter is the first to use a novel approach for coding conflict event data derived from primary sources. The idea is to use computer algorithms for language processing to emulate the way humans would code conflict data. The core unit of analysis is a sentence in a newspaper clipping. For each sentence key pieces of information are obtained, namely the subject-, verb- and object, along with the time and location that an event took place.[23]

$$\underbrace{\text{Two unidentified terrorists}}_{\text{subject}}\;\overbrace{\text{massacred}}^{\text{verb}}\;\overbrace{\text{six members of a family}}^{\text{object}}$$

$$\text{and}\;\overbrace{\text{left}}^{\text{verb}}\;\overbrace{\text{a seventh}}^{\text{object}}\;\overbrace{\text{injured}}^{\text{verb}}$$

$$\overbrace{\text{at Mangnar Top, Poonch district, on December 31, 2001.}}^{\text{time \& location}}$$

To illustrate, consider the above example of a sentence. The routine identifies, for every verb, its underlying subject, object and surrounding meta-information, such as time and locations, which are indicated by prepositions or due to their syntactic position in a sentence. With this processing step achieved, the data can be further refined. In the above case, we may want to label the perpetrator ("two unidentified terrorists") of the act of "massacring" to be "terrorists" and the subject ("six members of a family") to be "civilians". This allows a further study of the targets of violence and an analysis of casualty figures.

This routine is an improvement compared to what the existing literature does. Firstly, this approach allows the study of a myriad of "acts" that are reported. Many

[21] As the SATP presents only data from English language sources, there may be a systematic selection problem as indicated by Gawande et al. (2012). For the purpose of this chapter this is not a concern unless the selection is correlated systematically with rainfall shocks over time.

[22] There is a multitude of research papers that have separately hand coded subsets of the primary SATP newspaper clippings covering various Indian states or various time-periods, see Dasgupta et al. (2014); Gomes (2012); Hoelscher et al. (2012); Gawande et al. (2012); Khanna and Zimmermann (2013); Rana (2013); Shrivastava (2014); Vanden Eynde (2011); Buhaug and Wischnath (2014).

[23] Language processing algorithms, developed for the English language, but increasingly for many other languages as well, achieve very high accuracy rates in providing a correct syntactic analysis of a sentence, see Fetzer (2013) for a detailed discussion.

hand-coding approaches would restrict the analysis ex-ante to a set of verbs that are indicative of violent activities, such as "to kill". In the automated approach, this can be done ex-post. This way, one is able to include terrorist acts that did not involve casualties, such as attacks on infrastructure, where the word "to kill" would not have appeared. Furthermore, there is no limit to the geographic scope. Some authors have restricted the analysis ex-ante to cover only certain Indian states to keep the hand coding manageable or they have searched for district names.[24] A third concern that this approach addresses is human subjectivity. As the routine relies on natural language processing algorithms it removes any subjectivity bias from the coding that may emerge. A third advantage is the scalability of the routine. For the purpose of this paper, I construct two main dependent variables. The first is an indicator variable, that is simply a dummy variable that is one, in case there has been any incident, be it violent or non-violent, in a district in a given time period. I will refer to this as the incidence of violence. The second is the number of incidents that occurred in a given time period, which is a broader measure. I will refer to this as the violence intensity. The resulting dataset is a balanced district level panel covering the time period from 2000 to 2012.[25]

The spatial unit used in this chapter is an Indian district. I use district definitions from the 2001 census; since then, many districts have been carved out of existing ones or renamed. I map these to the 2001 district boundaries. I study conflict in the whole of India, excluding the Kashmir region. Out of this region, there are 222 districts that experience variation in conflict intensity over the sample period. 130 districts are classified as being left-wing extremist affected by the Ministry of Home Affairs between 2000 and 2005. Of the remaining 92 districts, 45 are located in the North East. The remainder are districts spread across the whole of India. Districts classified as being under left-wing extremist influence account for the bulk of 52% of all conflict events recorded. The states Assam and Manipur where insurgency movements have close ties with the Naxalites or where there is significant Naxalite presence, account for 41% of all other conflict events.

Appendix 1.A.4 provides an example of how the algorithm constructs an incident count based on individual newspaper clippings, while Appendix 1.A.5 compares the resulting dataset to the Global Terrorism Database. The insight is that the semi-automatically retrieved dataset performs extremely well, compared with other violence datasets and even with manually coded data drawn from the same newspaper clippings. In addition to this newly created conflict data, I complement this chapter by studying crime as well. A particular emphasis is on violent crimes and crimes against public order. These data are collected at the district level by the local police and reported to the National Crime Records Bureau. In order to establish the link between Monsoon rainfall and agricultural incomes, as a proxy for the livelihoods, I collect wage- and agricultural output data at the district level. The next section

[24] This is problematic as a common problem in India is that there is a multitude of spelling variations for similar district names. This could result in significant coding errors or omissions.

[25] The data collection through the SATP began only in mid 2000.

describes this data.

Agricultural Production and Wages In order to test whether NREGA had an impact on the relationship between Monsoon rainfall and agricultural wages or agricultural production, I constructed two datasets to measure these. I use Agricultural Wage Data from the Agricultural Wages in India (AWI) series which has been published by the Indian Ministry of Agriculture since 1951. It is unique in offering monthly wage rates by district (sometimes even containing multiple locations per district), and separate wage series for several categories of labour and by gender. The quality of the data is very poor however, with a large number of observations being missing or simply flat wages being reported throughout. In order to increase the signal to noise ratio, I average the data to generate an annual wage series. I detail some of the issues with this dataset in appendix 1.A.9. The resulting dataset is an unbalanced annual panel dataset at the district level covering the time period between 2001-2010. More reliably measured is agricultural production. I use data on annual district level production collected and published by the Directorate of Economics and Statistics with the Ministry of Agriculture. This data is reported at the financial year level, which ranges from April- to March in the subsequent calendar year. I match the year to the calendar year to ensure the largest period of overlap.[26] For every district, I only consider crops that have been consistently planted on at least 1000 hectares for the period that the state reports data. I use state-level harvest prices to construct a district level measure of agricultural output. The resulting dataset is an unbalanced panel dataset covering the period from 2000-2009. The exogenous variation in this chapter comes from a measurement of the intensity of rainfall during the Monsoon season. As rainfall reporting from ground measurements is potentially problematic in developing countries, I invoke some novel remote-sensed rainfall dataset.

Rainfall data This chapter uses data from the Tropical Rainfall Measuring Mission (TRMM) satellite, which is jointly operated by the National Aeronautics and Space Administration (NASA) and the Japan Aerospace and Exploration Agency (JAXA). The satellite carries a set of five instruments to construct gridded rainfall rates at very high spatial and temporal resolution. Due to the high spatial and temporal resolution it is providing more consistent rainfall estimates than any other available ground based observations and is considered the highest quality remote sensed rainfall dataset with global coverage that is currently available (see Li et al., 2012 and Huffman et al., 2007). Its adequacy to pick up the spatial heterogeneity in precipitation has been highlighted and verified in the Indian context by Rahman and Sengupta (2007), who have shown that it outperforms e.g. the Global Precipitation Climatology Centre (GPCC) rain gauge analysis data that has been used extensively in economics research.[27] The data has the advantage of using a consistent methodology and most

[26]This data is available on `http://apy.dacnet.nic.in/cps.aspx`, accessed 14.08.2013.

[27]For example by Miguel et al. (2004), Ferrara and Harari (2013) and Kudamatsu et al. (2014). My results are robust to using the GPCC data (Schneider et al. (2011)).

importantly, time invariant sources of input data derived from the instruments that are carried by the satellites. This could be important as in appendix 1.A.6 I present some evidence, suggesting that the number of ground based measurements that feed into the GPCC could be systematically varying with levels of violence. The daily rainfall from 2000 to 2012 comes at a fine spatial resolution of 0.25 by 0.25 degree grid-cell size, which is converted into overall monthly rainfall in mm. For the identification, I will focus on the Monsoon season rainfall, which I define based on the principal crops grown using the state specific Indian crop calendar.[28] The Monsoon period varies from state to state as the typical onset dates are early May for the north east of India, while the onset may be as late as late June for central India. For most states, the Monsoon-period ranges from June to September. The Monsoon period rainfall accounts for at least 70% of the annual rainfall and most of the volatility.

The last important piece for the empirical analysis consists of NREGA take-up and participation data. As I argue NREGA provides insurance against income shocks, I study how program participation and expenditure is affected by rainfall shocks.

NREGA Participation Data I use the NREGA participation data derived from the so-called Monthly Progress Reports (MPR) from before 2011 and from the Management Information System (MIS) from 2011 onwards. The key variables I study are extensive margin participation as the share of households in a district that participate under NREGA in a given financial year, the days worked per household and the total person days generated. I match the financial year ranging from April to March to the nearest calendar year to be consistent throughout. I also obtained data on the number and total cost of ongoing projects, where I classify projects for road construction and land development specifically.[29] I study three major margins of NREGA take-up. Firstly, extensive margin participation as the share of households in a district who are employed in a year. Secondly, intensive margin participation as the log of the number of days worked per household. Last but not least I consider total expenditures per district and financial year. Table 1.1 presents some summary statistics suggesting that NREGA participation is most widespread in districts that received NREGA in earlier phases. There participation is almost twice as high, around 40% of households participate. The expenditure per capita in districts that received the program the earliest is also significantly higher, standing out with around 480 INR per year and person in districts in the first two phases, compared to only 247 INR per capita and year for the richest districts that received the program in the last round.

I the next section, I present the empirical strategy before presenting the core results.

[28] In particular the key reference is the crop specific calendar maintained by the Indian Food Security Mission, available via `http://nfsm.gov.in/nfsmmis/RPT/CalenderReport.aspx`, accessed on 12.05.2013.

[29] Refer to Appendix 1.A.10 for further discussion of the available NREGA participation data.

1.3 Empirical Strategy

The empirical strategy of this chapter consists of linking three variables: monsoon rainfall, income and conflict. I study these relationships before and after the introduction of the NREGA workfare program.

In the first step I analyze the effect of Monsoon season rainfall on agricultural output and wages. I do so by estimating the relationship between agricultural output, wages and Monsoon season rainfall using an unbalanced panel covering the time-period before NREGA was introduced. This ensures that the estimates are not affected by the impact that NREGA may have. I focus on Monsoon season rainfall. The rain falling in this season is most important for India's agricultural productivity. The estimating equation is:

$$log(y_{dprt}) = a_d + b_{prt} + \theta \times log(R_{dprt}) + \epsilon_{dprt} \quad (1.1)$$

where the indices d stands for district, p stands for NREGA implementation phase ranging from 1-3, r indicates region and t indicates time. The regressions include two sets of fixed effects. First, there are district fixed effects a_d which absorb any time-invariant district characteristics that may explain levels of agricultural productivity. These are characteristics, such as soil characteristics, elevation or terrain ruggedness. The second set of fixed effects are time-effects b_{prt}. These time fixed effects are region- and NREGA implementation phase specific and thus remove region specific time shocks that affect districts that received NREGA in the first round differentially from districts that received the program in rounds two and three.[30] The coefficient of interest is θ. This coefficient measures the elasticity between Monsoon season rainfall and agricultural wages or GDP. For the specifications with agricultural wages, I include a set of state by NREGA phase specific linear time trends.

In the second step I empirically establish the link between Monsoon season rainfall and conflict. I study two margins: conflict incidence and conflict intensity. Conflict incidence is an indicator of whether there was any conflict event reported in a district and year. Conflict intensity is the number of all incidences reported in a district and year. The dataset is a balanced panel covering all mainland Indian districts with the exception of Kashmir.[31] The specification using conflict incidence is a linear probability model with the estimating specification being:

$$A_{dprt} = a_d + b_{prt} + \eta \times log(R_{dpr,t-1}) + \epsilon_{dprt} \quad (1.2)$$

The fixed effects are as before. The coefficient that measures the link between conflict incidence and Monsoon rain is η. This coefficient is interpreted as the rate

[30] I define three regions: the North-East, comprised of Arunachal Pradesh, Assam, Manipur, Meghalaya, Mizoram, Nagaland, Tripura and Sikkim, the Naxalite Red Corridor, formed by the states Andhra Pradesh, Bihar, Chhattisgarh, Jharkhand, Karnataka, Maharashtra, Orissa, and West Bengal, and a third region comprised of the remaining states consisting of Gujarat, Himachal Pradesh, Haryana, Kerala, Punjab, Rajasthan Tamil Nadu and Uttaranchal.

[31] I exclude Kashmir as this conflict exhibits strong inter-state dimensions.

of change by which changes in Monsoon season rainfall translate into changes in conflict incidence.

The choice of the empirical design follows closely the existing literature on conflict in India, which has found a lagged effect of income or proxies of income on the intensity of the Naxalite conflict (see Vanden Eynde, 2011 or Gawande et al., 2012). The choice of timing can be rationalized on grounds of the cycles of agricultural activity. Periods of peak labor demand are the planting season at the onset of the Monsoon, when e.g. rice plant seedlings are transplanted to the field and the harvest season that stretches from October to December (see Figure 1.12). Incomes are only realized at the end of the year. This holds true particularly for the self-employed farmers, which account for 58% of rural employment across India.[32] A good Monsoon creates the chance for a second crop to be grown during the dry season between November to May in the subsequent year. All these imply that household income (and expectations) at the beginning of each calendar year depend strongly on the previous Monsoon season rains.

The second empirical specification studies the intensity of violence using a conditional fixed effect Poisson model as in Silva and Tenreyro (2006). This type of model accounts for the count nature of the number of conflict events as dependent variable.[33] The specification is:

$$\mathbb{E}(A_{dprt}) = \alpha_d \exp\left(b_{prt} + \eta \times log(R_{dpr,t-1}) + \epsilon_{dprt}\right) \tag{1.3}$$

The coefficient of interest is again η which is interpreted as an elasticity. The specification is estimated using a balanced panel for the whole of India. However, districts that do not experience any variation in the dependent variable do not contribute to the estimation of the coefficients. For that reason, the tables will present a varying number of districts across specifications. Specifications 1.2 and 1.3 are reduced forms. It is also possible to perform an instrumental variables analysis to establish the causal link between Monsoon rainfall and conflict.[34] The use of lagged rainfall as an instrument for lagged agricultural output alleviates some direct concerns about the validity of the instrument. However, I focus mainly on the reduced form in this analysis.[35]

The above empirical analysis will be presented in one condensed table presenting the impact of Monsoon rainfall on agricultural output per capita, wages and conflict together. The main hypothesis is that before the introduction of NREGA, there is a strong relation between lagged Monsoon season rainfall and conflict. The subse-

[32]See Planning Commission (2005).

[33]I use a Pseudo Maximum Likelihood Poisson (PPML) estimator as implemented by Silva and Tenreyro (2006). This estimator overcomes some of the numerical problems in common implementations in statistical packages such as Stata (see Silva, 2011). The PPML estimator does not require the data to have equi-dispersion. It is consistent, so long as the conditional mean is correctly specified. The estimator is even optimal if the conditional variance is *proportional* to the mean, hence over dispersion is not an issue.

[34]See Table 1.14. The table also presents results for OLS and Negative Binomials as main specification.

[35]An additional problem that arises in particular for the post-NREGA period is the lack of balance in the panel on agricultural output and wages. The data stop in 2009 or 2010 respectively, so there is missing data for many districts that experience conflict both before and after.

quent part will study how the introduction of NREGA lead to a moderation of this relationship.

I do this in the same empirical setup by adding an interaction term to the previous specifications, where Monsoon rainfall is interacted with an indicator $T_{dprt} = 1$ in case a district d receives NREGA from point t onwards. Note that eventually all districts would receive NREGA.

$$T_{dprt} = \begin{cases} 1 & \text{if NREGA available in district } d \text{ from time } t \text{ onwards,} \\ 0 & \text{else.} \end{cases}$$

The identifying assumption for this chapter is that the timing of the introduction of NREGA was not correlated with other omitted variables that could affect the relationship between Monsoon rain and conflict. This identifying assumption is valid even if the timing of the introduction of NREGA was endogenous to levels of violence.

I proceed by studying the moderating effect of NREGA following the same steps. I first focus on the relationship between Monsoon rain and agricultural output and wages. The specification with agricultural output and wages becomes:

$$log(y_{dprt}) = a_d + b_{prt} + \eta \times log(R_{dpr,t}) + \theta \times T_{dprt} \times R_{dpr,t} + \epsilon_{dprt} \tag{1.4}$$

Note that the simple treatment dummy T_{dprt} is perfectly collinear with the region-by NREGA phase time fixed effects. This specification does not attempt to estimate a level effect due to the endogeneity of the roll out.[36] The specification asks whether the way rainfall translates into agricultural wages or GDP changes with the introduction of NREGA. As NREGA employment is available when households demand it, it provides an alternative source of income for households. The stabilization of household incomes should materialize also in stabilized agricultural wages as NREGA effectively creates a wage floor and can directly stabilize labor markets. Bringing it back to the regression the interest is on the joint significance of the estimated coefficients $\hat{\eta} - \hat{\theta}$.

While the actual employment provision under NREGA make it reasonable for there to be a direct effect on wage rates as determined by the agricultural labor market, it is not clear if and whether there should be an impact on the relationship between Monsoon rain and agricultural output in the short run as well.[37] In the longer run it is well possible that NREGA makes agricultural production less sensitive to Monsoon season rainfall as NREGA aims to create infrastructure, e.g. for drought proofing or micro-irrigation.

Turning to studying the relationship between Monsoon rainfall shocks and con-

[36]In Appendix section 1.A.2, I explore the level effect of the program as well; however, the identification of a level effect is much more difficult.

[37]NREGA could create income that is used for fertilizer and other agricultural inputs in the short run, which could improve agricultural output. It is questionable however, whether the use of such additional inputs could directly weaken the link between Monsoon rainfall and output.

flict, the empirical specifications are analogous:

$$A_{dprt} = a_d + b_{prt} + \eta \times log(R_{dpr,t-1}) + \gamma \times T_{dprt} \times log(R_{dpr,t-1}) + \epsilon_{dprt} \tag{1.5}$$

$$\mathbb{E}(A_{dprt}) = \delta_d \exp\left(b_{prt} + \eta \times log(R_{dpr,t-1}) + \gamma \times T_{dprt} \times log(R_{dpr,t-1}) + \epsilon_{dprt}\right) \tag{1.6}$$

Its easiest to think of the Monsoon rain and conflict relationship after the introduction of NREGA as an inward rotation: the relationship between Monsoon rainfall and conflict becomes less steep. This means that after NREGA is introduced, Monsoon shocks of similar magnitude may still translate into conflict, but by a smaller amount in comparison to before NREGA. In the extreme case, the relationship between Monsoon rainfall and conflict becomes entirely flat, suggesting that Monsoon shocks cease to have an effect on conflict. A common concern with difference-in-difference type estimators is to ensure that common trends hold. I verify this by transforming the treatment variable into a district-specific time variable measuring the time to the introduction of NREGA. This results in fifteen time-steps.[38] I then estimate a separate coefficient for the Monsoon rainfall and conflict relationship for each year. The specifications are:

$$A_{dprt} = a_d + b_{prt} + b_{pct} + \sum_{t=1}^{15} \eta_t \times log(R_{dpr,t-1}) + \epsilon_{dprt})$$

$$\mathbb{E}(A_{dprt}) = \delta_d \exp\left(b_{prt} + \sum_{t=1}^{15} \eta_t \times log(R_{dpr,t-1}) + \epsilon_{dprt}\right)$$

The estimated coefficients η_t can be plotted out along with confidence bounds. The expected coefficient patterns are such that the η_t's are negative for the period before NREGA and become insignificant for the period after NREGA was introduced.

The argument of this chapter is that NREGA breaks the link between Monsoon rainfall and conflict due to NREGA's moderating impact on household income. Household income becomes less responsive to Monsoon rainfall shocks because households can earn income through NREGA when facing adverse conditions. The study of agricultural wages and output is already one indication. The stabilization of agricultural wages should happen through increased NREGA participation. Whether NREGA functions as insurance is an empirical question. I study NREGA utilization along two margins: overall program expenses and total person days of employment provided in a district over time. The latter is broken up into extensive- and intensive margin participation. Since NREGA is available on a per-household level, participation is measured at that level. Extensive margin participation is mea-

[38] Note that this is longer, even though the panel only ranges from 2000 - 2012. The reason is simple: districts in the first phase have a shorter pre-treatment period, but a longer post-treatment period, while for districts in later phases, this is reversed.

sured as the share of households in a district that participate in the program, while intensive margin participation measures the days per household worked. Let these measures be denoted by P_{dprt}, the specification estimated is:

$$P_{dprt} = \delta_{a_k} + b_{prt} + \phi \times log(R_{dpr,t-1}) + \epsilon_{dprt} \tag{1.7}$$

The set of fixed effects used are similar: b_{prt} are region and NREGA implementation phase specific time-fixed effects, δ_{a_k} are district fixed effects that I allow to change for the period from 2011 onwards.[39] The coefficient of interest from these regressions is ϕ. I expect this coefficient to be negative, which indicates that good Monsoon realizations are correlated with lower levels of NREGA participation. This highlights how NREGA take-up is responsive to Monsoon shocks. This allows NREGA to function as a stabilizer to incomes following adverse shocks, breaking the direct relationship between local Monsoon shocks and income. This is key to explaining why the link between Monsoon rainfall, wages and conflict disappears. The next section presents the results from this analysis and highlights that they are robust to various ways of studying the data.

1.4 Results

The results are presented in the same sequence as presented in the empirical strategy. I proceed by establishing the relationships between Monsoon season rainfall, agricultural output, wages and conflict before the introduction of NREGA. In the second step, I present results pertaining to the whole period to study how the introduction of NREGA has affected the relationship between Monsoon rainfall, agricultural output, wages and conflict.

1.4.1 Before NREGA: Agriculture, Wages and Conflict

The first section presents the results pertaining to the relationship between Monsoon rainfall, agricultural output, wages and conflict. The regression results are presented in Table 1.2. Column (1) presents the results using agricultural output per capita as dependent variable. The coefficient on contemporaneous Monsoon rainfall is interpreted as an elasticity indicating that a 1% deficient Monsoon reduces agricultural output per capita by 0.362%. This indicates a strong dependence of agricultural production on Monsoon rainfall.[40] The strong relationship between agricultural output and Monsoon season rainfall is particularly relevant for self-employed farmers as they are directly hit by adverse shocks. Self-employment in agriculture accounts for at least 58% of rural employment (see Planning Commission, 2005) which establishes

[39] This is because the underlying data-sources for the NREGA participation data changes from 2011 onwards, which creates jumps in the NREGA participation data that are specific to each district. Refer to appendix 1.A.10 for more details.

[40] This finding is in line with a long standing literature that has studied the relationship between Monsoon rainfall and the welfare of rural households in the Indian context (see e.g. Rosenzweig and Binswanger (1993), Burgess et al. (2011), Cole et al. (2010))

a direct link between agricultural production and household incomes. In appendix Table 1.12 I show that the relationship between Monsoon rainfall and output is driven by production of staple crops.

The second source of income in rural areas is casual wage employment in the agricultural sector. A second step is to analyze whether agricultural productivity shocks in form of rainfall shocks translate into lower wages. This generates a second margin along which productivity shocks can depress household incomes. This is studied using data on agricultural wages in the second column of table 1.2. Again the coefficient is interpreted as elasticity. The effect is small but significant: a 1% decrease in Monsoon rainfall decreases agricultural wages by 0.058%. The effects of Monsoon rainfall on the agricultural labor market are highly complex, yet, this finding relates well with the existing literature (see Jayachandran (2006)). In appendix Table 1.13 I study wages in the planting- relative to the harvest season. The two findings are two important pieces of information that establish links between rural incomes and rainfall variation.

In the next step I address the question on whether Monsoon rain variation explains conflict. The results are presented in columns (3) and (4). Column (3) presents a linear probability model studying the incidence of conflict in a given year. The coefficient is negative and significant: a good Monsoon translates into a lower probability of conflict in a district. A 20% deficiency would increase the probability of conflict by 0.7 percentage points. Given that 17.6 % of district years exhibit conflict, this is an increase by 3.9%. Column (4) presents the results from the Poisson regression. Note that the regression is estimated using the whole balanced panel, however, only 141 districts provide time-variation in the dependent variable and thus, contribute to the estimation of the coefficients. The coefficient is interpreted as elasticity, indicating that a 1% reduction in Monsoon rainfall translates into an increase in conflict by 0.846%. This coefficient compares very well with estimates of previous studies, in particular, Vanden Eynde (2011) and Gawande et al. (2012) who study the Maoist conflict. This establishes a direct link between Monsoon season rainfall and conflict.[41]

The three relationships between Monsoon rainfall, agricultural output, wages and conflict can also be studied in a non-parametric manner to highlight non-linearities (see Hsiang et al. (2013)). The idea of this non-parametric approach is to obtain local estimates of the relationship studied and to visually display them.[42] Panel A on top

[41] In appendix tables 1.14 I perform a set or robustness checks highlighting that I obtain similar results using an instrumental variables approach and that the results are robust to the choice of empirical model.

[42] The procedure has two steps. First, the data is demeaned by the fixed effects. This ensures that the scales are identical. A loess regression of the residuals of the weather variable and the dependent variable of interest is then estimated repeatedly using a bootstrapping procedure. The residuals for the horizontal axis, in this case the residuals for (lagged) Monsoon-rainfall, are subdivided into a set of 200 grid points. Each bootstrapped regression is evaluated at the grid-points for the horizontal-axis. This results in a set of fitted values for each grid point along the horizontal axis. In the second step the fitted values are plotted. For each horizontal grid point, a kernel density is estimated. The colouring is related to two things. First, the overall color intensity at each horizontal axis grid point is related to the overall mass of data that accrues there. This color is then streched out vertically in relation to the density of the fitted values. 95% confidence bounds are plotted as dashed lines.

in Figure 1.3 presents the relationship between Monsoon rain and agricultural output per capita. This relationship seems fairly monotone. The second graph depicts the relationship between wages and Monsoon rainfall; again, a fairly monotone relationship emerges, as indicated by the linear fit. The conflict non-parametric exhibits some non-linear patterns. The estimated coefficients are positive for Monsoon rain deficits, indicating that deficient Monsoon is correlated with higher levels of violence. Large and positive Monsoon rain deviations are generally correlated with less conflict; the relationship is appears to be bending back up suggesting that extremely positive rainfalls can also induce conflict. This is not unsurprising: the agricultural output non-parametric seems to be bending down at very high Monsoon realisations.[43] In the next section I study how these relationships are affected by the introduction of NREGA.

1.4.2 After NREGA: Moderation of Monsoon and Conflict Relationship

The focus of this chapter is how NREGA changes the slope linking Monsoon rainfall and conflict. The instrumental link through which moderation is achieved is through stabilizing agricultural incomes. I proceed by presenting results on the relationship between Monsoon rainfall, agricultural output, wages and conflict as in the previous section. The only addition here is that I allow the slope linking Monsoon rainfall with each of these three variables to be different after the introduction of NREGA as discussed in the empirical strategy.

The results are presented in Table 1.3 in a condensed form. Column (1) presents the results using agricultural output per capita as dependent variable. The focus is on the relationship between the coefficient on Monsoon rainfall with the coefficient on the interaction between the NREGA treatment dummy and Monsoon rainfall. While the former is positive and significant, the latter is negative and insignificant at conventional levels. This suggests that the relationship between rainfall and agricultural output per capita has not changed, at least up to the year 2009 when the agricultural output data stops. This is not surprising. NREGA aims to produce sustainable local infrastructure which could, in the longer run, increase agricultural output and make it more resilient to Monsoon rainfall shocks. Nevertheless these should not have an immediate effect on this deeply structural relationship. It is thus not surprising that agricultural output is still very much a function of Monsoon rainfall. Column (2) presents the results for agricultural wages. NREGA is a major intervention in the agricultural labor market; in periods in which otherwise, agricultural wages would have been depressed due to an adverse weather shock, NREGA provides an outside option which should stabilize agricultural wages making them less responsive to rainfall shocks. The regression indicates this to be the case. The coefficient on the interaction term is positive and significant. The sum of the two coefficients is insignificant as indicated by the F-test. This implies that agricultural wages cease

[43] A general linear fit appears adequate given the data. The slight non-linearity for positive extremes, if anything, implies that I underestimate the steepness as the regression line is pulled back up.

to be a function of Monsoon rainfall following the introduction of NREGA, indicating that NREGA can stabilize agricultural wages for households who do not directly participate in the program through a general equilibrium effect. This finding is complementing the existing empirical research that has documented that NREGA lead to increased wage levels (Azam, 2011; Berg et al., 2012; Imbert and Papp, 2015; Zimmermann, 2012). The last two columns study the relationship between conflict and Monsoon rainfall. The coefficients in both columns indicate that the Monsoon rain conflict relationship for conflict incidence (column (3)) and conflict intensity (column (4)) has moderated dramatically. In order to assess the degree of moderation statistically, I perform an F-test on the joint significance of the two Monsoon rainfall coefficients. The F-test is insignificant with a p-value of 0.267. This suggests that there remains a negative relationship indicating that low Monsoon rainfall translates into conflict; the relationship is however, a lot weaker and statistically insignificant. This suggests that NREGA has completely removed the rainfall dependence of conflict. This relates well with findings in historical China, where the introduction of the drought resistant sweet potato has moderated the impact of weather shocks on peasant revolts (Jia, 2014).[44]

It is important to highlight that this finding is distinct from empirical approaches that aim to identify a level effect of NREGA on conflict levels. Poor districts and districts that experienced conflict received NREGA earlier, making districts that received NREGA later a poor counterfactual. My identification strategy steers clear of this concern. Nevertheless it becomes instructive to estimate level effects of the introduction of NREGA. In appendix 1.A.2 I estimate level effects using a simple difference-in-difference strategy. The estimated coefficients suggest that levels of conflict are 30%-50% lower. This estimated effect has two sources: first, there may be a direct level effect of NREGA as NREGA has lead to an increase in wage levels independent of weather shocks. In the context of classical opportunity cost based models of conflict, this can be seen as pushing out the participation constraint. However, the second margin is an insurance effect which prevents wages to drop in case a bad state is realized. The results on the level effect map well into the findings of Dasgupta et al. (2014). They use a difference-in-difference estimator to estimate the level effect of NREGA. The identifying variation this relies on is coming from the sequential roll out of NREGA and hence, is essentially coming from just two years of data. Khanna and Zimmermann (2013) obtain different results. They address the endogeneity of the roll out directly and use a fuzzy regression discontinuity design for identification. They reverse engineer the NREGA roll out algorithm to identify districts that were near the cutoff of being assigned into an earlier or later phase. Districts close to the cutoff serve as counterfactual. Their results indicate that the introduction of NREGA lead to an increase in levels of conflict in the short run. They argue that this is due to an increase in civilian collaboration with security forces. This induces more violence in the short run, but a moderation in conflict levels in the longer run. A concern with

[44]Related is the recent work by Kung and Ma (2014), who show that cultural norms seem to moderate the effect of adverse shocks on revolts in historical China as well.

the research design is that it is driven by relatively few districts near the cutoff that experience conflict variation. While this finding does not square with the insurance effect that I document in this chapter, it is at odds with the levels effect estimated.

In order to study the dynamics of this effect, an event study analysis can highlight the degree to which the relationship between Monsoon rainfall and conflict has evolved prior to the introduction of NREGA. The treatment variable is transformed into a district-specific time variable measuring the time to the introduction of NREGA. Note that this is longer, even though the panel only ranges from 2000 - 2012. The reason that districts in the first phase have a shorter pre-treatment period, but a longer post-treatment period, while for districts in later phases, this is reversed.

The result is a sequence of estimated coefficients $\hat{\eta}$ that are best visually presented, see Figure 1.2. I only plot coefficients that are estimated using districts from all three phases as otherwise the picture would be distorted due to a compositional effect. I will decompose the effect by NREGA implementation phase to study the effect for the different phases separately. The vertical line around zero refers to the point in time that NREGA was introduced. The dashed blue lines indicate the regression coefficients obtained from the baseline specification. The estimated coefficients suggest a consistently negative relationship between Monsoon rainfall and conflict before the introduction of NREGA. With the introduction of NREGA, the relationship disappears. It does so not instantaneously but gradually with the coefficient becoming insignificant only one after roughly two years after NREGA was introduced. The results suggest that common trends do hold. As the non-parametric analysis suggested that the relationship exhibits some non-linearities, it is important to explore this for the post-NREGA period as well. The results are presented in Panel B of Figure 1.3. Its best to directly compare Panel A from before NREGA with Panel B from after NREGA. Note that the scales are identical from both graphs, allowing a direct comparison. The apparent patterns are very similar to what the linear regressions indicate. The relationship between Monsoon rainfall and agricultural output per capita has remained very similar. However, the relationship between agricultural wages and Monsoon rainfall has rotated inward: following the introduction of NREGA, there appears no statistically significant relationship linking Monsoon rainfall and agricultural wages. This indicates that the stable outside option NREGA provides serves as a cushion for wages determined in the labor market. The relationship between Monsoon rainfall and conflict follows in the last column. The non-parametric paint a very suggestive picture: the relationship becomes flat. While in the period before NREGA, the relationship was weakly U-shaped indicating that weather extremes translate into conflict, following the introduction of NREGA, the relationship has disappeared. This indicates that since the introduction of NREGA, Monsoon rainfall variation ceases to have an effect on conflict.

Before exploring heterogeneity and mechanisms through which this moderation in the Monsoon conflict relationship was achieved, I perform various robustness checks to highlight that the core result is robust.

1.4.3 Robustness Checks

I categorize the robustness checks into three sets: first checks, involving adding more control variables and time varying fixed effects. The second set using different data sets or measures of Monsoon rain, while the last set includes some placebo tests.

The first set of robustness checks are exploring the robustness of the results to the inclusion of different sets of fixed-effects or adding controls. They are presented in Table 1.4. The first three columns study conflict incidence, while columns (4) - (7) study conflict intensity. I discuss the corresponding conflict incidence and conflict intensity specifications together. Column (1) and column (4) explore the robustness of the results to the inclusion of a set of time-invariant district characteristics from the 2001 census interacted with a set of year fixed effects. This is in the spirit of the semi-parametric difference in difference analysis as developed in Abadie (2005).[45] The estimated coefficient changes sign and becomes insignificant for the conflict incidence regressions; however, the results for the conflict intensity specification is robust. The fact that the linear probability model becomes insignificant is not too concerning, as the there is too little variation in the dependent variable. In column (2) and column (6) introduce time-varying district fixed effects where I estimate a separate set of district fixed effects for the period before and after the introduction of NREGA. The inclusion of such fixed effects would capture any district specific level effect that the introduction of NREGA may have had on conflict. While it is unlikely that the presence of the scheme triggers a dramatic conflict response, these fixed effects shut down this channel. The insight is that despite the inclusion of the fixed effects, the estimated coefficients change slightly for the conflict intensity specification; again, the interaction for the incidence regressions becomes insignificant, which is not surprising given that the time varying fixed effects effects explain most of the variation in the dependent variable. Column (3) and (7) introduce state by year fixed effects, while controlling for the NREGA treatment dummy. These fixed effects are absorbing a lot of the variation in Monsoon rainfall. As law- and order is in the domain of the Indian states, rather than the Union government, it is reassuring that the coefficient pattern remains similar. Lastly, column (4) studies only districts that have experienced conflict before the introduction of NREGA. The estimated effect for this subset of districts is very similar to the main specification. This ensures that the effect is not driven by an expansion of the geography of conflict that is correlated with the introduction of NREGA and Monsoon rainfall.

The second set of robustness checks concerns the measurement of local Monsoon shocks. I present robustness checks using different rainfall data or agricultural productivity proxies. These exercises are presented in Table 1.5. Again the analysis is

[45] The district characteristics are: terrain ruggedness, elevation, rural population share, tribal population share, scheduled caste share, illiteracy rates, household size, share of population younger than 6 years, population growth rate from census 1991 to 2001, gender gap, share of villages in district with primary school, share of villages in district with mud road approach, share of households in district that live in permanent housing, share of villages in district with primary health care facilities, share of villages with electric power, share of villages with a bus stop and the share of villages with a postal office.

presented on conflict incidence and intensity margins.

Columns (1) and (7) present results from a specification where district level rainfall is normalized by its standard deviation. This is problematic, as the 14 year time period that the TRMM data is available may be too short for a stable estimate of the mean volatility to emerge. Nevertheless, the results are very similar using this measure of rainfall. Columns (2) and (6) present results when studying a different rainfall dataset. I present results based on the GPCC data, which is based on ground level rain gauge measurements.[46] While I believe that the satellite based TRMM data is likely to be better, as rainfall reporting may be endogenous to conflict (see appendix 1.A.6) it is nevertheless reassuring that using a different dataset, I obtain very similar results. As both the GPCC data as well as the TRMM data has been processed using climatology algorithms, a general concern is "error propagation" (see Leung et al. (2005) and Burnicki et al. (2007)). As the raw data is transformed in the analytical process, simple small measurement errors may be propagated due to the mathematical and numerical transformations. This could generate spurious correlations that could affect the results. A simple way to address this is to instrument one rainfall dataset with the other one. This removes any systematic and non-systematic measurement error and ensures that the results that I obtain from the two different datasets are driven by the same underlying variation. The results are presented in columns (3) and (7) and are very similar to the core finding. Monsoon rainfall is only one proxy for local weather shocks. The analysis of the agricultural output relationship suggests that it is absorbing a lot of the variation in agricultural productivity. Nevertheless, it may make sense to explore a different measure. A candidate that has been proposed by Gawande et al. (2012) is a vegetation index measuring the degree of photosynthetic activity. Photosynthetic activity is driving plant growth and thus, agricultural output and is obviously affected not only by rainfall, but by other climatic conditions as well. Columns (4) and (8) present the results when studying the lagged Normalised Vegetation Index (NDVI) and its interaction with NREGA. The coefficient pattern is very similar as in the preferred specification. However, NDVI contains a lot less variation, which makes it more difficult for the coefficients to gain significance. Nevertheless, the results are very similar using this proxy. These results render me confident that I am genuinely picking up an effect of Monsoon rainfalls impact on conflict. I now turn to a third set of robustness checks, which explore different placebo tests. I perform two main placebo tests. The first one is simply following the standard approach in difference in difference type empirical setups: moving the reform to an earlier date. This allows me to test whether the relationship between Monsoon rainfall and conflict began to change already before NREGA was introduced. The previous exercise studying the effect of Monsoon rainfall over time already indicates that this is not the case. Moving the treatment to an earlier date should wash out the estimated NREGA effect. The second placebo check is testing whether rainfall outside the main grow-

[46]The data has been widely used in the economics literature and in the conflict literature in particular. Some prominent references are Ferrara and Harari (2013), Miguel et al. (2004), and Kudamatsu et al. (2014) and many others.

ing season has some effects on conflict. If Monsoon rainfall is capturing an effect on rural livelihoods, than rainfall outside the Monsoon season should not interact with NREGA in any systematic way explaining conflict. The results are presented in Table 1.6. Column (1)-(3) and columns (5)-(7) present results when moving the NREGA introduction three, two and one years ahead of time respectively. Moving NREGA treatment to an earlier date should wash out the estimated NREGA effect, since in truth the rainfall conflict relationship was present for the time period between the placebo NREGA introduction date and its true introduction date. Conversely, the closer the placebo is moved to the true NREGA introduction date, the more likely should we observe an NREGA effect for the placebo. This is the exact pattern that emerges in the data when studying the regression coefficients for conflict intensity in columns (5)-(7). The conflict incidence regressions have positive and significant coefficients on the Placebo interactions. If I restrict the analysis to only the period prior to the true NREGA introduction, the coefficients become insignificant and flip signs. This suggests that the positive and significant coefficients on the Placebo reform is driven by the period after the actual NREGA was introduced. Column (4) and column (8) present results when studying rainfall outside the Monsoon season. As expected, the coefficients are insignificant, though they share the coefficient pattern for conflict intensity. This is likely due to rainfall outside the Monsoon season being positive correlated with rainfall during the Monsoon. These exercises suggest that the introduction of NREGA fundamentally changed the relationship between Monsoon rainfall and conflict. In the next section I show that this also applies more generally to violent crimes, suggesting that NREGA also has indirect effects on the dynamics of crime more generally.

1.4.4 Effects on Crime

The existing literature highlights that there exists a relationship between weather and crime in the context of India. The common observation is that adverse rainfall shocks drive crimes against vulnerable populations, in particular crimes against populations from scheduled caste and scheduled tribes (see Sekhri and Storeygard (2013), Iyer and Topalova (2014)) and violent as well as property crimes in general (see Blakeslee and Fishman (2014)). I obtained the same district-level crime data used by these authors for the period 2002-2012 to study the impact of NREGA on the relationship between Monsoon rainfall on crime. I present results using the main specification as in the previous section with lagged Monsoon rainfall. As dependent variable I use the log of the number of crimes reported by broad categories.[47] The results are presented in Table 1.7. The coefficients suggest a moderation of the crime and rainfall relationship most prominently for violent crimes as well as for disruptions of public order. This crime category includes incidences of rioting and arson. These results map well into

[47]Poisson models as used in the rest of the chapter yield very similar results and are available upon request. I follow the categorisation of Iyer et al. (2012), Appendix 1. The only modification is that I do not consider only murders, but violent crimes includes the crime categories: murder, attempted murder, kidnapping and hurt.

my findings on insurgency related violence. Some of the events captured in conflict data may also be measured in the violent crime data. The results are driven by districts under left-wing extremist influence and districts that received NREGA in earlier phases. This maps very well into the results found using the conflict data.[48] A key concern with the crime data is that, especially for property crimes, data from cities is highly overrepresented. To illustrate: the state of West Bengal has about 91 million inhabitants, and the capital Calcutta has roughly 4.6 million inhabitants, accounting for about 5 % of the population. In the raw crime data, Calcutta accounts for 21% of all thefts in West Bengal in 2005 and about 3% of the murders. This suggest a strong urban bias for the crime reporting.[49] This is in itself not a problem but it could wash out the variation in the crime data attributable to Monsoon shocks if the relationship between crime and Monsoon shocks is different for cities relative to rural areas.[50] For violent crimes and crimes against public order, in particular murders and rioting, the reporting bias is weaker as these are highly visible crimes. A non-parametric analysis for the violent crime data suggests a pattern very similar to the one observed for the conflict data: before the introduction of NREGA, extremes in Monsoon rainfall in the preceding year translate into more violent crimes. Following the introduction of NREGA, this relationship is moderated significantly: the U-shape becomes flat (see Figure 1.6). Estimating the effect of lagged Monsoon rain over time suggests a similar response: negative coefficients for the period before NREGA was introduced and insignificant coefficients following NREGA introduction (see Figure 1.7).

In the next steps, I study some heterogeneity in the estimated NREGA effects. In particular, I study which NREGA implementation phases seem to drive the estimated effect. The expectation is that this is coming from districts that received the program earlier, as they are more vulnerable in general. Secondly, I highlight that the estimated effects are strongly driven by districts that are considered by the Ministry of Home Affairs to be affected by left-wing extremism. Lastly, I study who is a target of violence and how this changes with the introduction of NREGA.

[48]My results differ from those presented in Iyer and Topalova (2014) who do not find any systematic moderating effect of the NREGA on the crime- and rainfall relationship. This can be due to a set of differences in the two papers. Firstly, they use different rainfall data for their paper and control for contemporaneous, rather than lagged Monsoon rainfall. Secondly, specifications are estimated on a longer panel for which TRMM rainfall data is not available. This however, however, comes at the cost of loosing spatial variation in the rainfall measure as balancing the panel requires merging of districts to reflect district boundary changes. Last, but not least, the differences could also be due to the fact that they use different sets of fixed effects.

[49]Unfortunately, one can not "clean" the data by removing crime reported in cities falling into a district since the threshold city size was later changed to include only cities with at least 1 million inhabitants.

[50]There is some evidence that this is the case. An analysis for the 75 cities that report crime data for the period under study suggests an insignificant relationship between lagged Monsoon rainfall and violent crime, but a significant and positive relationship for property crimes. These results are available upon request from the author.

1.4.5 Heterogeneity

Effects by Implementation Phase NREGA was introduced in three distinct phases. This allows me to estimate effects by NREGA implementation phase. This becomes insightful in determining which districts seem to be driving the overall observed effects. I present results as before in studying the relationship between lagged Monsoon rainfall and conflict over time. The results are presented in Figure 1.4. The key observation is that the moderation of the Monsoon rainfall and conflict relationship is driven by districts which received NREGA in earlier phases. The estimated coefficient on Monsoon rainfall is negative before NREGA and becomes generally insignificant for the period after NREGA. This pattern is visible for districts that received NREGA in phases 1 and 2, while for districts in phase 3, there is no statistically discernible effect. As districts for phases 1 and 2 where poorer relative to the rest of India, the observed patterns are quite reasonable. Districts in earlier phases are most vulnerable. For districts that received NREGA in the third phase, the conflict incidence regressions suggest a statistically insignificant relationship, while the conflict intensity results suggest a mixed result. Around the NREGA introduction date, the Monsoon rainfall and conflict coefficients are negative and significant, but become insignificant towards the end of the sample period. This highlights that the effect is driven by poorer and more vulnerable districts. For districts that are, on average, richer, the results are less conclusive.

As it has been pointed out in the context discussion, NREGA was more likely to be introduced early in districts that are under the influence of Maoists. This makes it reasonable to study the distinct effect of NREGA specific to districts that have been classified of being under left-wing extremist influence: districts that form the Red Corridor.

Effects by Conflict The roll-out of NREGA suggests that districts that are under the influence of left-wing extremist were more likely to receive the program in earlier phases. This makes it important to study the impact of the scheme on the Monsoon rainfall and conflict relationship for districts under left-wing extremist influence relative to the rest of India. I present results plotting out the estimated coefficients of Monsoon rainfall over time constraining the analysis to districts classified as being under left-wing extremist influence and separately, for the rest of India. The results are presented in Figure 1.5. The top panel studies districts under left-wing extremist influence. These account for 130 of the 222 districts that experience conflict variation over time. The coefficient pattern strongly follows the suggested pattern, indicating that Monsoon rainfall is a predictor of conflict in these districts before the introduction of NREGA. This relationship has become a lot weaker since the introduction of NREGA. For non-left wing extremist affected districts that experience conflict, the results are less clear. The main variation is here coming from Assam and Manipur which captures 41% of all conflict events in the data. It is not clear whether one should study conflict in these states separately from the Maoist insurgency, as at least in Assam there is significant Maoist presence. For conflict incidence there is no relationship to begin with: there is hardly a district in these states that does experience no conflict. For conflict intensity, the coefficients are negative around the NREGA introduction date but then become insignificant in more recent years. An analysis of the agricultural output relationship for the North East is less straightforward as non-linearities in the Monsoon rainfall and output relationship are much more pronounced in the North East. A logarithmic transformation may not do the agricultural output relationship justice. Unfortunately, the production data for the North East is particularly thin, making it difficult to dig deeper into the underlying nature of the relationship. The results from this analysis suggest that the bulk of the NREGA effect is coming from districts that are classified of being under left-wing extremist influence by the Ministry of Home Affairs between 2000-2005.

The next section provides evidence that the NREGA moderation is driven by less (targeted) violence against civilians. This suggests that NREGA could help take civilians out of the line of fire.

NREGA and Targets of Violence The relationship linking Monsoon rainfall and conflict could be heterogeneous by who is the subject of violence. Vanden Eynde (2011) argues that civilians, facing an income shock, find themselves torn between becoming paid police informers. This comes at a cost, as insurgents react with more violence against civilians. As NREGA primarily stabilizes rural incomes, violence targeted against civilians may become less responsive to conflict. The conflict data allows a rough classification of the subject of violent activities into groups: civilians, security forces and terrorists. In Fetzer (2013) I highlight how this is done with the aid of humans to classify ambiguous cases. I proceed as before, except that now I change the dependent variable as being the number of conflict events in a year, where the subject of the event has been classified to be either civilian, security forces or insurgents. The results are presented in Table 1.8.

Columns (1) to (3) of table 1.8 performs the analysis of the NREGA effect studying conflict incidence, while columns (4) to (6) study intensity. The coefficient pattern that emerges that all types of violence are responsive to lagged Monsoon. However, the moderating effect of NREGA is most strongly seen for violence targeted against civilians in column (1) and (4). Violence against security forces in column (2) and (4) also exhibits a NREGA effect. The sum of the two coefficients actually is positive but insignificant, which could suggest that violence against security forces is starts to become positively correlated with Monsoon rainfall. The third column looks at incidences where the subject of the incidence was a terrorist. There appears to be only a weak moderating effect of NREGA. This evidence suggests that NREGA moderates the relationship between Monsoon rainfall and violence, with the bulk of that effect coming from less violence against civilians. This indicates that NREGA may help bring civilians out of the line of fire.

The key concern for identification is whether the NREGA treatment timing was correlated with other policies or events that could have been correlated with Monsoon rainfall and through that, affect the relationship between Monsoon rainfall and conflict. In Appendix 1.A.1, I rule out a whole range of alternative explanations for the observed moderation in the Monsoon rainfall and conflict relationship. The first set rules out other development schemes that could moderate the Monsoon rainfall and conflict relationship either directly or indirectly. This includes the Pradhan Mantri Gram Sadak Yojana Rural Road Construction Scheme (PMGSY) and the Integrated Action Plan, which channels additional funds into left-wing extremist affected districts. The second concern relates to a major military intervention "Operation Green Hunt", that has been underway since early 2011. Due to a lack of troop deployment data, I can not address this directly. Lastly, I rule out that the NREGA is not capturing a structurally different Monsoon rainfall and conflict relationship that is due to mining activity in a district. This could arise as the timing of the introduction of NREGA is correlated with a commodity price boom. Lastly, there could be an indirect way through which road construction under NREGA affects conflict. Fearon and Laitin (2003) highlight that guerrilla warfare thrives in places that are difficult to access by the state. Insurgency movements may have an incentive to prevent development of public infrastructure, in particular, roads in rural areas as they could lead to more government presence. That is to say, road construction could trigger more conflict. If road construction under NREGA was correlated with lagged Monsoon season rainfall, then this could explain part of the observed effects. In Appendix 1.A.3 I provide evidence that NREGA road construction may be correlated with conflict levels; I show that this effect is unlikely to go through Monsoon rainfall.

The chapter thus far has focused on how NREGA lead to an inward rotation of the Monsoon rainfall and conflict relationship. The implicit argument is that NREGA provides insurance against adverse income shocks. If the relationship between local Monsoon shocks and conflict was going through income, insulating income from adverse Monsoon shocks should break this relationship. The next section answers the question whether NREGA provides insurance and provides a rough quantification exercise for how much insurance is provided.

1.5 Does NREGA Provide Insurance?

The chapter has highlighted that the relationship between conflict and Monsoon shocks changes fundamentally with the introduction of NREGA. The implicit argument is that NREGA provides insurance, mitigating adverse weather shocks and thus, providing a cushion for incomes of rural households. By cushioning incomes, the link between income and conflict and some forms of crime is broken. This relationship needs to be empirically verified by studying how NREGA participation and expenditure responds following adverse Monsoon shocks. The hypothesis is that NREGA take-up is responsive to adverse Monsoon shocks: the slope linking Monsoon rainfall and NREGA participation is downward sloping indicating that positive rainfall realizations translate into lower participation as in these situations, the workers outside options are better which makes employment under NREGA at minimum wages less appealing. I study how adverse Monsoon shocks in the preceding growing season translate into increased NREGA participation. Since the harvest season is towards the end of the year (November-December), NREGA employment should only respond in the following spring. The results from the baseline analysis are presented in Table 1.9. The fist column measures the log of total expenditure in a financial year for ongoing projects. The elasticity is negative, indicating that good Monsoon rainfall in the preceding growing season translate into low expenditure. A 1% decrease in Monsoon rainfall increases expenditures by 0.257%.

Columns (2)-(4) study margins of participation. Overall take-up is measured as the total number of days worked under the scheme in a district and year. Again, the coefficient is negative: with good Monsoon, there is less need for NREGA employment. The coefficient indicates that a 1% drop in Monsoon in preceding growing season increases participation by 0.21%. This is high, but not unreasonable given the large share of self-employed farmers in rural India. The overall take-up effect is decomposed into extensive- and intensive margin participation in columns (3) and (4). The extensive margin measures the share of households who participate, while the intensive margin is the log of the number of days per household. Since the program is provided on a per-household level, this is the correct way to measure extensive margin participation. The measure also indicates a negative relationship. The intensive margin coefficient suggests that a 1% decrease in Monsoon rainfall, increases the number of days worked under NREGA by 0.12%. This suggests that a significant share of the overall observed participation response in column (2) is driven by extensive margin participation.

Columns (5) and (6) consider the heterogeneity in extensive margin participation: the relationship between lagged Monsoon rainfall and NREGA participation by implementation phase (column (5)) and by whether a district is classified as being under left-wing extremist influence according to the Ministry of Home Affairs. The pattern suggests that the effect of Monsoon rainfall on extensive margin participation is strongest for districts in phase 1. Furthermore, in left-wing extremist districts, the relationship between NREGA take-up and Monsoon rainfall is a lot stronger as well. This maps well into the findings from the previous section. The moderation was mostly driven by districts in the first phase and districts that are vulnerable to left-wing extremist activity.[51]

This suggests that NREGA does function as insurance: NREGA employment and program participation is higher, following local Monsoon shocks. This effect is driven by districts most vulnerable to conflict: districts that received NREGA in earlier phases and districts classified as being under left-wing extremist influence. The open question is *how much insurance* is provided.

[51] Similar to the previous analysis, I also present results from a non-parametric approach to highlight potential non-linearities in the response of NREGA participation with regards to Monsoon rainfall. These are presented in Figure 1.8.

To what extent do NREGA expenditures offset income losses due to adverse Monsoon shocks? Ideally this question would be answered using a household panel dataset. Such a dataset covering many parts of India does not exist to date.[52] The best I can do in this chapter is to provide a rough quantification exercise for how much insurance is provided exploiting variation in agricultural output value per capita due Monsoon variation at the district level. I relate this variation with NREGA expenditure per district. This allows a crude quantification exercise to measure the extent to which district output losses are compensated through increased NREGA expenditures. The question is by how much a INR 100 loss in agricultural output per capita is offset by increased NREGA expenditure per capita. I study the relationships in levels of expenditure and agricultural output per capita which allows direct comparison. As there are significant outliers in both data, I trim the bottom and top 1% of observations from both variables.[53]

The results are presented in Table 1.10. Column (1) is the agricultural production function that links Monsoon rainfall with output value in levels. The coefficient on Monsoon rainfall is a semi-elasticity indicating that a 10% increase in Monsoon rainfall increases nominal agricultural output value by INR 54.1. This can be interpreted as the first-stage for the analysis. Column (2) estimates how lagged Monsoon rainfall translates into increased NREGA expenditure per district. The coefficient indicates that a 10% reduction in Monsoon rainfall in the preceding growing season increases NREGA expenditure per capita by INR 10.1. This suggests that there is partial insurance coverage: every INR 10 loss is accommodated by an increase in NREGA expenditure by INR 1.86. This can be further refined. Column (3) I present results from an instrumental variables exercise. Lagged agricultural output is instrumented with lagged Monsoon rainfall. The coefficient suggests that a INR 100 loss in agricultural output due to Monsoon variation translates into increased NREGA expenditure by INR 30.1. If we study only labor expenses the coverage is INR 21.6. The combined results suggest that a significant share of the risk in crop cultivation due to local weather variation is offset by NREGA expenditures in a district. The estimates should be taken with a grain of salt. First, the number is not adjusted for marginal leakage in the program. Especially for the first years that NREGA was introduced, leakage rates from the program were significant as studied in Niehaus and Sukhtankar (2013a,b). This would suggest that the measure is likely to be an upper bound. On the other hand, the results in the previous section suggest that there is significant indirect insurance due to stabilization of agricultural wages. This has an indirect insurance effect that is not captured by the mere NREGA expenditures flowing into a district.

[52] The India Human Development Survey is a candidate dataset; the first round of interviews was completed in 2005 and the second round of data was collected in 2011-2012. Unfortunately, the data are not released until early 2015.

[53] Please consult appendix 1.A.8 for a discussion of how agricultural output value is scaled to ensure that it comes closer to the true agricultural output using district domestic product for 2000.

1.6 Conclusion

This chapter has studied the impact of social insurance on conflict in India. The existing literature studying conflict has devised various identification strategies to exploit arguably exogenous variation in incomes to study the relationship between income and conflict. The findings of this literature have a direct policy implication: any measure that helps insulate household incomes following adverse shocks should moderate the relationship between these exogenous productivity shocks and conflict. This chapter has taken up this question and evaluates the impact of the introduction of a public employment program established through the National Rural Employment Guarantee Act in India on the relationship between local Monsoon shocks and conflict and crime. The key findings suggest that the introduction of the public employment program has eliminated the link between Monsoon shocks and conflict and some forms of crime. Even after the introduction of the public employment program, productivity shocks continue to affect rural areas. However, with the public employment program in place, these shocks cease to translate into conflict. As conflict is a phenomena with a lot of persistence, removing the link between productivity shocks and conflict can lead to persistently lower levels of conflict. This seems to be the case in India. The insurance value delivered by the public employment scheme is significant. A simple quantification exercise suggests that roughly one third of district level income losses due to adverse Monsoon conditions are directly offset through increased expenditure under the program. This only captures the direct transfers. The indirect benefits due to reduction in the pass through of productivity shocks on wages are an added indirect insurance benefit.

The chapter has important implications for policy makers. Incomes in developing countries are much more volatile, leaving households exposed to a lot more risk in comparison to developed countries. Climate models suggest that erratic weather events could become even more pronounced. This has lead to concerns about increasing conflict in the future (see Hsiang et al. (2013)). Yet, many developing countries have not been able to devise policies to provide adequate protection. This chapter highlights that social insurance taking the form of a public employment program may be a policy that can be adopted in other developing countries as well.

Figures for Main Text

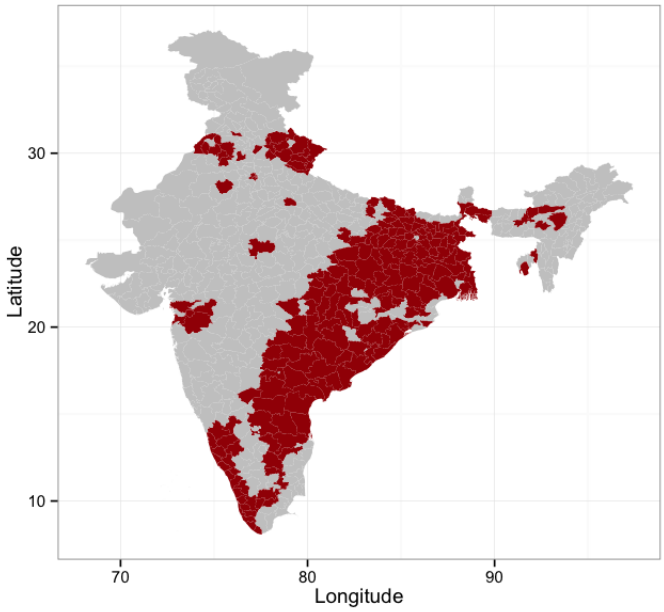

Figure 1.1: Districts Affected by Left-Wing Extremism According to Government of India.

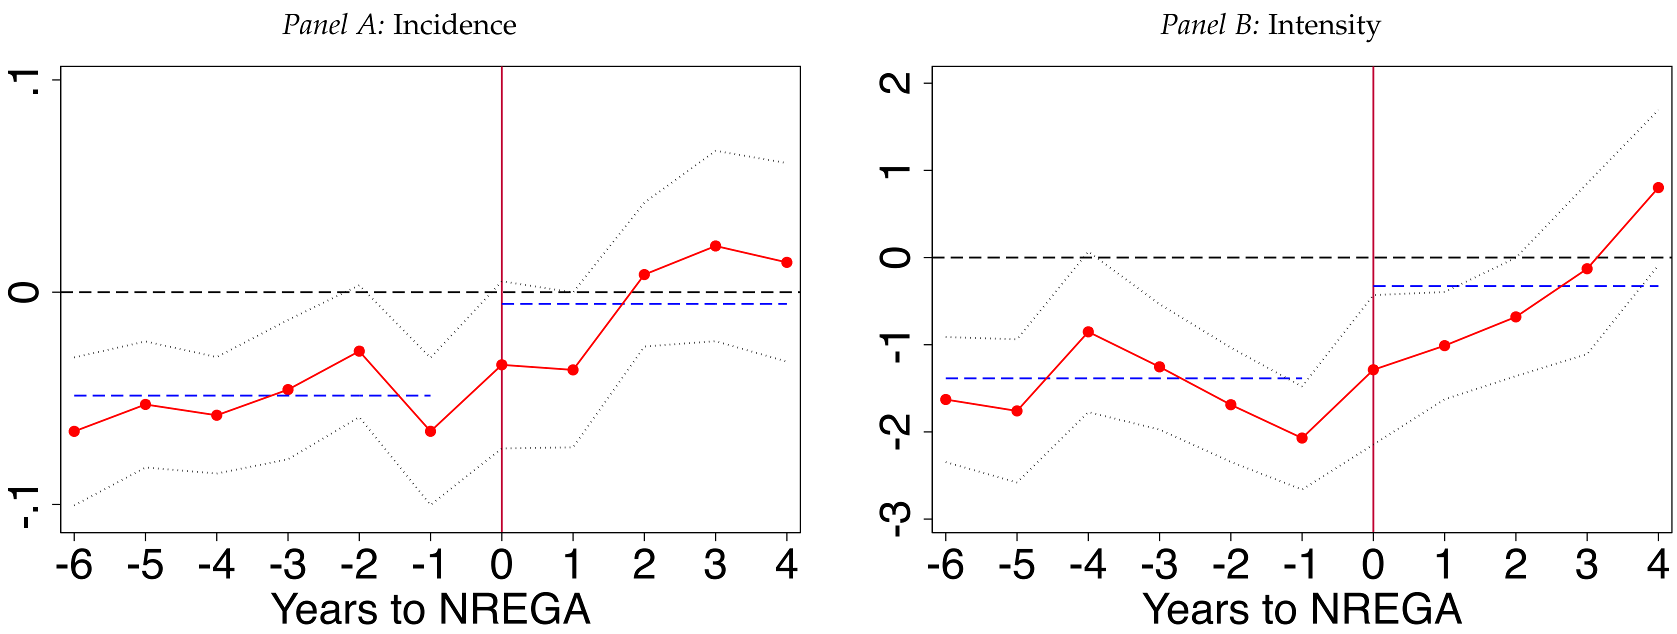

Figure 1.2: Effect of Monsoon Rain on Conflict Over Time. The vertical line indicates the NREGA introduction date. The blue dashed lines indicate the coefficients obtained from a simple regression interaction lagged Monsoon rainfall with the NREGA treatment indicator. The red line are each point estimates of the relationship between lagged Monsoon rainfall and conflict. 95% confidence bands are indicated as dotted black lines.

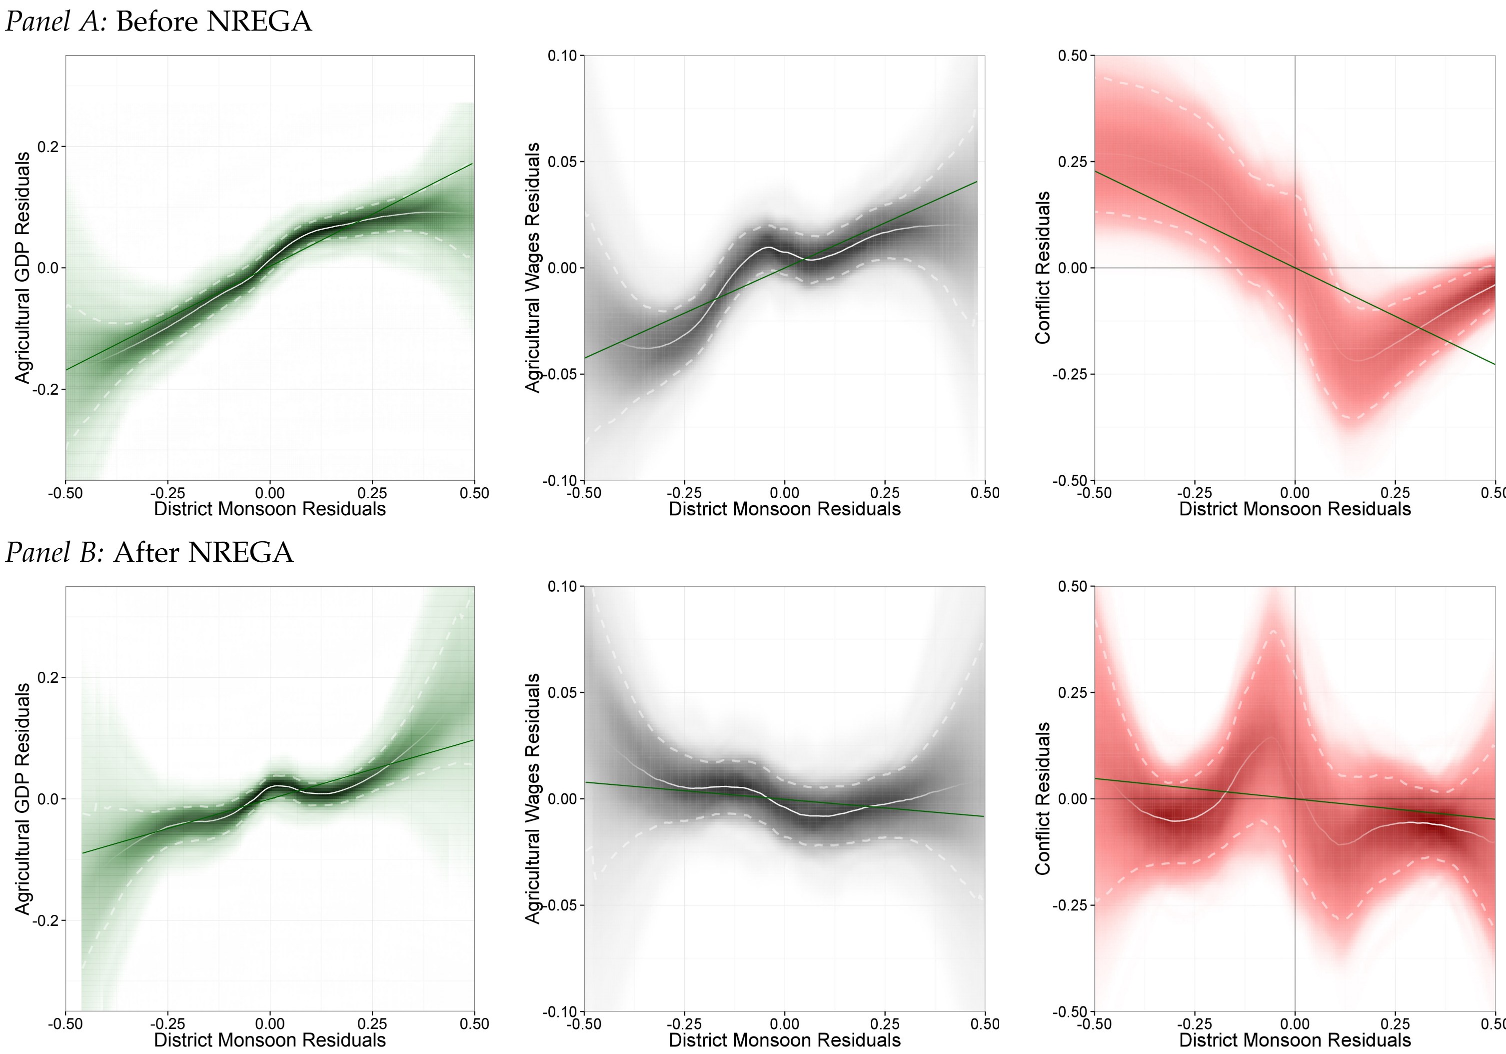

Figure 1.3: Non-Parametric Watercolor Regressions as in Hsiang et al. (2013): Effect of Monsoon Rain on agricultural output per Capita, Wages and Conflict Before and After Introduction of NREGA. 95% confidence bands are indicated as dashed lines. The color shading is related to the overall density of Monsoon rainfall realizations along the horizontal axis and to the density of fitted values from loess regressions along the vertical axis.

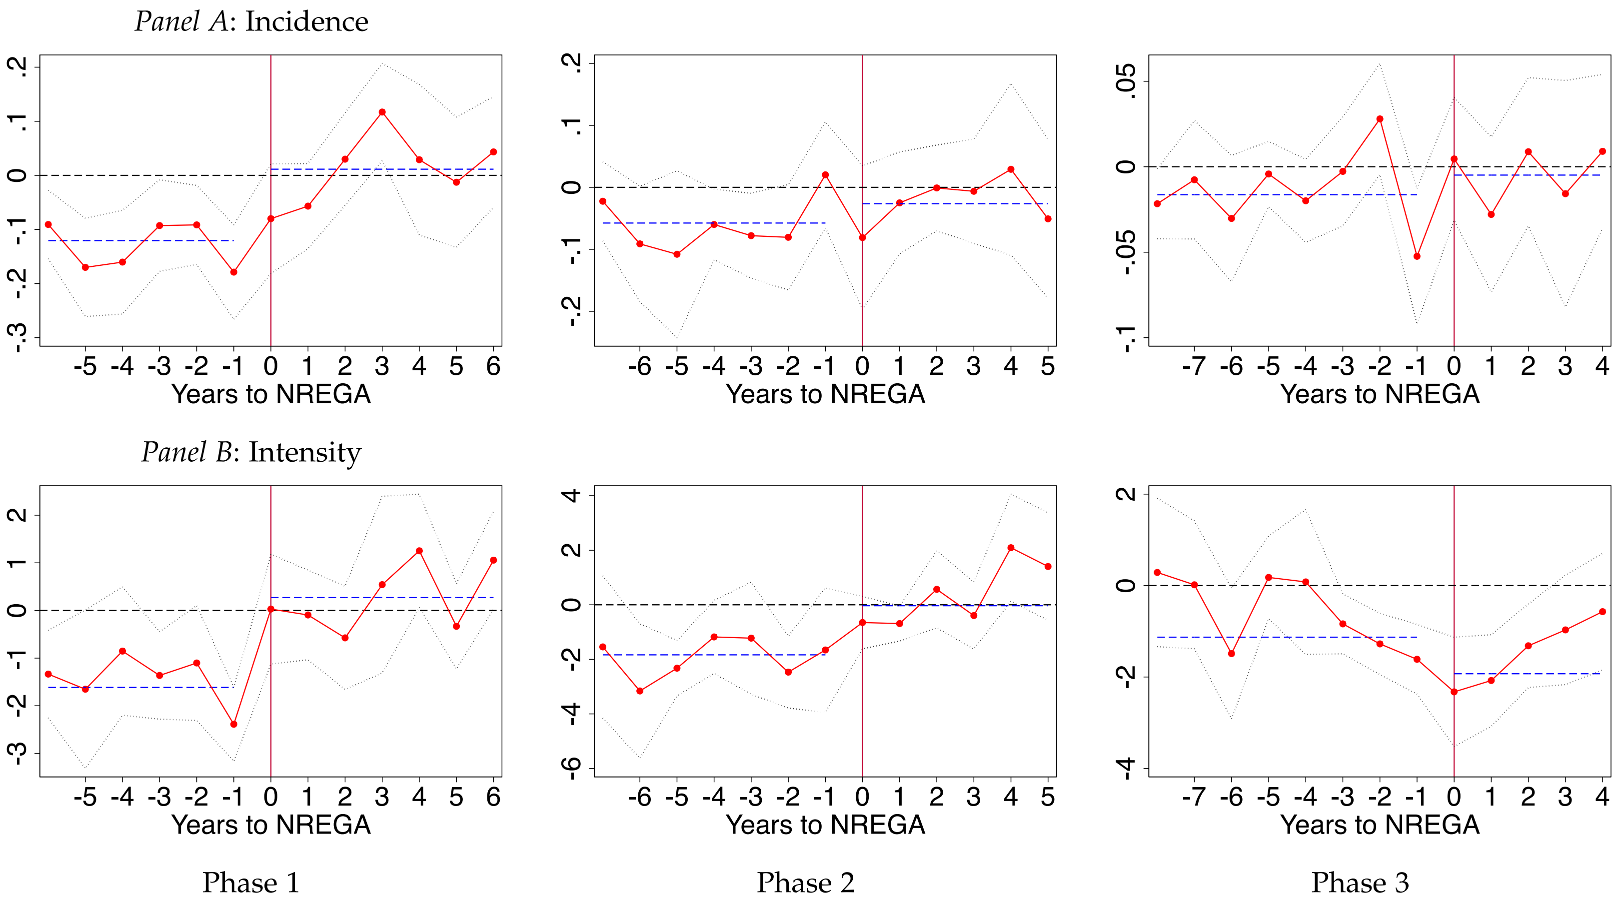

Figure 1.4: Effect of Monsoon Rain on Conflict Over Time By NREGA Implementation Phase. The vertical line indicates the NREGA introduction date. The blue dashed lines indicate the coefficients obtained from a regression interaction lagged Monsoon rainfall with the NREGA treatment indicator. The red line are each point estimates of the relationship between lagged Monsoon rainfall and conflict. 95% confidence bands are indicated as dotted black lines.

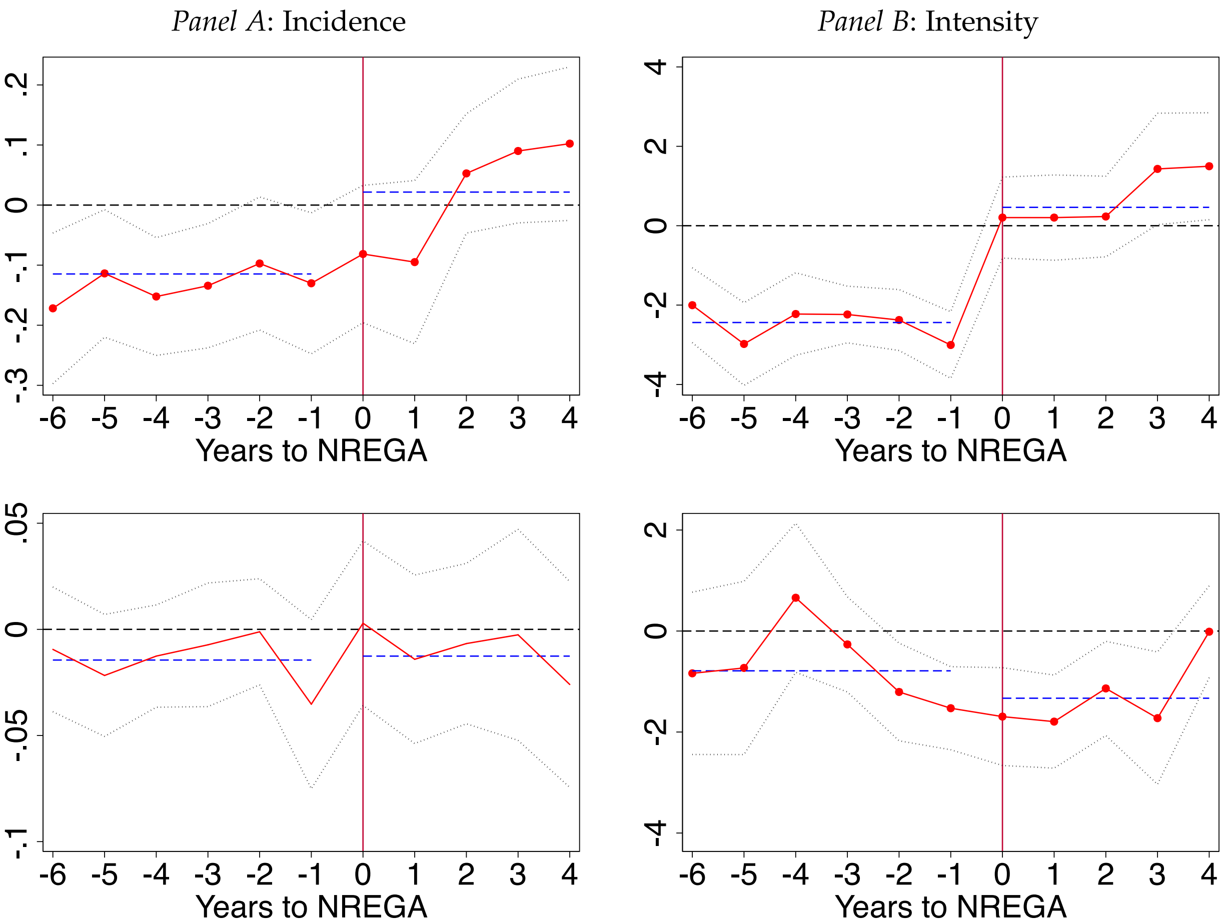

Figure 1.5: Effect of Monsoon Rain on Conflict Over Time in Red Corridor (top) and the Rest of India (bottom). The vertical line indicates the NREGA introduction date. The blue dashed lines indicate the coefficients obtained from a regression interaction lagged Monsoon rainfall with the NREGA treatment indicator. The red line are each point estimates of the relationship between lagged Monsoon rainfall and conflict. 95% confidence bands are indicated as dotted black lines.

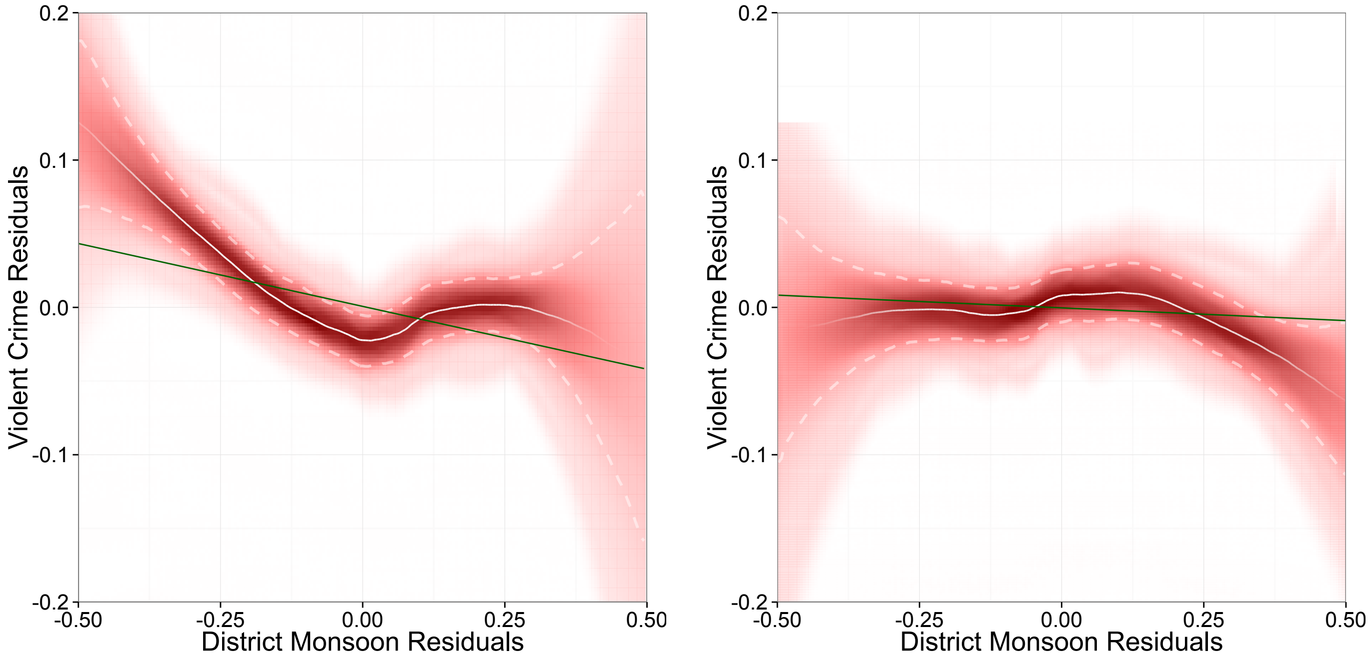

Figure 1.6: Non-Parametric Watercolor Regressions as in Hsiang et al. (2013): Relationship between lagged Monsoon rainfall and violent crime before (left) and after (right) the introduction of NREGA. 95% confidence bands are indicated as dashed lines. The color shading is related to the overall density of Monsoon rainfall realizations along the horizontal axis and to the density of fitted values from loess regressions along the vertical axis.

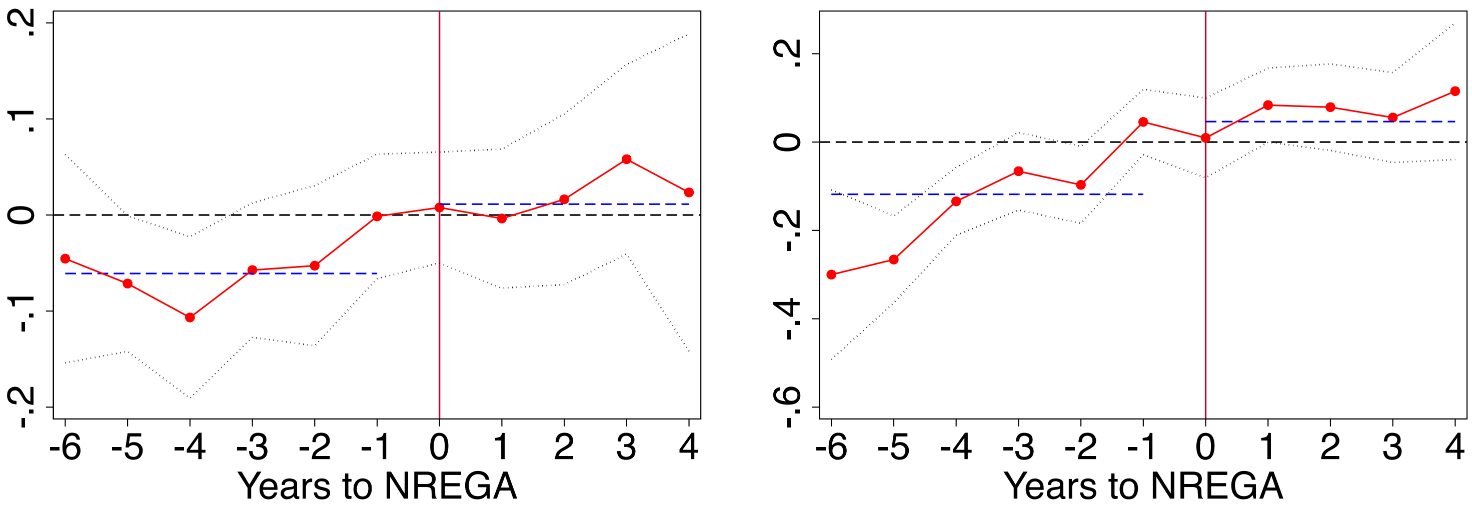

Figure 1.7: Effect of Monsoon Rain on Violent Crime (left) and on Crimes Against Public Order (right) over Time. The vertical line indicates the NREGA introduction date. The blue dashed lines indicate the coefficients obtained from a regression interaction lagged Monsoon rainfall with the NREGA treatment indicator. The red line are each point estimates of the relationship between lagged Monsoon rainfall and conflict. 95% confidence bands are indicated as dotted black lines.

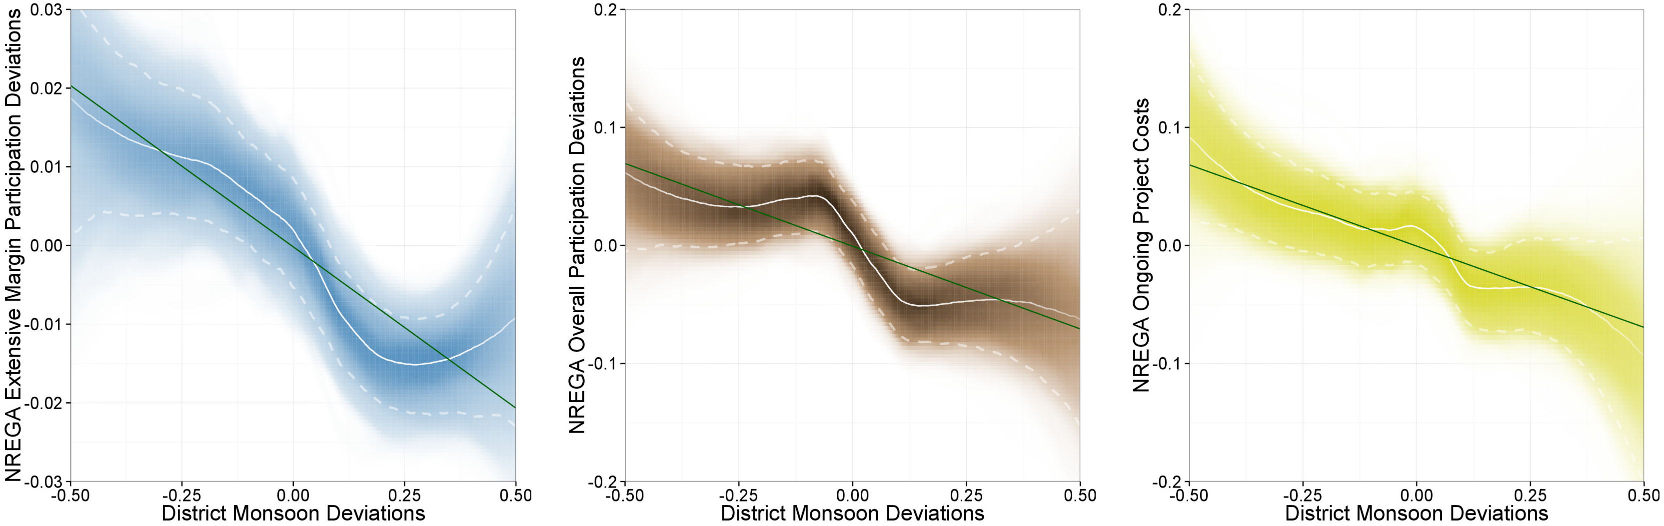

Figure 1.8: NREGA Take-Up and Lagged Monsoon Rainfall

Tables for the Main Text

Table 1.1: Summary Statistics and Socio-Economic Characteristics of Districts Before NREGA by NREGA Implementation Phase

		Phase 1	Phase 2	Phase 3
Conflict	Left-Wing Affected	0.563	0.372	0.233
	Any Violence	0.270	0.166	0.105
	Conflict Events	1.643	1.001	0.634
Income	Agricultural Output Value per Capita [INR]	2746.668	2962.016	4792.635
	Agricultural Wages [INR]	53.628	62.165	77.525
	Share of District Night Lights	0.383	0.463	0.666
Weather	Average Monsoon Season Temp [Degrees]	23.753	24.401	23.610
	Annual Rainfall [mm]	1333.619	1446.971	1258.493
	Monsoon Season Rainfall [mm]	1028.777	1052.379	878.994
	NDVI Index	0.483	0.512	0.491
Terrain	Elevation	476.021	415.810	418.082
	Ruggedness	47.760	54.268	67.933
Socio-Economic	Rural Population [share]	0.853	0.808	0.715
	Tribal Population [share]	0.226	0.163	0.112
	Scheduled Caste [share]	0.154	0.151	0.149
	Illiterate Population [share]	0.525	0.472	0.414
	Household Size [persons]	5.400	5.515	5.414
	Population younger than 6 [share]	0.262	0.253	0.239
	Population Growth 1991-2001	21.652	24.445	21.213
	Gender Gap [per 1000 inhabitants]	25.114	21.568	20.414
Infrastructure	Primary School [share]	0.810	0.820	0.857
	Mudroad [share]	0.679	0.657	0.575
	Permanent Housing [share]	0.356	0.434	0.566
	Primary Health Care [share]	0.322	0.374	0.457
	Electricity [share]	0.678	0.784	0.909
	Bus Stop [share]	0.329	0.401	0.561
	Post Office [share]	0.368	0.467	0.601
NREGA	Expenditure per Capita [INR]	436.936	523.060	247.880
	Labor Expenditure per Capita [INR]	302.197	358.869	183.402
	Days per Household	47.974	42.496	40.486
	Share of Households Participating	0.442	0.382	0.204

Notes: Socio-economic and district Infrastructure statistics based on the 2001 Census for India. Infrastructure statistics is the share of villages in a district with access to a particular type of infrastructure.

Table 1.2: Before the Introduction of NREGA: Reduced Form Relationship between Rainfall, Agricultural Production, Wages and Violence

	Agricultural Income		Conflict	
	(1)	(2)	(3)	(4)
	ln(Output/Capita)	ln(Wage)	Incidence	Intensity
log(Monsoon)	0.362***	0.058***	-0.030**	-0.897***
	(0.086)	(0.018)	(0.015)	(0.309)
Observations	3239	1419	3843	932
Number of Districts	471	314	543	144
Estimation	OLS	OLS	OLS	Poisson

Notes: All regressions include region by NREGA phase and time fixed effects and district fixed effects and constrain the analysis to the period before NREGA was introduced. Columns (1) and (2) study agricultural production and wages on an unbalanced annual district level panel, using contemporaneous Monsoon rainfall as independent variable. Column (2) also controls for state- by NREGA implementation phase linear time trends. Columns (3) and (4) are estimated on a balanced district level annual panel. Column (3) is a linear probability model using a dummy variable as dependent variable indicating whether a district experienced any conflict events in a given year. Column (4) estimates a Poisson regression with the dependent variable being the number of conflict events per district and year. Note that conditional fixed effect poisson models drop districts which do not have any variation in the dependent variable. For columns (1)-(3) standard errors are adjusted to reflect spatial dependence as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cutoff of 500 km. Poisson regressions present standard errors clustered at the district level, stars indicate *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 1.3: After the Introduction of NREGA: Reduced Form Relationship between Rainfall, Agricultural Production, Wages and Violence

	Agricultural Income		Conflict	
	(1) ln(Output/Capita)	(2) ln(Wage)	(3) Incidence	(4) Intensity
log(Monsoon)	0.374***	0.062***	-0.049***	-1.386***
	(0.078)	(0.019)	(0.018)	(0.291)
NREGA x log(Monsoon)	-0.132	-0.086***	0.043***	1.058***
	(0.083)	(0.016)	(0.014)	(0.392)
F-Test	6.36	1.68	.06	2.02
p-value	.01	.19	.8	.16
Observations	4480	2455	7059	2760
Number of Districts	471	336	543	222
Estimation	OLS	OLS	OLS	Poisson

Notes: All regressions include region by NREGA phase and time fixed effects and district fixed effects. Columns (1) and (2) study agricultural production and wages on an unbalanced annual district level panel from 2000-2009 and 2001-2010 respectively, using contemporaneous Monsoon rainfall as independent variable. Column (2) also controls for state- by NREGA implementation phase linear time trends. Columns (3) and (4) are estimated on a balanced district level annual panel. Column (3) is a linear probability model using a dummy variable as dependent variable indicating whether a district experienced any conflict events in a given year. Column (4) estimates a Poisson regression with the dependent variable being the number of conflict events per district and year. Note that conditional fixed effect poisson models drop districts which do not have any variation in the dependent variable. For columns (1)-(3) standard errors are adjusted to reflect spatial dependence as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cutoff of 500 km. Poisson regressions present standard errors clustered at the district level, stars indicate *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 1.4: Robustness to Adding Controls: Moderating Effect of NREGA on Conflict

	Incidence			Intensity			
	(1) Controls	(2) NREGA x FE	(3) State x Year FE	(4) Prev Conflict	(5) Controls	(6) NREGA x FE	(7) State x Year FE
log(Monsoon)	-0.019	-0.030**	-0.044**	-1.360***	-0.800***	-0.930***	-0.455
	(0.019)	(0.014)	(0.018)	(0.305)	(0.240)	(0.295)	(0.362)
NREGA x log(Monsoon)	-0.008	0.011	0.008	0.975**	0.951***	1.133***	0.250
	(0.014)	(0.028)	(0.013)	(0.415)	(0.327)	(0.385)	(0.340)
Observations	7059	7059	7059	1794	2760	2217	2580
Number of Districts	543	543	543	144	222	222	222

Notes: Data is on a balanced panel of conflict events from 2000-2012. Column (1)-(3) are linear probability models using a dummy variable as dependent variable indicating whether a district experienced any conflict events in a given year. Column (4) - (7) estimates a Poisson regression with the dependent variable being the number of conflict events per district and year. Note that conditional fixed effect poisson models drop districts which do not have any variation in the dependent variable. Column (1) and (5) include controls interacted with year fixed effects. The district characteristics are: terrain ruggedness, elevation, rural population share, tribal population share, scheduled caste share, illiteracy rates, household size, share of population younger than 6 years, population growth rate from census 1991 to 2001, gender gap, share of villages in district with primary school, share of villages in district with mud road approach, share of households in district that live in permanent housing, share of villages in district with primary health care facilities, share of villages with electric power, share of villages with a bus stop and the share of villages with a postal office. Columns (2) and (6) interact the NREGA treatment dummy with district fixed effects. Columns (3) and (7) control for state by year fixed effects. In column (4) I constrain the analysis to the districts that experienced conflict prior to NREGA. For columns (1)-(3) standard errors are adjusted to reflect spatial dependence as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cutoff of 500 km. Poisson regressions present standard errors clustered at the district level, stars indicate *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 1.5: Robustness to Weather Measures: Moderating Effect of NREGA on Conflict

	Incidence				Intensity			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Normalized Monsoon	-0.013**				-0.225***			
	(0.006)				(0.043)			
NREGA x Normalized Monsoon	0.006				0.157***			
	(0.005)				(0.057)			
log(GPCC Rain)		-0.031**				-1.437***		
		(0.015)				(0.277)		
NREGA x log(GPCC Rain)		0.040***				1.695***		
		(0.012)				(0.352)		
Fitted log(Monsoon)			-0.042*				-1.208***	
			(0.023)				(0.401)	
NREGA x Fitted log(Monsoon)			0.047***				0.981**	
			(0.014)				(0.441)	
NDVI				-0.287*				-2.918
				(0.172)				(2.987)
NREGA x NDVI				0.154***				2.222*
				(0.042)				(1.176)
Observations	7059	7059	7059	6516	2760	2760	2760	2536
Number of Districts	543	543	543	543	222	222	222	219

Notes: All regressions include region by NREGA phase and time fixed effects and district fixed effects. Notes: Data is a balanced panel of conflict events from 2000-2012. All weather measures are lagged by one year. Column (1)-(4) are linear probability models using a dummy variable as dependent variable indicating whether a district experienced any conflict event in a given year. Column (5) - (8) estimates a Poisson regression with the dependent variable being the number of conflict events per district and year. Note that conditional fixed effect poisson models drop districts which do not have any variation in the dependent variable. Columns (1) and (4) present results where Monsoon rainfall is normalized by its standard deviation. Columns (2) and (5) use the GPCC rainfall data as alternative rainfall data source. Columns (3) and (6) instrument the TRMM rainfall data with the GPCC data to remove measurement error. Columns (4) and (8) use the Modis Vegetation index as measure of photosynthetic activity available from 2000-2011. For columns (1)-(4) standard errors are adjusted to reflect spatial dependence as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cutoff of 500 km. Poisson regressions in columns (5) - (8) present standard errors clustered at the district level, stars indicate *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 1.6: Robustness to Treatment Timing: Moderating Effect of NREGA on Conflict

	Incidence				Intensity			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log(Monsoon)	-0.048***	-0.050***	-0.050***		-1.166***	-1.043***	-1.002***	
	(0.018)	(0.018)	(0.018)		(0.310)	(0.334)	(0.374)	
NREGA Placebo 1 x log(Monsoon)	0.029*				0.627*			
	(0.015)				(0.370)			
NREGA Placebo 2 x log(Monsoon)		0.030**				0.409		
		(0.014)				(0.378)		
NREGA Placebo 3 x log(Monsoon)			0.026*				0.310	
			(0.014)				(0.385)	
log(Outside Monsoon)				-0.002				-0.216
				(0.008)				(0.264)
NREGA x log(Outside Monsoon)				-0.001				0.162
				(0.004)				(0.316)
Observations	7059	7059	7059	7059	2760	2760	2760	2760
Number of Districts	543	543	543	543	222	222	222	222

Notes: All regressions include region by NREGA phase and time fixed effects and district fixed effects. Column (1)-(4) are linear probability models using a dummy variable as dependent variable indicating whether a district experienced any conflict event in a given year. Column (5) - (8) estimates a Poisson regression with the dependent variable being the number of conflict events per district and year. Note that conditional fixed effect poisson models drop districts which do not have any variation in the dependent variable. Placebo 1 -3 move the NREGA treatment indicator 1, 2, 3 years ahead of time, respectively. For columns (1)-(4) standard errors are adjusted to reflect spatial dependence as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cutoff of 500 km. Poisson regressions present standard errors clustered at the district level, stars indicate *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 1.7: Extended Results: Effects on Overall Crime

	(1) All	(2) Violent	(3) Property	(4) Public Order	(5) Women
log(Monsoon)	0.009	-0.061**	0.033	-0.118***	0.059*
	(0.021)	(0.031)	(0.029)	(0.042)	(0.032)
NREGA x log(Monsoon)	0.028**	0.072***	-0.016	0.165***	0.039
	(0.014)	(0.019)	(0.020)	(0.044)	(0.029)
Observations	5356	5356	5356	5356	5356
Number of Districts	537	537	537	537	537
F-test	2.64	.09	.2	.92	6.15
p-value	.1	.76	.65	.34	.01

Notes: All regressions include region-phase-time fixed effects and district fixed effects. Monsoon rain is the previous growing season's Monsoon rainfall realisation. The dependent variable is the log of the number of reported crime incidents in the category given in the column head per district and year from 2002-2012. Standard errors are adjusted to reflect spatial dependence as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cutoff of 500 km. District distances are computed from district centroids. stars indicate *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 1.8: Explaining the NREGA Effect: Monsoon Rainfall and Targets of Violence

	Incidence			Intensity		
	(1) Civilian	(2) Security	(3) Terrorist	(4) Civilian	(5) Securitry	(6) Terrorist
log(Monsoon)	-0.061***	-0.021**	-0.019	-1.532***	-0.950**	-1.196***
	(0.017)	(0.009)	(0.012)	(0.309)	(0.392)	(0.346)
NREGA x log(Monsoon)	0.044***	0.026***	0.022*	1.352***	1.317***	0.622
	(0.013)	(0.010)	(0.012)	(0.402)	(0.500)	(0.407)
Observations	7059	7059	7059	2490	1673	1893
Number of Districts	543	543	543	203	137	153

Notes: All regressions include region by NREGA phase and time fixed effects and district fixed effects. Columns (1)-(3) are linear probability models using a dummy variable as dependent variable indicating whether a district experienced any conflict events in a given year. Columns (4)-(6) estimate Poisson regressions with the dependent variable being the number of conflict events per district and year. Conflict events are categorized into whether the subject of a conflict event was a civilian, security force or terrorists. Note that conditional fixed effect poisson models drop districts which do not have any variation in the dependent variable. For columns (1)-(3) standard errors are adjusted to reflect spatial dependence as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cutoff of 500 km. Poisson regressions present standard errors clustered at the district level, stars indicate *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 1.9: Explaining the NREGA Effect: Monsoon Rainfall and NREGA Participation

	Costs	Participation			Heterogeneity	
	(1) Projects	(2) Overall	(3) Extensive	(4) Intensive	(5) Phase	(6) LWE
log(Monsoon)	-0.257**	-0.212***	-0.055***	-0.114**	-0.078***	-0.030**
	(0.101)	(0.075)	(0.014)	(0.056)	(0.024)	(0.014)
Phase 2 x log(Monsoon)					0.046*	
					(0.025)	
Phase 3 x log(Monsoon)					0.038	
					(0.025)	
LWE Affected x log(Monsoon)						-0.056**
						(0.024)
Observations	2872	3060	3060	3060	3060	3060
Number of Districts	504	537	537	537	537	537

Notes: All regressions include region-phase-time fixed effects and district fixed effects. Monsoon rain is the previous growing season's Monsoon rainfall realisation. Column (1) studies the log of total cost of achtive NREGA projects, columns (2)-(4) study participation as log of total person days employed in column (2), share of households in a district in column (3) and the log of number of days employed per household in column (4). Columns (5) and (6) study heterogeneity in extensive margin participation. Standard errors are adjusted to reflect spatial dependence as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cutoff of 500 km. District distances are computed from district centroids. stars indicate *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 1.10: Insurance Value of NREGA: Monsoon Rainfall, Output Losses and NREGA Expenditures

	Output Value/Capita	NREGA Expenditure/Capita	
	(1) OLS	(2) OLS	(3) IV
log(Monsoon$_t$)	541.335***		
	(160.249)		
log(Monsoon$_{t-1}$)		-101.219***	
		(32.984)	
Output Value/ Capita$_{t-1}$			-0.308***
			(0.103)
First Stage			12.2
Observations	4086	3059	1465
Number of Districts	438	537	410

Notes: All regressions include region-phase-time fixed effects and district fixed effects. Column (1) relates Monsoon rainfall with agricultural output per capita. Column (2) studies lagged Monsoon rainfall and its effect on levels of NREGA expenditure in a district per capita. Column (3) is an instrumental variables exercise, instrumenting lagged agricultural output value per capita with lagged Monsoon rainfall. Standard errors are adjusted to reflect spatial dependence as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cutoff of 500 km. District distances are computed from district centroids. stars indicate *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

1.A Appendix

1.A.1 Other Explanations

This section rules out a range of alternative explanations or policies that could explain why the Monsoon rainfall and conflict relationship has become weaker. Most notably is the Pradhan Mantri Gram Sadak Yojana Rural Road Construction Scheme (PMGSY) that was implemented around the same time as NREGA was devised and introduced. Other confounders are the Integrated Action Plan, which channels additional funds into left-wing extremist affected districts. I also address a concern that large mineral sectors are driving the observed moderation. Lastly, I address issues concerning a Military Operation that has been underway since early 2011.

Pradhan Mantri Gram Sadak Yojana Rural Road Construction Scheme A concern with the analysis is that the Indian government has put forth many other development programs, whose implementation may affect the relationship between Monsoon and conflict at the same time and may be correlated with the roll-out of NREGA. In this case, the results would falsely attribute the observed inward rotation of the Monsoon-rainfall and conflict relationship to the employment guarantee scheme. The most prominent developmental scheme that was implemented around the same time is the Pradhan Mantri Gram Sadak Yojana (PMGSY). This scheme was introduced in 2000 and aims to provide improved road access for rural households. The scheme in particular aimed to provide roads to all villages with at least 1000 inhabitants by 2003, with a population of 500 and more by 2007 and had special provisions for tiny villages with at least 250 inhabitants for the hill states, tribal areas and desert areas. These were to be connected by 2007. As early NREGA districts are among the poorest and least urbanised, they are more likely to have received treatment through the PMGSY as well, which could partly explain my reduced form findings.

The crucial role that transport infrastructure may have in mitigating adverse weather shocks has been highlighted in Burgess and Donaldson (2010) and Donaldson (2010). Aggarwal (2014) evaluates the impact of the PMGSY using a difference-in-difference design and finds that the scheme increased incomes by increasing the potential market size for locally produced agricultural commodities; in addition, there is less price dispersion across market centers. I use her data to see whether the PMGSY moderates the relationship between Monsoon rainfall and conflict. I construct two variables: first, the share of all unconnected habitats connected in a year and second, the cumulative share of habitats among the unconnected habitats that received road access by the end of each year. The former measure may pick up direct effects from road construction on violence, while the latter variable, in its interaction with rainfall, could pick up the more persistent effects of this scheme by connecting previously unconnected villages.

The empirical design is identical to the main analysis, except that I now add these controls and interaction terms to the main specification. The results are presented in Table 1.15. Column (1) and (2) study violence intensity, while column (3) and (4) look at incidence. Panel A presents the results for contemporary road construction, while Panel B looks at cumulative connectivity. Columns (1) and (3) look at the rural connectivity and its interaction with rainfall by itself, while column (3) and (4) are a type of horse race. In neither specifications do the road construction interactions with rainfall achieve predictive power. This renders me confident that my results genuinely reflect the effect of the workfare scheme on the dynamics of conflict.

Integrated Action Plan A second important policy aimed to tackle the Naxalite conflict is the Integrated Action Plan (henceforth, IAP). The plan was presented in 2010 and provides special funding for for districts that are considered to be severely affected by left-wing extremism. Originally it was designed for 33 districts, but since then, it expanded to provide additional funding for 82 districts. The money is to be spend on projects such as roads and other public infrastructure to improve rural livelihoods; some projects are specifically aimed to improving the way NREGA is made accessible in these districts: some IAP funding may be used to complement NREGA projects. Another margin through which the IAP may have a distinct level effect on conflict is provided as money may be used to reinforce police stations to expand the states' presence in rural areas.

Investment in infrastructure funded by the IAP could moderate the rainfall dependence of income and thus, on conflict. I don't think that the IAP would have the effects described in this paper, as its implementation would have to correlate meaningfully with lagged Monsoon rainfall. Since the grants are block grants, this is unlikely to be the case. Nevertheless I study this and the results are presented in Table 1.16. In any case, there are three simple things I can do to rule out effects of the IAP driving my results. Firstly, I can drop the 33 districts which received the scheme from 2010 onwards.[54] The results from this is presented in columns (2) and (5). The interaction term becomes smaller and size and statistical significance, especially for the conflict intensity regressions. This is not implausible as the districts that receive the IAP are ones with most variation in conflict. In second exercise, I can restrict the analysis to the period from 2000 - 2010. Again, the estimated coefficient on the post NREGA period become weaker, but the core result is still there. In the last exercise I study IAP fund expenditures, which measures utilisation of the disbursal amounts. Column (1) indicates that IAP expenditures are not correlated with lagged Monsoon rainfall. Column (3) and (6) study the effect of IAP expenditures on conflict. There appears to be a positive relationship between the two. The estimated coefficient on the NREGA interaction remains the same, thus rendering the core result robust.

[54] The districts translate into 30 districts according to the 2001 Indian census district definitions, they are: Aurangabad (Bihar), Arwal, Balaghat, Bastar, Bokaro, Chatra, Dantewada, Deogarh, Gadchiroli, Gajapati, Gaya, Garhwa, Gondiya, Gumla, Hazaribagh, Jamui, Jehanabad, Khammam, Lohardaga, Midnapore, Nabarangpur, Palamu, Pashchim singhbhum, Purba singhbhum, Rajnandgaon, Rayagada, Rohtas, Sambalpur, Sonbhadra, Surguja, Malkangiri.

Operation Green Hunt A major military operation to tackle Naxalite violence has been underway since late November 2009. The operation involves the deployment of Central Armed Reserve Police force to aide state governments tackle Naxalite threat. If deployment of troops is correlated with lagged Monsoon season rainfall, this could explain some of the observed patterns. It is not clear in which direction the effect should be. If military deployment was correlated with lagged rainfall, increased military deployment following an adverse shock could either lead to a conflict escalation or a reduction in conflict. Unfortunately, data on military deployment is not available. As with the integrated action plan, I can limit the analysis to the period before 2010 or by removing a set of districts that likely, were the primary target for a military operation. The second main concern is the relationship between the local prevalence of mineral resources. Mineral resources are a natural stabiliser to district level income, as the resource revenues are less likely to depend strongly on Monsoon rainfall. As the period around 2007 saw a major commodity boom, this could have boosted mineral resource revenues. While it is not clear that this was correlated with Monsoon rainfall in a systematic manner, it is still an important to assess the relevance of this channel as it relates this paper back to the existing literature.

Mineral Resources Another concern is that the NREGA interactions may be picking up moderation of rainfall shocks due to a sectoral shift away from agriculture to the mining sector, which is less affected by rainfall variation. Vanden Eynde (2011) shows that districts with a large mining sector see a smaller elasticity between rainfall and conflict. If the introduction of NREGA is correlated with a sectoral shift towards the mineral resource sector, the NREGA interactions could be picking up this effect. This is not entirely implausible as the mid 2000s saw a commodity price boom which could have induced a lot more investment in the mining sector. In order to control for this I construct a share of a district's income that is due to the mining sector.[55]

Again, the specifications I present are very similar, adding a simple interaction with the mining sector share in district domestic product interacted with the Monsoon season rainfall. The results are presented in Table 1.17. Column (1) presents the results on violence intensity without the NREGA interactions. It becomes evident that districts with a larger share of the mining sector experience a weaker relationship between violence and Monsoon rain. This maps into the findings of Vanden Eynde (2011). Once including the NREGA interaction, the coefficient on the Mining sector interaction becomes insignificant with a p-value of 12%. More importantly, the NREGA interaction remains strongly significant. This suggests that the NREGA effect seems not to be picking up a moderation in the Monsoon shock and conflict relationship due to the presence of a large mineral resource sector.

The next section provides tentative results of a level effect of NREGA on conflict levels.

[55] I use district domestic product data for years between 1998 to 2005 available from `http://planningcommission.nic.in/plans/stateplan/index.php?state=ssphdbody.htm`, accessed on 21.06.2014. The district domestic product construction is discussed in detail in Katyal et al. (2001).

1.A.2 Level Effect of NREGA

The preceding results suggested that NREGA does have a moderating effect on the cyclical nature of violence, in particular, the violence targeted against civilians. However, the existing literature evaluating the economic impacts of NREGA also indicate strong increases in wage levels.[56] An increase in wage levels can be seen as an increase in the returns to labour in both, good- and bad states of the world. This does have an independent level effect on conflict. It is challenging to identify a level effect due to the endogeneity of the roll-out. Nevertheless, in this section I provide an estimate of the level effect of NREGA. I estimate specifications with less demanding time-fixed effects that vary by region. This ensures that the coefficient on the NREGA treatment dummy is not collinear with the time effects and can thus, be interpreted. The specification I estimate is as follows:

$$\mathbb{E}(A_{dprt}) = \delta_d \exp\left(b_{rt} + \alpha T_{dprt} + \eta R_{dpr,t-1} + \gamma T_{dprt} \times R_{dpr,t-1} + \epsilon_{dprt}\right) \tag{1.8}$$

where b_{rt} are now region by time fixed effects, rather than region by phase and time fixed effects. This set of fixed effects allows the estimation of the parameter α, which can be interpreted as the level effect of NREGA if we are willing to assume that the roll-out of NREGA was exogenuous. This is not a classical difference in difference estimator with one set of treated and one set of untreated locations since eventually, all districts receive NREGA. The coefficient α is estimated off the time variation due to the sequential roll out of NREGA. This implies that the NREGA treatment indicator is estimated soley from the variation for the years in which some districts had already received NREGA relative to other districts that did not have NREGA yet; that is, the coefficient solely lives off the variation in differences in conflict across phases for the years 2006 and 2007. When adding interaction terms with Monsoon rainfall and NREGA, the interaction term becomes, in addition, a heterogenous effect for the level effect of NREGA in these two years. In order to get the average treatment effect I demean Monsoon rainfall variable for these regressions. I estimate three versions of the above specification. First, imposing the constraint that $\eta = \gamma$. In this case, I force the effect of rainfall to be the same before and after the introduction of NREGA. I also estimate a specification with the constraint $\eta = \gamma = 0$, which effectively means not controlling for rainfall. The key question is how this will affect the estimated coefficient $\hat{\alpha}$. In both cases, the coefficient $\hat{\alpha}$ should overstate the effect of NREGA in absolute value.

The results are presented in table 1.18. The first column presents the constrained regression where I do not control for rainfall. The level effect coefficient is negative and statistically significant. This coefficient is a mixture of the level effect and the implied effect due to a reduced rainfall and conflict elasticity. In the second column, I control for rainfall, which renders the coefficient slightly larger in absolute value.

[56] See Zimmermann (2012), Berg et al. (2012), Imbert and Papp (2015) and Azam (2011).

The third column is the unconstrained coefficient, allowing the functional relationship between rainfall and conflict to change with the introduction of NREGA. The interesting observation is that the coefficient on the level effect goes down and is estimated relatively imprecisely, moving from a p-value close to 0.001 to p-value of 0.45. This suggests that the dynamic effect of NREGA, operating by mitigating income shocks, is being partially captured in estimates of $\hat{\alpha}$, when one does not explicitly control for this important economic channel through which NREGA operates. When comparing column (2) and column (3), this suggests that at least 1/3 of the estimated reduction in violence levels is due to the reduced rainfall dependence of conflict.

This paper provides evidence that NREGA functions as insurance. This suggests that the correct way to evaluate NREGA is through its dynamic effect through program participation. Nevertheless recently, Dasgupta et al. (2014) and Khanna and Zimmermann (2013) separately estimated level effects of the introduction of NREGA. They arrive at different conclusions. Khanna and Zimmermann (2013) use a regression-discontinuity design relying on a reverse engineered NREGA roll-out algorithm to identify districts that were close to the cutoff of being assigned into either an earlier, or a later phase. They argue that this provides a good counterfactual for a fuzzy regression discontinuity design and estimate the effect of the NREGA treatment. They find that NREGA increased conflict levels in the short-run. Dasgupta et al. (2014) use a difference in difference estimator as I discussed above. This design lives off variation in differences in conflict arising due to the gradual roll-out. I provide some evidence of level effects in this paper, estimating a similar difference-in-difference specification as in Dasgupta et al. (2014). The results are presented in Table 1.19. The first column presents the basic level effect estimate of contemporaneous treatment. The second column adds lagged effects of the NREGA treatment indicator, suggesting that the first lag is highly significant. The point estimate suggest that the introduction of NREGA reduced levels of violence by between 30% to 50% for average Monsoon rainfalls. Columns (4)-(9) explore the heterogeneity of the estimated effect by interacting the treatment indicator with a set of district-characteristics. The district characteristics are demeaned for ease of interpretation of the marginal effects. The results suggest that the level effect is weaker for districts with a high scheduled tribe share, but stronger for districts with higher scheduled caste share. Indicative is the coefficient on average household size. This suggests that the level effect is significantly weaker for districts with a larger average household size. Since the NREGA program provides an allowance for 100 days of work *per household*, larger households are disadvantaged in that respect. Column (8) interacts the treatment indicator with the the mean level of agricultural output per capita before 2005 expressed in INR 1000. The coefficient is negative and significant, suggesting that richer districts saw a stronger drop in conflict. While the results on the dynamics of conflict do not square with Khanna and Zimmermann (2013), the estimated level effects do stand at odds with the ones estimated in their paper but map well into the findings of Dasgupta et al. (2014). That being said, as NREGA is aimed to provide insurance in bad states of the world, this insurance value should be driving the change in conflict levels. This is captured in this paper through the changing slope linking Monsoon rain and conflict.

1.A.3 NREGA Road Construction and Conflict

NREGA aims to "create durable assets which have potential to generate additional employment in the years to come in rural areas."[57] While the analysis of the agricultural production function did not suggest a dramatic change in the relationship between Monsoon rainfall and agricultural output in the short run, asset construction under NREGA could still affect the dynamics of conflict. There is anecdotal evidence suggesting that Naxalites oppose road construction under the scheme (see Banerjee and Saha (2010)). The anecdotal accounts suggest that this is for fear that roads could provide easier access for police and military. There is anecdotal evidence suggesting that Naxalites have taken road construction contractors hostage or killed them, suggesting that road construction could drive conflict.[58] There are two ways that road-construction could affect the results here. First, road construction may itself be correlated with lagged Monsoon rainfall and through that affect the dynamics of conflict in a way that is correlated with Monsoon rainfall. There could also be an independent effect from road construction that affects conflict levels. This section shows two things. I show that road construction is correlated with lagged Monsoon rainfall; however, this relationship is not present for districts for which the moderation in the rainfall and conflict relationship is strongest. There appears to be a distinct effect of road construction on conflict that is not related to Monsoon season rainfall. Districts in which the share of overall NREGA funds allocated to road construction in the years since NREGA was introduced is higher, experience more conflict in recent years. This effect is a mere correlation due to the endogeneity of NREGA road construction. It is however, indicative for further research.

Monsoon Rainfall and Road Construction If road construction itself was correlated with lagged Monsoon rainfall, this could explain the finding of the inward rotation of the relationship between Monsoon rainfall and conflict. The argument is quite simple. Before NREGA, good Monsoon rains would reduce conflict. With NREGA available, strong Monsoon rainfalls may be associated with increasing road construction to repair mud roads that have been damaged due to the Monsoon. This may lead to more conflict, reversing the previously existing relationship. This is a genuine concern and is studied in this section.

A brief look at summary statistics is already quite telling. Studying districts that have been categorized as being under left wing extremist influence or have seen some conflict for the period prior to NREGA suggests that these districts see a significantly

[57] See `http://www.nrega.nic.in`, accessed 12.02.2014.

[58] See for example `http://www.satp.org/satporgtp/detailsmall_news.asp?date1=7/16/2011&id=5`, accessed 11.10.2014.

lower share of NREGA expenditure going to road construction. The total expenditure weighted share in 2010 is around 31.3%, while it is 36.3% for the other districts. This is an important insight, as it suggests that the types of assets created under NREGA may reflect local preferences.

Table 1.20 presents results studying how the share of overall NREGA expenditures in road construction (columns 1-3) or for land development (column 4-6) in a financial year respond to Monsoon rainfall in the previous season. The relationship suggests that lagged Monsoon rainfall predicts an increasing share of road construction (column 1), but not for projects that can be classified for land development (column 4). This gives rise to the genuine concern that rainfall may drive road construction which, in turn, is driving conflict as suggested. I rule out this explanation by studying the heterogeneity across NREGA implementation phase, by whether districts are classified as being under left-wing extremist influence and by studying whether it is excessive Monsoon rains that drive this effect.

Columns (2) and (4) study this relationship for districts by implementation phase: it appears that the positive Monsoon rainfall and road construction relationship is driven by districts that receive NREGA in early phases. Columns (3) and (6) study the responsiveness of NREGA asset construction for districts classified under left-wing extremist influence. Neither expenditures for land-development (column 6) nor road construction (column 3) meaningfully correlate with Monsoon rainfall for these districts. This is reassuring given that the moderating effect of NREGA on the Monsoon conflict relationship is coming mostly from these districts.

A non-parametric analysis further suggests that it is positive rainfall that correlates with road construction for phase 1 districts (see Figure 1.9). This is an important insight as the non-parametric analysis of the Monsoon rainfall and conflict (or crime) relationship indicated that NREGA's moderating effect on that relationship is due its impact on below normal Monsoon rainfall. This renders me confident that, while road construction may have an independent effect on conflict, this effect is not confounding the moderation in the Monsoon rainfall and conflict relationship studied in this paper.

This is also conceptually reasonable: road construction due to excessive Monsoon may simply repair and replace already existing mud roads, which are most prevalent in districts that received NREGA in the first phase (see Table 1.1). This is qualitatively a lot different from new roads being constructed. Places, in which a lot of roads are constructed, may experience a change in the conflict dynamic, as it is new roads that improve access to remote places for the military and not the improvement or repairing of existing roads. That is to say: places that receive a lot of road construction through NREGA independent of Monsoon rainfall may experience a change in the conflict dynamics that is, however, unrelated to the moderation in the Monsoon and conflict relationship studied here. This is highlighted in the next paragraph.

Independent Effect of New Road Construction In order to study the direct effect of road construction on conflict that is unrelated to Monsoon rainfall, I construct a measure of road construction intensity as the share of all funds devoted to road construction activity for all the post NREGA years. This overall measure may reflect local preferences for different development projects that is, due to the averaging, independent of Monsoon rainfall. Let this measure be denoted as ρ_d. I estimate the event study analysis interacting the time to treatment with the measure ρ_d and plot out the coefficients. The coefficients are estimated off the variation in NREGA road construction intensity across districts. The results are presented in Figure 1.10. The pattern that emerges suggests that NREGA road construction intensity is correlated with higher incidence and intensity of conflict following the introduction of NREGA.

Figures and Tables for Robustness Appendix

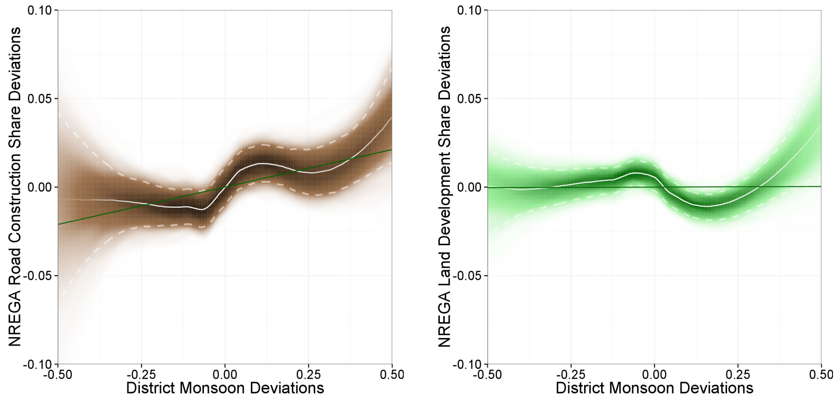

Figure 1.9: NREGA Infrastructure Expenditure Shares and Lagged Monsoon Rainfall for Phase 1 Districts

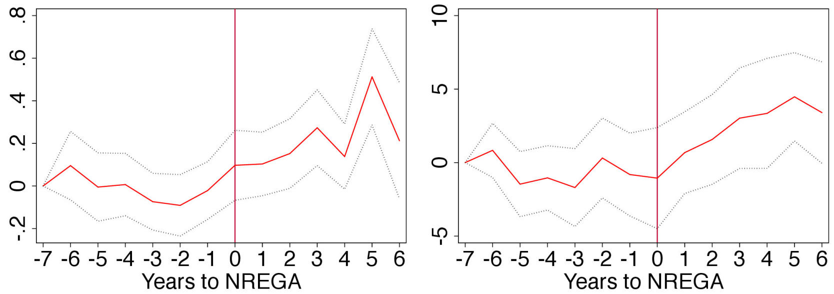

Figure 1.10: Effect of NREGA Cumulative Road Construction Expenditure Share on Conflict over Time.

Table 1.11: Socio-Economic Characteristics of Districts with Naxalite Presence

	LWE Affected	Other Districts
Number of Districts	206	365
Panel A: Demographic		
Rural Population	77.17	69.52
Tribal Population	9.47	7.49
Scheduled Caste	16.55	15.75
Illiterate	46.07	44.56
Population Age < 6	24.39	24.81
Permanent House	46.69	55.81
Panel B: Infrastructure		
Primary School	74.27	81.57
Mud Road	72.11	59.63
Primary Health Care	29.72	33.19
Electricity	65.41	84.86
Bus Stop	29.66	40.35
Post Office	34.03	46.83

Notes: Statistics derived from the 2001 Census for India. Panel A presents demographic indicators as shares of the overall population. Panel B presents Infrastructure indicators derived from the share of villages that have access to a particular type of infrastructure.

Table 1.12: Before the Introduction of NREGA: Robustness of Relationship between Weather Variables and Agricultural Output

	log(Output Value/Capita)			log(Grain Value/Capita)		
	(1) Outside Monsoon	(2) Temperature	(3) Controls	(4) Outside Monsoon	(5) Temperature	(6) Controls
log(Monsoon)	0.364***	0.364***	0.424***	0.369***	0.369***	0.335***
	(0.086)	(0.087)	(0.080)	(0.076)	(0.076)	(0.073)
log(Outside Monsoon)	0.122**	0.124**	0.129***	0.114***	0.115***	0.090**
	(0.051)	(0.051)	(0.046)	(0.043)	(0.043)	(0.042)
Temperature		-0.063	-0.030		-0.028	-0.053*
		(0.051)	(0.031)		(0.040)	(0.030)
Observations	3239	3239	3239	3196	3196	3196
Number of Districts	471	471	471	464	464	464

Notes: All regressions include region-phase-time fixed effects and district fixed effects. Temperature measures the average temperature during the Monsoon months. Columns (1)-(3) study agricultural output value per capita, while columns (4) - (6) study the value of grain production encompassing ragi, rice, wheat, bajra, jowar, maize, pules and barley. Columns (3) and (6) add a set of district characteristics interacted with a set of year fixed effects. The district characteristics are: terrain ruggedness, elevation, rural population share, tribal population share, scheduled caste share, illiteracy rates, household size, share of population younger than 6 years, population growth rate from census 1991 to 2001, gender gap, share of villages in district with primary school, share of villages in district with mud road approach, share of households in district that live in permanent housing, share of villages in district with primary health care facilities, share of villages with electric power, share of villages with a bus stop and the share of villages with a postal office. Standard errors are adjusted to reflect spatial dependence as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cutoff of 500 km. District distances are computed from district centroids, stars indicate *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 1.13: Before the Introduction of NREGA: Robustness of Relationship between Weather Variables and Agricultural Wages

	log(Annual Wage)			log(Seasonal Wage)	
	(1) Outside Monsoon	(2) Temperature	(3) Controls	(4) Harvesting	(5) Planting
log(Monsoon)	0.058***	0.058***	0.057***	0.046*	0.010
	(0.019)	(0.018)	(0.020)	(0.028)	(0.017)
log(Outside Monsoon)	0.006	0.006	0.011		
	(0.011)	(0.011)	(0.015)		
Temperature		0.001	0.021		
		(0.015)	(0.013)		
State by NREGA Phase Trend	Yes	Yes	No	Yes	Yes
Observations	1419	1419	1419	1387	1195
Number of Districts	314	314	314	318	260

Notes: All regressions include region-phase-time fixed effects and district fixed effects. State by NREGA-Phase Trend are linear trends at the State by NREGA implementation phase level. Temperature measures the average temperature during the Monsoon months. Data is an unbalanced district level panel of annual agricultural wages in India. Columns (1)-(3) study agricultural wages, while columns (4) and (5) study wages at the planting stage compared to wages at harvesting stage towards the end of the year. Column (3) adds a set of district characteristics interacted with a set of year fixed effects. The district characteristics are: terrain ruggedness, elevation, rural population share, tribal population share, scheduled caste share, illiteracy rates, household size, share of population younger than 6 years, population growth rate from census 1991 to 2001, gender gap, share of villages in district with primary school, share of villages in district with mud road approach, share of households in district that live in permanent housing, share of villages in district with primary health care facilities, share of villages with electric power, share of villages with a bus stop and the share of villages with a postal office. Standard errors are adjusted to reflect spatial dependence as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cutoff of 500 km. District distances are computed from district centroids, stars indicate *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 1.14: Before the Introduction of NREGA: Robustness of Relationship between Monsoon Rainfall and Conflict

	Robustness to Choice of Empirical Model			Controls and Interactions		
	(1) Poisson-IV	(2) Neg Bin	(3) OLS	(4) Weather	(5) Controls	(6) Interactions
log(Fitted Output Value/Capita$_{t-1}$)	-2.362**					
	(0.986)					
log(Monsoon$_{t-1}$)		-0.830***	-0.313**	-0.960***	-0.811**	-1.192***
		(0.230)	(0.147)	(0.346)	(0.328)	(0.422)
log(Outside Monsoon$_{t-1}$)				-0.384*		
				(0.205)		
Temperature$_{t-1}$				0.504		
				(0.314)		
log(Monsoon$_t$)				-0.056		
				(0.257)		
Temperature$_t$				0.096		
				(0.310)		
District Controls	No	No	No	No	Yes	No
Monsoon Rain Interactions	No	No	No	No	No	Yes
Observations	646	932	3843	932	932	932
Number of Districts	117	144	543	144	144	144

Notes: All regressions include region-phase-time fixed effects and district fixed effects. Column (1) presents the results of an IV regression, instrumenting lagged agricultural output per capita with lagged Monsoon rainfall. Column (2) is a negative binomial, while column (3) presents OLS results. Column (4) includes temperature during the Monsoon season as well as contemporaneous weather. Column (5) interacts a set of district controls with a set of year fixed effects, while column (6) interacts Monsoon rainfall with the demeaned district characteristics. The district characteristics are: terrain ruggedness, elevation, rural population share, tribal population share, scheduled caste share, illiteracy rates, household size, share of population younger than 6 years, population growth rate from census 1991 to 2001, gender gap, share of villages in district with primary school, share of villages in district with mud road approach, share of households in district that live in permanent housing, share of villages in district with primary health care facilities, share of villages with electric power, share of villages with a bus stop and the share of villages with a postal office. Robust standard errors clustered at the district level are given in the parentheses with stars indicating *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 1.15: Alternative Mechanism: Rural Connectivity and Moderation of Rainfall and Conflict Relationship

	Incidence		Intensity	
	(1)	(2)	(3)	(4)
Panel A: Road Construction				
log(Monsoon)	-0.040**	-0.059***	-0.712***	-1.418***
	(0.020)	(0.022)	(0.212)	(0.358)
NREGA x log(Monsoon)		0.047***		1.093***
		(0.017)		(0.394)
Roads	-0.387	-0.622	5.127	4.153
	(0.571)	(0.571)	(5.249)	(5.203)
Roads x log(Monsoon)	0.073	0.108	-0.647	-0.521
	(0.092)	(0.091)	(0.753)	(0.756)
Panel B: Cumulative Road Construction				
log(Monsoon)	-0.049**	-0.055**	-0.927***	-1.325***
	(0.023)	(0.024)	(0.311)	(0.347)
NREGA x log(Monsoon)		0.044***		1.415***
		(0.016)		(0.354)
Cumulative Roads	-0.248	0.005	-7.788*	4.778
	(0.248)	(0.226)	(4.647)	(4.011)
Cumulative Roads x log(Monsoon)	0.040	0.002	0.917	-0.880
	(0.037)	(0.034)	(0.661)	(0.586)
Observations	5775	5775	2230	2230
Number of Districts	525	525	210	210
Estimation	OLS	OLS	Poisson	Poisson

Notes: All regressions include region-phase-time fixed effects and district fixed effects. Monsoon rain is the previous growing season's Monsoon rainfall realisation. The dependent variable is the number of violent incidences per quarter in columns (1) and (2) and an indicator whether there was any violent incidence in columns (3) and (4). Panel A studies the effect of contemporaneous road construction on violence, while Panel B studies the impact of rainfall through the overall share of unconnected habitats that became connected up to 2012. Standard errors are clustered at district level in column (1) and (2), while in column (3) and (4) they are adjusted to reflect spatial dependence as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cutoff of 500 km. District distances are computed from district centroids. stars indicate *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 1.16: Alternative Mechanism: Integrated Action Plan Disbursals and the Moderation of Monsoon Rainfall and Conflict Relationship

	IAP	Incidence			Intensity		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
log(Monsoon)	0.407	-0.045***	-0.044***	-0.048***	-1.257***	-1.170***	-1.352***
	(0.426)	(0.017)	(0.015)	(0.018)	(0.279)	(0.271)	(0.286)
NREGA x log(Monsoon)		0.037***	0.046***	0.042***	0.642*	0.880**	1.008***
		(0.014)	(0.016)	(0.014)	(0.366)	(0.424)	(0.384)
IAP Expenditure				0.011			0.084***
				(0.010)			(0.029)
Observations	184	6669	5973	7059	2354	2152	2760
Number of Districts	72	513	543	543	192	205	222
Estimation	OLS	OLS	OLS	OLS	Poisson	Poisson	Poisson

Notes: All regressions include region-phase-time fixed effects and district fixed effects. Monsoon rain is the previous growing season's Monsoon rainfall realization. Column (1) studies IAP expenditure as a function of lagged Monsoon rain. The dependent variable in columns (2)-(4) is an indicator whether there was any conflict event in a district and year, while it is the number of conflict events per year in columns (5)-(7). Columns (2) and (5) remove the 33 districts that received the IAP originally. Columns (3) and (6) restrict the analysis to the period 2000-2010. Columns (4) and (7) control for IAP expenditure. Standard errors in column (1)-(4) are adjusted to reflect spatial dependence as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cutoff of 500 km. District distances are computed from district centroids. Errors in columns (3) and (4) are clustered at the district level, with stars indicating *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 1.17: Alternative Mechanism: Mining Sector Share, Commodity Boom and Moderation of Rainfall and Conflict Relationship

	Incidence		Intensity	
	(1)	(2)	(3)	(4)
log(Monsoon)	-0.042**	-0.054***	-0.947***	-1.797***
	(0.018)	(0.018)	(0.199)	(0.283)
NREGA x log(Monsoon)		0.049***		1.369***
		(0.015)		(0.389)
Mining Sector Share	0.011	0.012	4.554*	3.820*
x log(Monsoon)	(0.315)	(0.313)	(2.348)	(2.256)
Observations	6552	6552	2504	2504
Number of Districts	504	504	204	204
Estimation	OLS	OLS	Poisson	Poisson

Notes: All regressions include region-phase-time fixed effects and district fixed effects. Monsoon rain is the previous growing season's Monsoon rainfall realisation. The dependent variable is an indicator whether there was any conflict event in columns (1) and (2) and the number of violent incidences per year in columns (3) and (4). Mining Sector Share is the share of the districts domestic product that is generated in the Mining sector based on data between 1998 and 2005. Standard errors in column (1) and (2) are adjusted to reflect spatial dependence as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cutoff of 500 km. District distances are computed from district centroids. Errors in columns (3) and (4) are clustered at the district level, with stars indicating *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 1.18: Dynamic Versus Direct Level Effect of NREGA

	(1) $\eta = \gamma = 0$	(2) $\eta = \gamma$	(3) Unconstrained
NREGA	-0.437***	-0.498***	-0.300*
	(0.165)	(0.178)	(0.170)
log(Monsoon)		-0.859***	-1.758***
		(0.237)	(0.356)
NREGA x log(Monsoon)			1.457***
			(0.369)
Observations	2886	2886	2886
Number of Districts	222	222	222

Notes: All regressions include region-by-time fixed effects and district fixed effects. Monsoon rain is the previous growing season's Monsoon rainfall realisation; Monsoon rainfall is demeaned for ease of interpretation of the interaction terms. All regressions are estimated using Poisson models with the dependent variable being the number of conflict events per district and year. The first column does not control for Monsoon rainfall, while the second column constraints the rainfall coefficient to be the same before, and after the introduction of NREGA. Robust standard errors clustered at the district level are given in the parentheses with stars indicating *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 1.19: Level Effect of NREGA

	Level Effect Estimates			Heterogeneity of Level Effect				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
NREGA	-0.300*	-0.225	-0.533***	-0.724***	-0.297*	-0.263	-0.695***	-0.930***
	(0.170)	(0.150)	(0.197)	(0.198)	(0.173)	(0.171)	(0.234)	(0.232)
NREGA$t-1$		-0.470**						
		(0.201)						
Heterogeneity: NREGA $\times$								
Scheduled Tribe			1.051***					0.540
			(0.400)					(0.448)
Scheduled Caste				-5.974***				-2.651
				(1.539)				(2.166)
Illiteracy					0.564			-0.708
					(1.439)			(1.469)
Householdsize						0.461***		0.250
						(0.178)		(0.188)
Agricultural GDP Before 2005							-0.224**	-0.181*
							(0.096)	(0.093)
NREGA Dynamic Effect								
Monsoon Rain	-1.758***	-1.725***	-1.708***	-1.852***	-1.760***	-1.596***	-1.883***	-1.812***
	(0.356)	(0.348)	(0.351)	(0.333)	(0.355)	(0.373)	(0.420)	(0.428)
NREGA x Monsoon	1.457***	1.489***	1.381***	1.565***	1.459***	1.401***	1.589***	1.585***
	(0.369)	(0.372)	(0.361)	(0.360)	(0.368)	(0.373)	(0.454)	(0.464)
Observations	2886	2886	2886	2886	2886	2886	2496	2496
Number of Districts	222	222	222	222	222	222	192	192

Notes: All regressions include region-by-time fixed effects and district fixed effects. The time period is restricted to the period before NREGA was introduced. All regressions are estimated using Poisson models with the dependent variable being the number of conflict events per district and year. Robust standard errors clustered at the district level are given in the parentheses with stars indicating *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 1.20: Explaining the NREGA Effect: Monsoon Rainfall and NREGA Infrastructure Construction

	Road Construction			Land Development		
	(1)	(2)	(3)	(4)	(5)	(6)
Monsoon	0.033**	0.040*	0.059***	-0.020	-0.056**	-0.042**
	(0.014)	(0.021)	(0.017)	(0.013)	(0.025)	(0.021)
Phase 2 x Monsoon		-0.006			0.031	
		(0.030)			(0.030)	
Phase 3 x Monsoon		-0.016			0.084***	
		(0.027)			(0.030)	
LWE Affected x Monsoon			-0.056***			0.050*
			(0.022)			(0.026)
Observations	2894	2894	2894	2894	2894	2894
Number of Districts	504	504	504	504	504	504

Notes: All regressions include region-phase-time fixed effects and district fixed effects. Monsoon rain is the previous growing season's Monsoon rainfall realisation. The dependent variable is the share of NREGA expenditures in a district that go into road construction relative to land development. Standard errors are clustered at district level in column (1) and (2), while in column (3) and (4) they are adjusted to reflect spatial dependence as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cutoff of 500 km. District distances are computed from district centroids. stars indicate *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

1.A.4 Conflict Data

Empirical research on the economics of conflict almost always suffer from severe data limitations. This lies in the nature of the subject of study, that typically places that exhibit conflict are only weakly institutionalised with little official report of violence and little press and media coverage. Blattman and Miguel (2010)'s review cites that the correlation across different civil war datasets ranges from 0.42 to 0.96, which may be the reason why empirical results are often not reproducible using similar identification strategies, but different datasets or variable definitions (e.g. Ciccone (2011)).

There exists no broad conflict dataset that covers India or South East Asia as a whole. This gap was filled through the violence dataset introduced in Fetzer (2013). This paper documents the process through which in the Indian context 28,638 newspaper reports were transformed into a workable conflict dataset using both machine-learning, semi-automated coding techniques and scalable manual hand-coding methods.[59] This section sketches the semi-automated process through which the daily newspaper clippings are transformed (more details are provided in Fetzer (2013)). A typical sample may look as follows:

> Two unidentified terrorists massacred six members of a family and left a seventh injured at Mangnar Top, Poonch district, on December 31, 2001. Local residents refused to cremate the bodies of the slain victims, insisting that a Union Minister should visit the area and take notice of the increasing terrorist violence there.

The semi-automated routine defines a terrorist-incident as an Event-tuple, $E = \{L, T, V, S, O\}$ defined by a location L, a date or time of the event T, a verb V that indicates the type of violent act, and the verb's associated subject S, the perpetrator of the act and the object O that was subjected to the act V. The semi-automated routine tries to fill all these elements of the tuple for each sentence using common machine-learning algorithms implemented in natural language processing packages.

I work with the following set of Trained Natural Language Processing Algorithms:

1. Sentence Detection to break up individual sentences.

2. Semantic Role Labelling (SRL) to tag the grammatical structure of words in relation to one another.

3. Named Entity Recognition (NER) to identify names (places, institutions, names) lives off spelling, preposition and gazetteer. Complemented with dictionary of 1,978 spelling variations.

4. Part of Speech Tagging (POS) to tag role of words (subject, verb, object)

[59] The raw material was a set of 28,638 newspaper clippings collected by the Institute for Conflict Management in New Delhi through the South Asian Panel on Terrorism (SATP) since 2001, see `http://www.satp.org`, accessed in October 2012.

These are together implemented in SENNA (Collobert et al., 2011), available as open-source in C. The sample output for the above sentence would look like:

Two			B-A0	B-A0	
unidentified			I-A0	I-A0	
terrorists			E-A0	E-A0	
massacred		massacred	S-V		
six			B-A1		
members			I-A1		
of			I-A1		
a			I-A1		
family			E-A1		
and					
left		left		S-V	
a				B-A1	B-A1
seventh				I-A1	E-A1
injured		injured		I-A1	S-V
at				I-A1	B-AM-LOC
Mangnar	B-LOC			I-A1	I-AM-LOC
Top	E-LOC			I-A1	I-AM-LOC
,				I-A1	I-AM-LOC
Poonch	S-LOC			I-A1	I-AM-LOC
district				I-A1	E-AM-LOC
,				I-A1	
on				I-A1	
December				I-A1	
31				I-A1	
,				I-A1	
2001				E-A1	

In the above text-snippet, only one sentence satisfies the requirement of all elements forming an event tuple $E = \{L, T, V, S, O\}$ being present. This yields:

$$\begin{aligned} E_1 \;=\; & \{'\text{Mangar Top Poonch}', '\text{December 31 2001}', \\ & '\text{massacre}', '\text{two unidentified terrorists}', \\ & '\text{six members of a family at Mangnar Top, Poonch district}'\} \end{aligned}$$

An incident is counted as long as all pieces of information can be deduced from the underlying sentence. This is essentially mimicking the process through which humans would code this data manually. An exhaustive list of verbs is used to spot events and a sentence is normalised to contain at most one event. The individual elements of the tuple E are then transformed by assigning labels to the snippets indicating whether the actor was a terrorist, security force or a civilian and similarly for who subjected to the act V. Note that in the sentence there exists a further event:

$$
\begin{aligned}
E_2 \quad = \quad & \{'\text{Mangar Top Poonch}','\text{December 31 2001}', \\
& '\text{left}','\text{two unidentified terrorists}', \\
& '\text{a seventh injured at Mangnar Top, Poonch district}'\}
\end{aligned}
$$

As described in Fetzer (2013), a sentence will be counted as containing information of at most one incident. The data has been evaluated in Fetzer (2013) and correlates very well with hand-coded data. The correlation between this automatically retrieved data and the hand-coded data for the Naxalite conflict used by Vanden Eynde (2011) is at least 93%.

1.A.5 Comparison of Results with Global Terrorism Database

This section highlights that the results obtained in This chapter can not be replicated when studying the conflict for India contained in the Global Terrorism Database (GTD) collected by National Consortium for the Study of Terrorism and Responses to Terrorism at the University of Maryland. This database has been used in more than 30 journal publications and thus, serves as an interesting testing ground. Unfortunately, the GTD database does not come at a district level spatial resolution. However, it provides the nearest big town to where the incident occurred. In order to be able to compare the datasets, I geo-code the locations of the nearest towns to obtain a similar district level count variable of the number of conflict events. I then estimate the main specifications using the number of terrorist incidences in the global terrorism database as a left-hand side. The results are presented in Table 1.21.

Table 1.21: NREGA Effect in the GTD and Fetzer (2013) dataset

	Fetzer (2013) Dataset			Global Terrorism Database		
	(1) Pre NREGA	(2) Dynamic	(3) Level	(4) Pre NREGA	(5) Dynamic	(6) Levels
Monsoon	-0.866***	-1.330***	-0.680***	-0.985	-1.338*	-1.062**
	(0.270)	(0.306)	(0.261)	(0.684)	(0.764)	(0.462)
NREGA x Monsoon		1.098***			0.359	
		(0.388)			(0.676)	
NREGA			-0.540***			-1.098
			(0.166)			(1.264)
Observations	2841	8868	10199	851	5268	5268
Number of Districts	148	217	217	57	186	186

Notes: All regressions are estimated using a pseduo-maximum likelihood estimator, whose moment conditions coincide with a Poisson model. Regressions in columns (1)-(2) and (4)-(5) include region-phase-time fixed effects as well as district fixed effects, while results for columns (3) and (6) come from a regression with time- and district fixed effects. The dependent variable is the number of incidences per district and quarter. Robust standard errors clustered at the district level are given in the parentheses with stars indicating *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Columns (1)-(3) study the dataset used in this paper, while columns (4)-(6) use the GTD database. In column (4) it becomes obvious that in the GTD data, there appears to be no statistically significant correlation between rainfall and conflict, while there is a strong documented in the Fetzer (2013) data in column (1). The geographic coverage of the GTD dataset is a lot more limited before the introduction of NREGA, with only 57 districts reported as having violent incidences before NREGA was introduced while there are almost three times as many districts reported in the other datasets. The moderating effect of NREGA is seen only in column (2), but not in column (5), albeit the coefficient is positive.

As the number of districts covered in the GTD database seems to increase dramatically when expanding the analysis to the whole time-period in column (5) it becomes instructive to study how the correlation between these two datasets has evolved over time. I regress the two datasets onto one another, allowing for there to be a separate coefficient for each year:

$$GTD_{dt} = \delta_d + b_{rt} + \sum_{t=2000}^{2010} \gamma_t A_{dt} + \epsilon_{dt}$$

The estimated coefficients γ_t are plotted out in Figure 1.11.

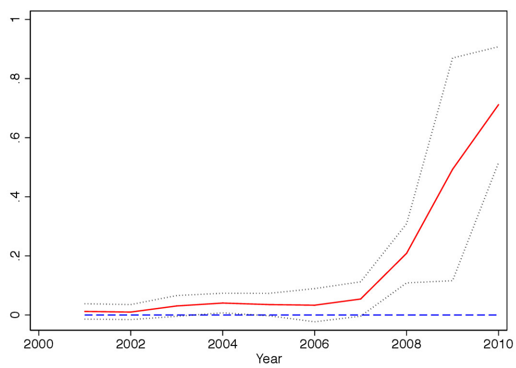

Figure 1.11: Relationship between Fetzer (2013) and GTD Data over Time

The specification, by using district- and region by time fixed effects takes out any fixed- conflict region and time varying reporting differences, while the district fixed effects remove any time-invariant district specific reporting biases. The coefficients paint a very stark picture: the datasets do not compare well at all before 2007. The good news is that the coefficients are consistently positive, suggesting that the overall correlation is positive. However, the point estimates are very small and only sometimes statistically significantly different from zero. This suggests that int he earlier years it is extremely unlikely for an incident captured in one dataset to appear in the other. In more recent years, the data become increasingly similar.

Why have the two datasets converged? It appears that the underlying data source in the GTD database has evolved significantly over time. Since 2008, the SATP reports feed into the GTD database, while before that the GTD database was mainly fed by newswire services. By 2010, more than 53% of the incidences in the GTD database were directly referenced with a report from the SATP newspaper clippings dataset.

This is clearly, a lower bound since for many reports in the GTD dataset one can manually find references in the SATP dataset, but not necessarily vice versa.

While the level of violence reported in the GTD database seems to be significantly lower for early years, it is important for the identification whether this mismatch in reporting is correlated with rainfall realisations.

In order to explore this, I measure the differences and the absolute value of the differences between the two datasets and run the three specifications from above again.

The results are presented in table 1.22. The coefficients suggest that a positive rainfall realisation in the preceding month is significantly correlated with a lower reporting difference, i.e. implying that the mismatch between the Fetzer (2013) dataset and the GTD dataset is smaller. This highlights that reporting is likely to be endogenous to past weather and thus, past income realisations. While this is something that can fundamentally, not be checked, I believe that this is more likely to be a problem for the GTD database, where reporting has been found to correlate with Foreign Direct Investment in Fetzer (2013). The introduction of NREGA appears to have further reduced the mismatch between the two datasets.

Table 1.22: Evolution of Reporting Differences between GTD and Fetzer (2013) datasets

	Reporting Difference			Absolute Value of Reporting Difference		
	(1) Pre NREGA	(2) Dynamic	(3) Level	(4) Pre NREGA	(5) Dynamic	(6) Levels
Monsoon	-0.078** (0.032)	-0.090** (0.036)		-0.107*** (0.030)	-0.136*** (0.034)	
NREGA x Monsoon		0.051 (0.042)			0.060 (0.043)	
NREGA		-0.398 (0.269)	-0.048 (0.055)		-0.503* (0.278)	-0.094* (0.050)
Observations	12657	25521	27693	12657	25521	27693
Number of Districts	543	543	543	543	543	543

Notes: All regressions are simple linear regressions with time- and district fixed effects. Robust standard errors clustered at the district level are given in the parentheses with stars indicating *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

If we take this and the previous results together, this suggests that there is some systematic differences to the GTD dataset which correlates with rainfall in a systematic way and the introduction of NREGA may have lead to a moderation of this reporting difference. Since the two datasets appear to be converging over time and the coverage of the GTD dataset actually expanding, it seems reasonable to conclude that the SATP data source on which the Fetzer (2013) dataset is a more consistent way to measure conflict.

1.A.6 TRMM Rainfall Data

This paper is the first one in economics to use data from the Tropical Rainfall Measuring Mission (TRMM) satellite, which is jointly operated by the National Aeronautics and Space Administration (NASA) and the Japan Aerospace and Exploration Agency (JAXA). The satellite carries a set of five instruments to construct gridded rainfall rates at very high spatial and temporal resolution.

The TRMM Multi-Satellite Precipitation Analysis provides daily rainfall from 1998 to 2012 at a fine spatial resolution of 0.25 by 0.25 degree grid-cell size. The data from the various instruments aboard the satellite are cleaned and calibrated using additional data from the accumulated Climate Assessment and Monitoring System (CAMS). The output of the algorithm are 3-hourly rainfall rates for that time-period. This is then scaled up to obtain monthly mean precipitation rates, which in turn are transformed into overall monthly rainfall.

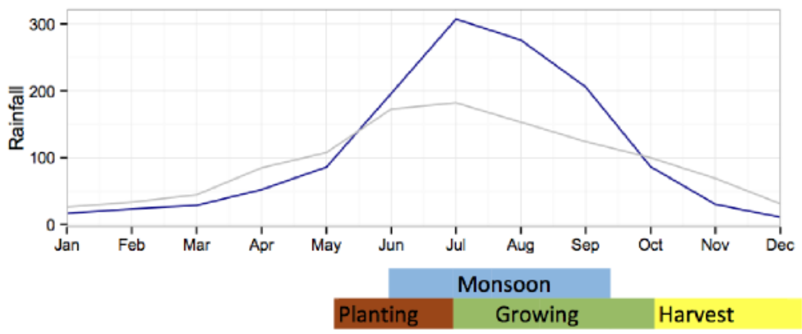

Figure 1.12: Rainfall and Growing Season for Andhra Pradesh

Remotely sensed weather data is an important source of data, in particular, for less developed countries, where observational data is scarce. This is particularly relevant in the case of India, where observational weather may vary in systematic ways. There are three main drawbacks. First, most observations come from rain gauges, where measurements are taken once a day. Climatologist are concerned about rain gauges in particular in tropical- or subtropical areas, since most rainfall is convective. Such convective rainfalls are highly local, generating intermittent and scattered rainfall, which may not be picked up using rain gauges, if the network is not spatially fine enough. The TRMM satellite orbits the earth every 90 minutes, thus providing multiple observations each day. An alternative is to consider data from weather radars. Rainfall radar may provide estimates for rainfall in a radius of 200 km around the station, however it is unreliable for distances in excess of 200 km. In the Indian case, rainfall radar data is not made available and would be problematic, since most reporting radar stations are clustered along the coast. The third general concern regarding observational weather data is the fact that reporting may be endogenous e.g. to violence or other variables that are correlated with the dynamics of violence. This has been highlighted recently by Smith et al. (2011), who show that Somalian piracy has generated a "black hole" in the Indian ocean, where observational weather data from merchant vessels is not available anymore, as vessels take routes avoiding piracy infested areas.[60]

[60] Another example is the case of Vanden Eynde (2011), who had to merge several districts together in

I prefer the TRMM data as it is less subject to systematic measurement error, as the underlying data source is consistent over time. This is not the case with rain gauge based data, such as the GPCC as used by Miguel et al. (2004), Ferrara and Harari (2013) and Kudamatsu et al. (2014) and many others. In the case of India, the number of reporting weather stations for the GPCC data set varies from year to year. In 2001 there were a total of 1197 stations that reported at least some data, while in year 2008 that number dropped to 978. On average, 15.7 % of the district-year observations have some rainfall station reporting data. This pattern varies systematically with violence as is shown in table 1.23. The table presents results from the same specification as in the main part of the paper, including region-by NREGA phase time fixed effects and district fixed effects. The dependent variable is an indicator whether any station reported data for that district and year. The regressor is either an indicator whether a district experienced any violent incident in the last year (column (1)) or the number of incidents in column (2).

Table 1.23: Weather Station Reporting in GPCC Varies with Violence

	(1)	(2)
Any Violence	-0.013	
	(0.009)	
Attacks		-0.002**
		(0.001)
Mean of DV	.157	.157
Observations	5440	5440
Number of Districts	544	544

Notes: All regressions are simple linear regressions with time- and district fixed effects. Robust standard errors clustered at the district level are given in the parentheses with stars indicating *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The coefficient on the violence indicator is insignificant, with a p-value of 18.5%. The coefficient on the number of attacks is significant at 5%, indicating that one additional attack per year decreases the probability of a rain gauge station reporting data in the subsequent year by 1.3% percent, when evaluating it against the mean of the dependent variable. Despite this general concern, my results are robust to using either the GPCC data (Schneider et al. (2011)) or the Indian Meterological Department data used in Vanden Eynde (2011).

1.A.7 Temperature Reanalysis Data

As a solution to the problem of limited data availability for ground measurements, I construct temperature readings from a gridded daily reanalysis dataset that uses remote sensing data and sophisticated climate models to construct daily temperature

order to obtain consistent rainfall estimates, since many stations simply fail to report rainfall estimates. Most of these stations are located in places with conflict or in newly created districts or states.

on a 0.75° (latitude) x 0.75° (longitude) grid (equivalent to 83km x 83km at the equator).[61] The ERA-Interim reanalysis is provided by the European Centre for Medium-Term Weather Forecasting (ECMWF).[62] As the grid is significantly coarser than the rainfall data, I construct inverse distance weighted daily mean temperatures for all grid points within 100 km of the geographic centre of each district. The weighting used is the inverse of the distance squared from the district centroid.

1.A.8 Agricultural Production, State Level Harvest Prices and District Domestic Product

For every district, I only consider crops that have been consistently planted on at least 1000 acres for the whole period that the state reports agricultural production to the data dissemination service of the Directorate of Economics and Statistics with the Ministry of Agriculture.[63] This leaves the following crops: bajra, barley, castorseed, chilly, cotton, gram, groundnut, jowar, jute, linseed, maize, mesta, potato, ragi, rapeseed, rice, sesamum, sugarcane, tobacco, tumeric, tur-arhar and wheat. These capture India's most important staple crops as well as cash crops. Underrepresented is production of fruits or other horticulture products.

For each of these crops, I obtained state-level farm harvest prices to compute a district level measure of the agricultural output value. Unfortunately, district level harvest prices were not available throughout or only for a limited number of crops that did not match well with the actual planted crops. For that reason, I stuck with the state-level prices. The resulting dataset is an unbalanced panel, since not all states consistently report data to the Ministry of Agriculture information systems.

For the quantification exercise on the insurance value, I scale up the district level agricultural output value to match the district domestic product for the year 2000. The district domestic product is an estimate of local area incomes that has been produced for the period 1998-2005, but is not available for more recent years.[64] It relies on a large set of input statistics, including the Annual Survey of Industry, the National Sample Survey and Crop Production Surveys. The district domestic product construction is discussed in detail in Katyal et al. (2001). I obtain a baseline measure of the agricultural output per capita from the district domestic product. This measure will be unambiguously larger than the computed agricultural output value derived from the crop production statistics, as I only include crops that have been consistently reported for the time period that a state reports data to the Directorate of Economics and Statistics. I compute for each district a scaling factor ω_d that measures the share of the agricultural output value per capita that is captured in the agricultural district domestic product. I then simply scale up the agricultural output value per capita by

[61] To convert degrees to km, multiply 83 by the cosine of the latitude, e.g at 40 degrees latitude 0.75 x 0.75 cells are 83 x cos(40) = 63.5 km x 63.5 km.

[62] See Dee et al. 2011 for a detailed discussion of the ERA-Interim data.

[63] This data is available on `http://apy.dacnet.nic.in/cps.aspx`, accessed 14.12.2013.

[64] The data is available from `http://planningcommission.nic.in/plans/stateplan/index.php?state=ssphdbody.htm`, accessed on 21.06.2014.

this scaling factor. This preserves the variation but likely gets the agricultural output value closer to the true. This scaled agricultural output value per capita will be used for the quantification exercise to evaluate how much insurance NREGA provides.

1.A.9 Agricultural Wages in India

This appendix describes the process of how the agricultural wage data was cleaned and put in shape for the analysis in the paper. The data is only source providing consistently reported wage data for the whole of India. The raw data gives monthly wages for male, female and children, broken into skilled- and unskilled agricultural labour and different types of labour. The types of skilled labour are blacksmith, carpenter and cobbler, while unskilled labour combines ploughman, reaper/harvester, sower, weeder, other agricultural labour. In some states, these separate unskilled labour categories are not reported[65], but rather, a category "Field Labour Wages" is reported. This is conceived to be an average of the different categories.

In some districts these wages are reported throughout the year, while in others the wages are reported only in the parts of the year, when particular activities are actually carried out (i.e. sowing wages in the early Kharif season of May, June and July), while harvesting wages are reported in the fall of a given year.

After digitising and entering the raw data, I proceed to construct an annual level agricultural field-labour wage as my main dependent variable. For each district, there may be multiple wage-observations in case there are multiple reporting centres. I generate a balanced panel requiring each quarter of the year to have at least one non-missing observation of agricultural wages belonging to the particular category of unskilled labour. I then construct the simple average across these wage-observations.

There are advantages and disadvantages to this approach. In particular, by construction, this implies that within a year, some field labour wage observations are noisier then others. This can be taken into account by adequately weighting the observations. As an alternative, I can impose the requirement that there be at least one observation for each different unskilled labour category within a quarter. This condition is very stringent, as it fails to recognise the types of agricultural activities that are pursued during a year. This approach reduces the number of districts significantly, but the results remain the same.

1.A.10 NREGA Data Sources and Roll Out

The data for the roll-out of NREGA come from the Ministry of Rural Development, which is responsible for administering the scheme. The sequence of roll-out was highly endogenous to a set of district level characteristics, such as the share of scheduled caste, scheduled tribe population, baseline agricultural productivity, literacy and existing levels of conflict. This becomes obvious when considering Figure 1.13. This picture highlights that a lot of districts in the east of India received NREGA in the

[65] The states for which this is the case are Andhra Pradesh, Karnataka and Maharashtra.

first round. A lot of these districts did suffer from Naxalite violence. As discussed in the main body, I do not require exogeneity of treatment to levels of violence for my empirical design. There are two main sources for data on NREGA take-up. These are the district-level monthly-progress reports (MPR) and data coming from the Management Information System (MIS). The latter is a completely non-paper based system that has only become mandatory to use in the financial year but was still not fully operational until 2010-2011.

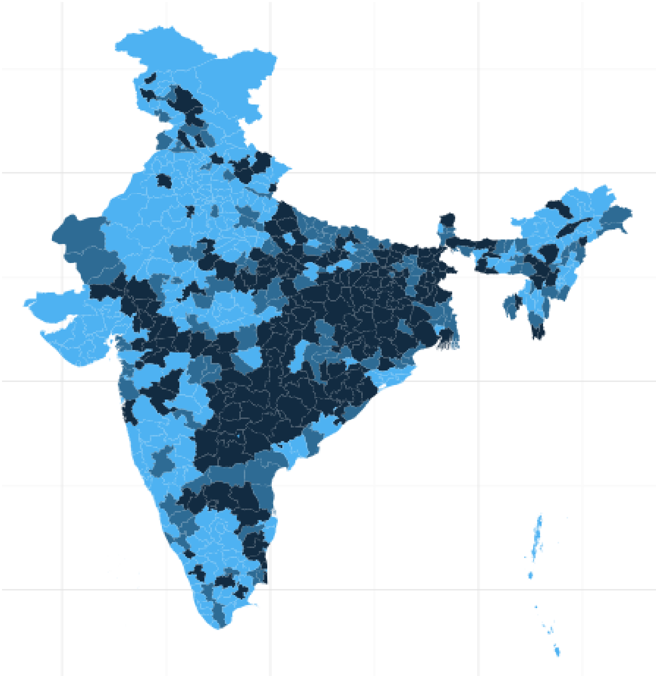

Figure 1.13: Phases of the NREGA Roll-out across India

There are a lot of issues regarding the reliability of either datasets, as there is quite some mismatch between the two datasets, especially in the earlier years when the MIS was introduced.[66] This may be due to partial compliance in the MIS after it had been introduced, but could be also because the MPR system is more subject to manipulation. It is difficult to asses the underlying divergence in the two databases.

The MPR data is available continually from 2006 to the financial year 2010-2011, from which point onwards I rely on data from the MIS.[67] The format of the reports has changed considerably, with the major break occurring in 2011. This is partly due to the evolving nature of NREGA. Ministry of Rural Development (2009) details that several programs by the Ministry of Water Resources are to be joined with the NREGA by 2011. An important part of this program are rural sanitation projects that are funded by the Ministry of Water Resources for a set of targeted districts. This implies that there are district-specific breaks in the NREGA data. In the empirical

[66]See for example mismatch between MIS data and National Sample Survey returns data highlighted by `http://www.indiatogether.org/2013/jun/gov-nregs.htm`,accessed on 12.06.2013.

[67]Thanks to Clement Imbert for sharing NREGA MPR data for the earliest years.

specifications which combine data from before and after 2011, I flexibly control for these breaks by allowing the district fixed effects to be different before and after 2011.

I focus on a set of variables measuring take-up, project expenditures and overall expenditures at the district level. For the take-up I study cumulative person days provided, cumulative number of (distinct) households provided employment as well as the number of days per household at the district level. I also look at the number of person days for scheduled caste and scheduled tribe populations, as well as the share of person days that accrue to females.

For the NREGA project measures, I study the total cost or number of ongoing projects at the end of each financial year.[68] For overall expenditures, I study total expenditure in a district and year or total labour expenditures.

Despite having access to NREGA for many months in a financial year, I only study the reported metrics at the end of each financial year (that is March of each calendar year). This becomes necessary as there are significant reporting delays which induce large jumps in the cumulative month on month measures which are less likely driven by participation, but more likely due to reporting issues.

I construct the NREGA take-up, participation and project data to match the Monsoon calendar as in the main exercises.

[68]The categories in the data that are consistently reported are: "Micro Irrigation Works","Drought Proofing","Water Conservation and Water Harvesting","Provision of Irrigation facility to Land Owned by Scheduled Caste/ Scheduled Tribe".

Chapter 2

The Welfare Cost of Lawlessness: Evidence from Somali Piracy

For centuries, piracy has posed a threat to ocean-going trade. In essence, it is organized private predation which thrives in locations in which law and order is weak, either because particular states provide a safe haven or due to poor international cooperation. And it has repercussions for worldwide trade.[1]

However, despite the long-standing importance of piracy, little is known about its economic costs.[2] The issue has been brought into sharp relief by the upsurge of piracy in the Gulf of Aden which poses a threat to one of the world's busiest shipping routes. Frequently attributed to the collapse of effective authority in Somalia, it has provoked an international response.

We match data on piracy attacks in the maritime area around Somalia to data on around 24,000 shipping contracts by constructing the closest navigable sea distance between each origin and destination port for which a ship has been chartered. This allows us to exploit the monthly time-series variation in the frequency of piracy attacks in the main areas affected by Somali piracy to estimate the impact of piracy on shipping costs. We then use these estimates to calibrate a model of the welfare cost of Somali piracy.

Figure 2.1 previews our findings by showing the relationship between piracy attacks in Somalia and a non-parametric estimate of the additional shipping cost paid on routes through the piracy area.[3] There is a visible association between the two variables. Both shift upwards in mid 2008 after the maritime area is declared a piracy risk area by the maritime insurance industry in May 2008.

[1] For example, North (1968) argues that a decline in piracy from 1600 to 1850 accounts for a significant proportion of the observed productivity increases in transatlantic shipping in this period.

[2] Bensassi and Martinez-Zazosa (2010) study the impact of piracy in the Strait of Malacca on trade costs. Most cited numbers are from the One Earth Future Foundation (2011) reports. Our direct approach is distinct from these reports. A recent World Bank (2013) report calculates the welfare effects with a gravity trade model but finds mostly insignificant effects of piracy on trade.

[3] We constructed Figure 1 by regressing shipping costs on route and time fixed effects and a set of time dummies for those trade routes going through the Somalia area. The coefficients on these dummies allow us to draw charter rate differentials across time. Figure 2.1 shows the rolling average of the estimated coefficients of this regression together with the rolling average of attacks.

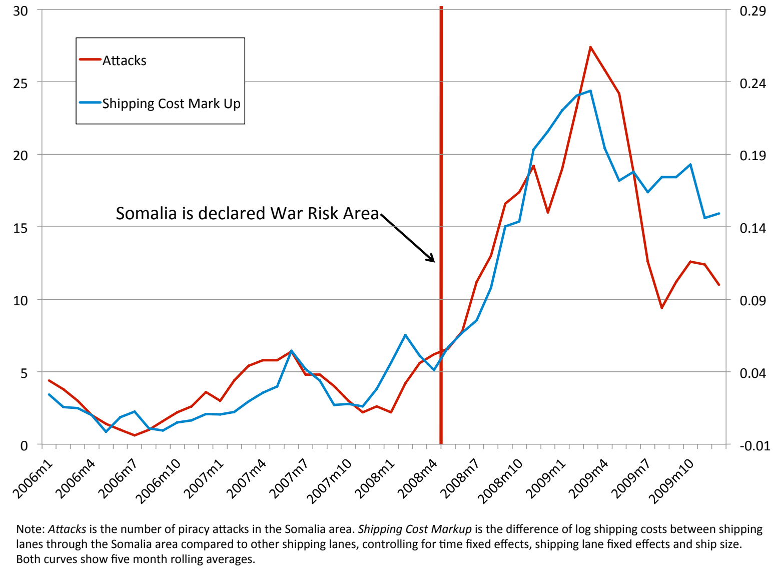

Note: Attacks is the number of piracy attacks in the Somalia area. Shipping Cost Markup is the difference of log shipping costs between shipping lanes through the Somalia area compared to other shipping lanes, controlling for time fixed effects, shipping lane fixed effects and ship size. Both curves show five month rolling averages.

Figure 2.1: Non-Parametric Visualization of Piracy Effect on Chartering Rates

Our regression results show that shipping costs for dry bulk goods rose by between 8 and 12 percent when pirate activity increased in Somalia. We also show that these larger shifts mask significant variation across months. Charter rates fluctuate by 18 percent between the most and least dangerous months. This seasonal pattern in shipping prices is absent prior to the upsurge in pirate activity in the region during 2008. Accounting for this seasonal variation highlights that the average shipping costs through the Somali area did not increase during the months in which weather conditions inhibit pirates from operating.

The extra shipping costs that we uncover are mostly due to higher insurance costs and the increased security measures that are needed to repel pirate attacks. These constitute a welfare cost to the extent that labor and resources are allocated from productive tasks towards protection. Our model compares the extraction of resources through pirate attacks to a tax on shipping which finances an equivalent transfer. This allows us to calculate the welfare loss caused by piracy. Our central estimate suggests that the resource costs incurred in transferring around 120 million USD annually to Somali pirates is well in excess of 630 million USD.

Studying Somali piracy provides a unique opportunity to measure the costs of economic predation. Moreover, the factors that lie behind the welfare costs in this context are generic. In particular, it is useful to reflect on why taxation is less costly than predation. Ideally, a state that levies taxes has the capacity to ensure compliance and to commit to providing security to those who pay those taxes. Economic predators typically lack both of these capacities. Somali pirates can extract resources only by attacking ships while ship owners only have the option to invest in defence or bear the cost of predation. We show empirically that, in this situation, large costs can be occurred even when the amount extracted from predation is fairly small.[4]

This article belongs to a wider literature on the value of establishing the rule of law and its role in securing trade and investment.[5] A traditional problem in weakly-institutionalized environments is that bringing goods to market is subject to predation and theft.[6] The consequences of the failure to establish and enforce property rights is a core theme in the development literature such as Knack and Keefer (1995) and Acemoglu et al. (2001). Piracy is a specific consequence of state failure because it creates a spill-over of insecurity from one country to a maritime region. We show that in the case of Somalia this has taken on striking dimensions with shipping through the whole Indian Ocean now affected. We show that the consequent predation generates sizeable costs relative to the revenues that it raises for pirates.

A recent literature has studied the economic effects of an extreme case of state failure, namely violent conflict.[7] Guidolin and La Ferrara (2007) provide the example of diamond mining companies benefiting from local conflict. Besley and Mueller (2012) provide a framework to capture the effect of expected violence on housing prices which we use for our estimation. Voors et al. (2012) show that violence in Burundi affected individual preferences permanently. In particular, they find that individuals that were exposed to violence became more risk seeking. Disruptive, high risk activities, like piracy, are therefore more likely to arise in a conflict setting.

Piracy poses a particular issue because of the difficulty of securing international agreement over the assignment of responsibility to deal with the problem and how the costs of such efforts are to be shared. Private solutions to increase security such as carrying guards aboard ships are inherently less efficient compared to dealing with the public good of security for all. Our calculation of the welfare cost gives a sense of the magnitude of this benefit.

Insecurity due to piracy leads to a rise in shipping costs which are an important part of total trade costs. In this respect, our paper relates to studies of the consequences of trade costs for trade patterns. In particular, it is related to Mirza and Verdier (2008) which studies how international terrorism affects trade costs.[8] Our model allows us to calculate the likely impact of the estimated increase in shipping costs on trade. For this purpose we use recent findings by Feyrer (2009) who studies the Suez Canal closure 1967-1975. It has been argued in the context of Somali piracy that it has reduced shipping and led to a re-routing of ships.[9] We show em-

[4]Our arguments are akin to the distinction between roving and stationary bandits in Olson (1993). Bandiera (2003) argues that fractionalized ownership reinforces this problem in the context of the Sicilian Mafia.

[5]See Dixit (2004) and Rose-Ackerman (2010) for literature overviews.

[6]Anderson et al. (2002) and Anderson and Bandiera (2006) study the link between predation and trade. Olken and Barron (2009) study predation in the context of trucking in Indonesia.

[7]See Blattman and Miguel (2010) for an review of the literature on civil war.

[8]For reviews of the extent of trade costs and their importance in explaining patterns of trade see Anderson and van Wincoop (2003), Behar and Venables (2011) and Hummels (2007). Donaldson (2010) is a recent study of the impact of a change in trade costs due to the construction of railroads in India.

[9]One Earth Future Foundation (2010) calculates large costs from re-routing around the cape of Good Hope. This cost is dropped in the One Earth Future Foundation (2011) report which argues that re-

pirically and theoretically that effects on trade volumes related to piracy have likely been small.

The remainder of the paper is organized as follows. In the next section, we discuss the background to both our piracy and shipping cost data. Section three presents our estimation procedure and discusses the results while section four provides a framework for thinking about the welfare loss and uses this, along with our estimates, to develop estimates of the welfare loss from piracy. Concluding comments are in section five.

2.1 Background and Data

In this section we discuss our data on piracy and shipping costs. We present potential channels for piracy to affect these costs. We also discuss how susceptibility to piracy can be matched to specific shipping routes.

2.1.1 Piracy Data

Our data on piracy attacks comes from the ICC International Maritime Bureau (IMB) annual reports which provide the exact position of the attack, details on the ship and its status (anchored or steaming) and the type of attack (attempted, boarded, fired upon, hijacked).[10]

We geo-code attacks and focus on the Somali area which we define as the rectangle spanned by the coordinates S11, E38.4 and N18.3, E74.7 depicted as the shaded area in Figure 2.2. We focus on this area because we believe that there are common factors driving piracy attacks within this zone, i.e. if pirates attack in some point along the Somali coast, it is informative about the likelihood of an attack elsewhere within the area. The crosses in Figure 2.2 represent the locations of the piracy attacks. Figure 2.2 also depicts a geographically narrower area in a darker shade, the Gulf of Aden, which we use as a robustness check on our main results below. Piracy in the Somalia area is a sophisticated crime with a large number of ships being hijacked. Pirates rely on external finance, political support and safe havens on the Somali coast to operate effectively.[11]

routing around the cape was unlikely to be an issue.

[10] We discuss our data in more detail in appendix 2.A. Table 2.8 provides summary statistics.

[11] A previous draft of this paper studied piracy in the broader Indonesia area. However, the type of piracy which takes place there is distinct from Somali piracy. It consists mainly of armed robbery, which takes place in ports. Hence, arguably its consequences are less severe and are easier to control.

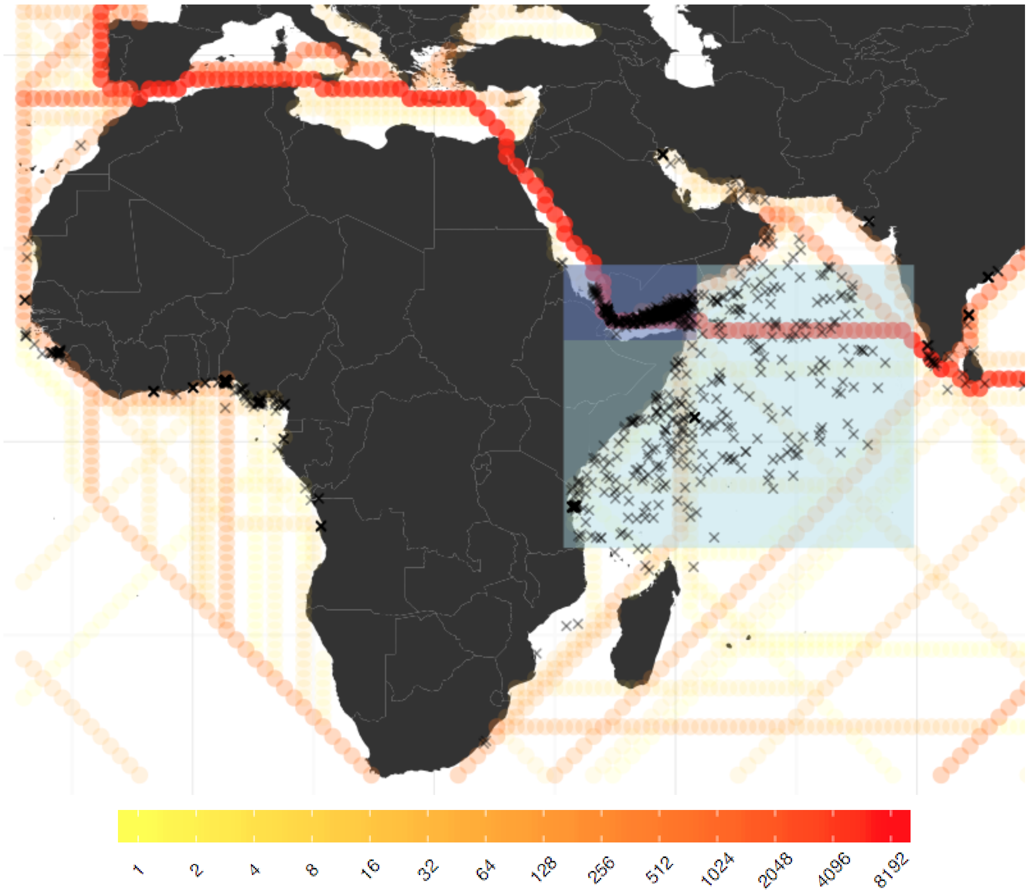

Figure 2.2: Calculated Shipping Lanes and Treatment Areas
Note: The light shaded rectangle is the "Somalia" treatment area, while the darker shaded area is the "Gulf of Aden" treatment area. The location of attacks is indicated by a cross. The circles indicate the shipping lanes, the colouring of which is proportional to the number of observation on each shipping lane according to the continuous colour scheme.

Figure 2.3 illustrates the time-series variation in piracy attacks, showing the upsurge in attacks during 2008. We exploit this to study the effect of Somali piracy on shipping costs. Interpreting this as an effect of piracy requires us to be sure that there was no change in amounts shipped due to piracy during 2008. We show in section 2.2.5 that, if anything, shipping through Somalia decreased during 2008 making it highly unlikely that changes in traffic patterns were responsible for the increase in pirate activity. There is a consensus among experts on Somali piracy that the origins of the increase in pirate activity lie in what happened on land rather than at sea. Hansen (2009), for example, argues that a key trigger for the increase in piracy attacks was the crisis in public finances in the Puntland government in Somalia which left it unable to pay the police. This, he argues, along with the generally weak state of law and order in Somalia, made it increasingly feasible for pirates to operate without sanction. Pirates had long masqueraded as coast guards protecting Somali territorial waters from illegal fishing. This cloaked a build up of organized violence which emerged strongly after May 2008.

The developments were closely observed by the maritime insurance industry. Ta-

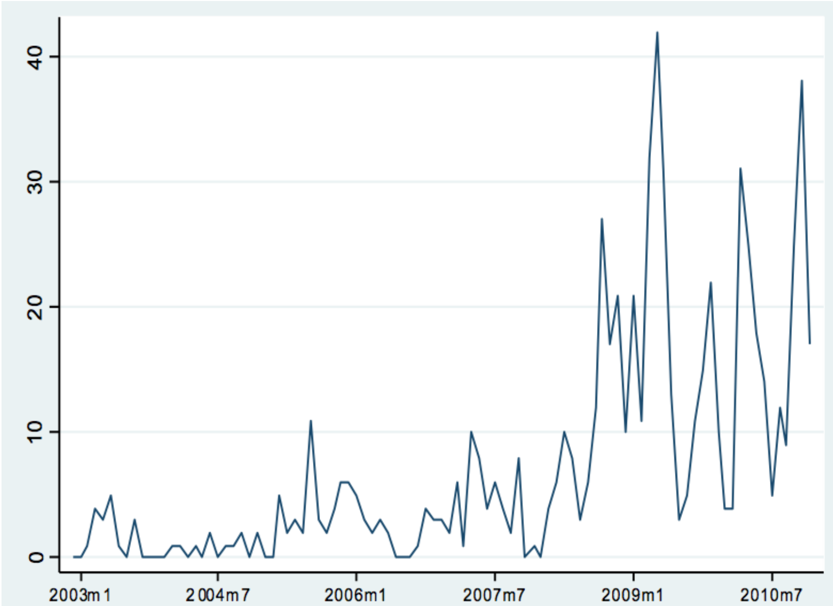

Figure 2.3: Time Series of Attacks in Somalia

ble 2.1 summarizes the piracy data around the date that the Somalia area was declared a *war risk area* by the maritime insurance industry (May 2008). The average number of attacks increased from 2.8 attacks per month before that date to 17.1 attacks per month from May 2008 onwards.

Aside from the structural break, seasonality induced by wind conditions plays a crucial role in the pattern of piracy, something which we will exploit in our empirical analysis. Most of the attacks are carried out using small vessels, known as "skiffs". These are typically between 7 and 10 metres long and at most two meters wide with a low freeboard. This renders them particularly vulnerable to wind and waves. The summary in Table 2.1 illustrates the resulting seasonal pattern. The post May 2008 column features a strikingly low piracy risk in the Monsoon months of July and August, for example. In these months the level piracy attacks is rather similar to pre May 2008 levels. The calm spring period is the most dangerous time with over 30 attacks in March and April. The close link between this seasonal pattern in attacks and wind speeds is discussed in more detail in Appendix 2.A.1.

2.1.2 Shipping Cost Data

Our shipping cost data comes from the web-site of N. Cotzias Shipping Consultants which provides monthly reports on the time charter market for the period November 2002 until December 2010.[12] The data is comprised of 33,529 individual charters in

[12]In early 2011, Cotzias merged with Intermodial (www.intermodal.gr). As of 25th July 2012, the Cotzias data was available on `http://www.goo.gl/g5d0c`. There are many shipping consultants, however, Cotzias consistently made data available for a long time period. The selection of a particular shipping consultant will only affect our results in case there is a time varying bias to reporting charter contracts on Somalia routes that is correlated with the onset or intensity of piracy. We do not believe this is the case.

the dry bulk cargo segment of the market. These are ships that transport primary commodities such as iron ore or agricultural products such as grain. These types of vessels constitute approximately one third of the tonnage of the global shipping fleet. Short term chartering agreements are typical for bulk carrier ships, due to the volatile nature of commodity markets. Since the starting point for these charter agreements are previous agreements ('last done'), shipowners and charterers take an active interest in reports of recent transactions.[13] The individual time charter agreements are also used to construct general shipping indices such as the Baltic Exchange Dry Index (BDI). Thus our data-set provides a window onto the wider shipping market.

In a time charter agreement the shipowner places his ship, with crew and equipment, at the disposal of the charterer and bears the costs of keeping the ship operational. The charterer pays a daily charter rate and decides the type and quantity of cargo to be carried and the ports of loading and discharging. The charterer is also responsible for paying for fuel (known as *bunkers*) and costs like port charges including the payments due, for example, for using the Suez Canal. The fact that time charter rates are provided on a daily basis makes them comparable across contracts of differing lengths.

The summaries made available on the web-site provide, among other information, the name of the ship, its deadweight tonnage (DWT) - a measure of ship size, the year it was built, the port or country of origin and the port or country of destination. From this information we construct our measure of shipping cost - the rate per day per DWT. We also use the origin and destination to assign the ship's voyage to countries (see Appendix 2.A.1). Our data set contains information on around 1600 distinct shipping routes. Most of the charters are from Asia with China making up the bulk of origin and destination locations.

2.1.3 Piracy Risks and Shipping Costs

There have been a number of private responses to the piracy threat. A variety of insurance arrangements have emerged to cover piracy risks with higher premia being paid to travel in areas deemed to be at risk. Ships increasingly carry armed guards and other preventive measures (mostly modifications to ship hulls) have become "best practice" which makes them relevant for insurance purposes.[14]

The costs to the shipping industry can be decomposed into five main categories: (i) damage to vessels (ii) loss of hire and delay to cargo delivery while a ship is held to ransom (iii) costs of defensive measures (iv) cost of ransoms and negotiators fees paid when a crew is kidnapped or a vessel is held (v) re-routing, speeding-up of vessels to avoid areas at risk (vi) extra wages paid to the crew compensate for the risk of being kidnapped. We discuss these cost factors in detail in Appendix 2.A.3.

[13] See Stopford (2009) for a detailed discussion of the time charter market.

[14] Best Practice manuals are published and updated regularly by the shipping industry. See `http://www.goo.gl/zL1Ut`, accessed on 10.04.2012.

Ship owners typically buy insurance to cover themselves against a number of these costs with insurance costs being sensitive to developments in the number of piracy attacks. Throughout the paper we assume a competitive insurance industry.[15]

We use shipping contracts to measure the cost of shipping. These reflect the consequences of piracy to the extent that costs of piracy are borne by the ship owner and passed on to the charterer. This is not unrealistic. The association of independent tanker owners, for example, provides model clauses for chartering agreements with regard to piracy risks, stating that:[16]

> "Charterers shall indemnify Owners against all liabilities costs and expenses arising out of actual or threatened acts of piracy or any preventive or other measures taken by Owners [...], including but not limited to additional insurance premiums, additional crew costs and costs of security personnel or equipment."

Hence, charterers have to compensate ship owners for extra costs created by piracy risk on the chartered route. However, it is still possible that some of the pirate costs are borne directly by the charterer which would result in us underestimating the cost of piracy. In section (2.2.5) we therefore discuss the sensitivity of our welfare estimates to the exact division of piracy costs between ship owners and charterers. Specifically, we calculate the welfare cost under the assumption that piracy costs are shared according to the General Average (GA) rule which is widely used in the shipping industry and is explained below.

2.1.4 Identifying Exposure to Piracy Risks

We assign a risk of exposure to piracy attacks to each shipping route by using the information on the origin and destination of the shipping contract. For example, a vessel with a destination in Germany and an origin in China is quite likely to travel through the Somalia area. However, there are some cases where it is not entirely clear whether the vessel would travel on a Pacific route or through the Indian Ocean and Atlantic using the Suez canal. In assigning piracy risk to a specific route, we employ a path algorithm to obtain an automatic coding of that route.[17] We are then able to see whether the shortest sea route passes though the piracy areas that we study. If it does, we suppose that the shipping contract is subject to a piracy risk based on the forecast number of attacks in the relevant region at a point in time.

Figure 2.2 provides a bird's-eye view of the trade-routes for the areas around Somalia based on our path algorithm. The points which are less opaque and more deeply shaded in red represent more ships going through a particular route. We suppose that a shipping route is more vulnerable to piracy attack if it crosses the rectangles in Figure 2.2. As a check on our core results, we construct a measure of

[15] There are debates about whether this assumption is reasonable. If it were not the case then markups in this industry would create a further potential welfare cost from piracy.

[16] Refer to `http://www.goo.gl/yShgs`, accessed on 10.04.2012.

[17] Details are discussed in the appendix 2.A.1.

a route being vulnerable to piracy attack based on it passing through the convex hull which is spanned by all such attacks up to each year. This measure is arguably more satisfactory since it takes into account the fact that the Somali pirates were able to expand their reach into the Indian Ocean since 2008. The empirical findings are similar when either method of assigning vulnerability to piracy attacks is used.

It is possible that some ships re-routed around the Cape of Good Hope to avoid exposure to piracy risks. We check for this possibility below and find no evidence for changes in either the extent of traffic through the Suez Canal or in the composition of ship size through affected areas after the upsurge in piracy attacks. Moreover, assigning piracy risk to routes allowing the possibility of re-routing when this would add relatively little distance to the journey, makes our results even stronger. This supports the view of other commentators, such as One Earth Future Foundation (2011), that re-routing around the Cape in response to piracy is not important.

We do not distinguish between attacks on different types of vessel (container, tanker, dry bulk, etc.) since all varieties of ship, including all sizes, have been attacked and hijacked in the piracy-affected area. The first successful hijack of a dry bulk ship took place as early as May 2008.[18] Attacks seem sufficiently random across of a range of ship types and so we do not to attempt to distinguish empirically between different bulk ships.

2.1.5 A Model of Piracy Attacks

To motivate the time-series variation in piracy attacks, consider the following simple theoretical model. Suppose that there are M active pirate ships and that in each period each pirate receives an opportunity to hijack a ship where V_{it} is the benefit and c_{it} is the cost.[19] Pirate i at date t will launch an attack if the expected benefit exceeds the cost:

$$\xi_t V_{it} \geq c_{it}$$

where ξ_t is the success probability, V_{it} is the value of a successful attack and c_{it} is the cost.

A key parameter is the cost-benefit ratio $\rho_{it} = c_{it}/V_{it}$. We suppose that ρ_{it} is drawn for each pirate ship i at date t from a uniform distribution with mean θ_t. Given M

[18]According to a Lloyds List report on July 2008 the ship was freed 41 days later for a ransom of 0.75 million USD.

[19]To endogenize M, suppose that there is a fixed cost becoming an active pirate. Then we would have that a pirate will enter if

$$E\left\{V_{it} - c_{it} : \xi_t\right\} > F_i$$

in which case we would also predict that M would be a function of ξ_t, i.e.

$$M_t = H\left(\xi_t\right).$$

So we would have

$$E\left[a_{rt}\right] = \xi_t H\left(\xi_t\right)$$

and the expected number of pirate attacks will still depend on ξ_t reflecting underlying law and order.

independent draws the expected number of pirate attacks at date t is given by:

$$E[a_t] = \xi_t M. \tag{2.1}$$

The variation in expected piracy attacks in equation (2.1) is then captured by ξ_t which we assume reflects two things. First, there can be short-term factors which shape piracy costs and benefits, including weather variation. Second, there can be persistent changes in law and order as we saw after after the break down in law order in Puntland in 2008 which lead to a permanent shift in the feasibility to conduct piracy. To capture these two factors we allow the success probability, ξ_t, to be related empirically to climatic conditions and the insurance evaluation of the industry which requires ships to insure against war risks since May 2008.

2.2 The Effect of Piracy on Shipping Costs

In this section, we present estimates of the effect of piracy attacks on shipping costs. We will begin with a comparison of mean shipping costs between regions affected by Somali piracy before and after the upsurge in attacks in 2008. We then present regression-based estimates.

2.2.1 Difference in Difference Estimates

We present a simple difference-in-difference estimate of the effect of Somali piracy on mean shipping costs by looking at the routes affected by piracy before and after May 2008 compared to all other routes. The result of this exercise is reported in Table 2.2 which gives the average rate per DWT on routes which pass through the Somalia area compared to other shipping routes before and after May 2008.

Column (1) shows that the average shipping costs were not significantly different between routes before May 2008. However, they diverge after that date with the mean cost per DWT being significantly above the rate for other routes by 0.074 USD per day per DWT. This represents an increase of around 15%. This result parallels the finding in Figure 2.1 which also compared affected routes before and after piracy began. The key identifying assumption is that the influence of other time-varying factors which are affecting shipping costs have a common impact on both sets of routes. In particular, the global recession which led to a fall in trade and shipping rates in winter 2008 is assumed to influence routes that are affected by piracy and those that are not to the same degree.[20]

We now turn to investigating how this finding holds up in regression evidence based on individual shipping contracts.

[20]We run a number of robustness checks with regard to changes in the economic environment in section 2.2.4.

2.2.2 Piracy Attacks and Shipping Costs

Our core regression specification assumes that the dry bulk shipping market is contestable so that pricing is based on the average cost per day for each voyage.[21] We would then expect prices in that market to reflect expected piracy attacks and any other factors that influence costs.

We denote the cost per dead weight ton (DWT) per day for a ship of size s on route d in month t as:

$$C\left(s, d, t, A_{dt}\right)$$

where A_{dt} is the forecast number of attacks affecting route d at date t.[22] An effect of piracy on costs is not unrealistic as the shipping conditions at so-called "choke points" (the straits of Hormuz and Malacca, the Suez and Panama canals, the Bosporus) are known to affect freight rates. Since there are scale economies in shipping, we expect this cost function to be decreasing in s.

For simplicity, we adopt the specification:

$$\log C\left(s, d, t, A_{dt}\right) = c\left(s, d, t\right) + \gamma A_{dt} + \beta x_{dst} + \eta_{dst} \tag{2.2}$$

where γ is the core parameter of interest, x_{dst} are other time varying controls and η_{dst} captures other idiosyncratic factors which are uncorrelated with A_{dt}.

The cost from piracy depends on the route that the ship takes. As we have already discussed, we construct a treatment indicator for each route depending on whether it passes through the area of Somalia. Denote this as a dummy variable where $\delta_d = 1$ if route d passes through piracy. Then:

$$A_{dt} = \delta_d \times a_t.$$

is our measure of the cost shock expected on route d where, in the core specification, a_t is the recorded level of pirate attacks in the Somali piracy area in month t. In the basic specification, we do not the effect of piracy attacks to vary with ship size, s, or route, d. However, we will allow for a heterogeneous effect in some specifications that we report below.

This baseline specification, in effect, supposes that the best estimate of piracy en route is the level of piracy attacks in the current month, i.e. $E\left[a_{t+1}\right] = a_t$. This is somewhat implausible to the extent that there are known seasonal patterns and other understandable features of the time series. Hence, below, we will consider some

[21]Shipping has the classic conditions for a perfectly contestable market: (i) no entry or exit barriers (ii) no sunk costs and (iii) access to the same level of technology (to incumbent firms and new entrants). This is essentially the model of the Bulk shipping market used in Kalouptsidi (2014) who also assumes competitive freight rates. See Behar and Venables (2011) for a discussion of the extent of contestability in shipping markets. In Appendix **??** we show that with constant pass-through we can identify the effect of piracy attacks on shipping costs from changes in rates, even if shipping markets are not perfectly competitive.

[22]Due to the absence of good monthly data on ship traffic for our period 2002-2010 we have to use A_{dt} as a measure of piracy risk. This disregards the fact that dense traffic makes journeys less risky for each ship.

alternative models for the expected level of piracy attacks.[23]

To reflect this discussion, our core empirical specification is:

$$z_{isdt} = \alpha_s + \alpha_d + \alpha_t + \gamma A_{dt} + \beta x_{dt} + \varepsilon_{isdt} \tag{2.3}$$

where z_{isdt} is the (log of) daily charter rate per DWT for contract i on a ship of size s, for route d in month t. The parameters $(\alpha_s, \alpha_d, \alpha_t)$ are fixed effects for ship size, route and month. The standard errors ε_{isdt} are adjusted for two-way clustering on origin- and destination country. Other controls in x_{dt} include the age of the ship and the ballast bonus per DWT (a bonus paid for empty return journeys).

Our key identifying assumption is that factors that drive piracy, the factors in ξ_t in equation (2.1) are orthogonal to other drivers of shipping costs, conditional on the controls that we use. Month fixed effects, α_t, for example, should capture changes in the operating costs which affect all routes. The fact that bulk shipping is a competitive world market makes the inclusion of these dummies particularly important.

The main parameter of interest is γ which we interpret as the additional shipping cost from anticipated piracy attacks. We are expecting that $\gamma > 0$. The empirical approach can be thought of as a difference-in-difference specification where ships that pass through a region where pirates are expected to attack are compared to ships using different routes over the same time period. This exploits monthly time-series variation in piracy attacks.

2.2.3 Core Results

Our core results are reported in Table Table Table 2.3 which uses the specification in (2.3). We normalize the piracy attacks variable in columns (1) to (2) such that the coefficients can be interpreted as the percentage point increase in shipping costs with the shift in pirate activity around May 2008.

In column (1), the only controls are fixed effects for route, time and ship size. For the latter, the omitted ship size category is "small" Capesize ships between 80,000 and 150,000 DWTs. There is a strongly significant positive coefficient on the expected number of attacks. The point estimate says that shipping costs were around 8.2 percent higher after the upsurge in piracy.

The ship size dummy variables show evidence of significant scale economies in shipping with the smallest ships being around 62 percent more expensive per DWT than the excluded category. The point estimates decline across the ship size categories. This is a feature of all the estimates that we show.

In column (2), we add the additional ship controls: ballast bonus payments and the vessel's age. We find a large variation in rates paid for younger compared to older vessels with chartering rates for older vessels being significantly lower. However, the point estimate on piracy attacks does not change much after adding these controls.

[23] We discuss the prediction of pirate attacks in detail in appendix 2.A.2.

As we discussed in section 2.1.1 piracy attacks after May 2008 were highly seasonal. We now ask whether this seasonal variation in attacks affects shipping costs. There are good reasons to believe that seasonal variation in risk is relevant for charter rates. Supplementary insurance to pass through high risk areas, for example, is priced based on specific weeks in high risk zones.[24] Other cost factors such as security crews and ship modifications are adjustable as well.

One way to exploit the seasonality in attacks is presented in column (3). Here we identify the effect of piracy attacks only with data after May 2008. The coefficient on piracy is still positive and significant but somewhat smaller in size. Thus, our findings in column (1) and (2) are not entirely driven by changes on routes through Somalia before and after 2008 but also by month-to-month variation within the years with pirate activity.

Declaring an area as a special war risk area is a significant event in the insurance industry and reflects risk perceptions at the time. So instead of using the level of piracy attacks, we can simply use these dates. The representative of the marine hull war insurance business in the London market, the Joint War Committee, added the Gulf of Aden in May 2008. We use a dummy variable to represent this event in equation (2.3) instead of the level of piracy attacks.[25] This specification is bound to capture the sharp increase in costs depicted in Figure 1. The result is in column (4) of Table 2.3. The coefficient on the war risk dummy suggests a 12.3 percent increase in shipping costs around May 2008.

A striking feature of the pattern of attacks is how closely they match with wind speed in the area. In order to exploit exogenous variation in wind speed we create an interaction term between the treatment dummy of column (4) with the monthly average wind speed in the Somalia area. We code the wind speed variable such that it goes from a value of 0 at the maximum wind speed in June to a value of 1 with minimum wind speed in March. In this way the coefficient can be interpreted as the difference in shipping costs between months with maximum and minimum wind speed. The resulting coefficient in the third row of column (5) suggests that charters through Somalia after May 2008 were about 18 percent more costly in March than in June.[26] Moreover, the coefficient on the war risk dummy itself is insignificant suggesting that it was not significantly more costly to charter ships through Somalia sea area in June than it was before the rise in piracy in 2008.

The interaction between the wind speed variable and the indicator that a route is susceptible to Somali piracy is negative. Thus, if anything, there has been the opposite seasonal pattern in charter rates on Somalia routes in the absence of piracy. The fact that pirate activity introduced a seasonal pattern that did not appear previously

[24]The absence of seasonal variation in charter rate differentials would provide opportunities for arbitrage in the insurance market.

[25]May 2008 is in the confidence interval based on our structural break analysis (presented below) but does not coincide with the break date that we found which is July 2008. The results are similar if we use a dummy variable that is equal to one at this slightly later date.

[26]There is a lag between windspeed and piracy attacks of one month. This implies that windspeed at the time of the charter is a good predictor of piracy attacks on the charter route. See Appendix 2.A.1 for details.

adds further credibility to the claim that piracy influenced shipping costs.

Figure 2.4 plots the fitted values from column (5) Table 2.3. It shows the shipping cost predicted by the Somalia war risk, wind speed and their interaction. The graph illustrates the sharp increase in seasonality in costs after May 2008; shipping costs are roughly twice as high when wind conditions favor piracy attacks after this date.

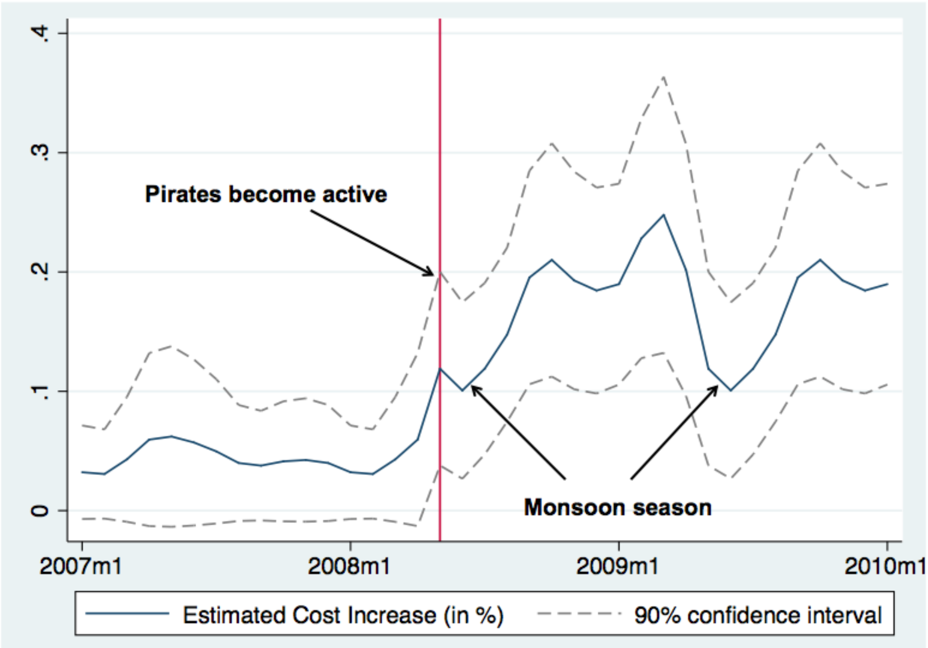

Figure 2.4: Shipping Cost Prediction of Pirate Activity and Monsoon Season

Overall, these results suggest that piracy in the Somalia area has a positive effect on the cost of shipping through this region. The effect is consistent with an average increase in shipping costs of between 8 and 12 percent in the period after piracy attacks increase off the coast of Somalia.

2.2.4 Robustness

In this section we look at the robustness of our results to alternative ways of forecasting piracy attacks and discuss additional controls for economic conditions. We also explore alternative definitions of exposure to piracy risk.

A Markov Chain Model for Piracy Attacks Our baseline specification, in effect, supposes that the best estimate of piracy en route is the level of piracy attacks in the current month, i.e. $E[a_{t+1}] = a_t$. As a more structural approach, we model the level of piracy attacks using a Markov switching model based on an underlying (latent) law and order state. This will have an advantage of picking up the persistence of the shift that occurs in the piracy data and captures some of the features of the structural break analysis we perform in section 2.2.5. In addition, the Markov Chain model allows for an intuitive way to integrate the discussed seasonality in attacks to make predictions of piracy. This will be discussed in the following section.

To motivate the switching model, we can return to the theoretical approach above and allow the probability of a successful pirate attack to depend on a latent state, $\ell \in \{S, W\}$ with $\xi(S) < \xi(W)$ where S stands for "strong" and W for "weak". We assume that the probability of successfully hijacking a ship and demanding a ransom is higher when law and order is weak. Using this in the model of piracy above, the mean number of pirate attacks in state ℓ is

$$\mu_\ell \equiv \xi(\ell) M, \; \ell \in \{S, W\}.$$

where $\mu_S < \mu_W$.

Dynamics across law and order states are modelled as a Markov chain governing the process of state transitions. This gives us a filter for emerging data on pirate attacks which can be used to construct a forecast for pirate attacks which can capture the sharp non-linear pattern in the data. We show in Appendix 2.A.2 that this model gives the following formula corresponding to equation (2.1) for the expected number of attacks at $t + 1$:

$$E[a_{t+1}] = \Omega + (\mu_W - \mu_S) \lambda P_t(\ell = W) \tag{2.4}$$

where Ω is a constant, λ is a measure of persistence of the process and $P_t(\ell = W)$ is the probability that the region is in the weak state at time t. The latter is the only time-varying factor in equation (2.4) and evolves according to the history of piracy attacks. By estimating the parameters of the underlying process, we can construct an empirical counterpart to equation (2.1).[27]

This type of model, first proposed in Hamilton (1989), has been popular among time series economists modelling the non-linear properties of business cycle fluctuations. The model's core parameters are estimated using the data on attacks using the Expectation Maximization (EM) Algorithm described in Hamilton (1990) which generates an estimate of the parameters by iteration and is easy to implement.

The abrupt swings in the forecast number of attacks are driven by changes in $P_t(\ell = W)$ between values that are close to zero and one while the impact of the estimated probability on expectations is driven by our estimate of $(\hat{\mu}_W - \hat{\mu}_S)\hat{\lambda}$. It is interesting to observe that the predictions made by our model are that the state shifted in April 2008 which is very much in line with the assessment of the Joint War Committee.

The results when (2.4) is used instead of a_t to estimate (2.3) is in column (1) of Table 2.4. The coefficient on Somali piracy remains significant. Moreover, the

[27] We discuss details of the estimation in appendix 2.A.2. Note that $P(\ell_t = W)$, is a function of the particular history of attacks in month t and the set of Markov chain parameters: two state-specific means, two persistence parameters which together determine λ and two state-specific variances. To forecast piracy attacks, we use the observed number of attacks in month t to calculate the probability $P(\ell_t = W)$ that the region is in a weak state given a set of known parameters. Equation (2.4) shows that if $P(\ell_t = W)$ increases then the expected value of attacks next month increases by $(\hat{\mu}_W - \hat{\mu}_S)\hat{\lambda}$. The estimate for $(\hat{\mu}_W - \hat{\mu}_S)\hat{\lambda}$ is 11.45 attacks.

estimated increase in piracy costs around May 2008 is similar with 9.2 percent which is only slightly higher than the estimate in column (2) of Table 2.3.

Column (3) of Table 2.4 entertains an alternative measure of expectations. We obtain data from Google search intensity for the term "Somalia Piracy". This may capture overall expectations about piracy as well. As the coefficient suggests it does predict shipping costs, but as table 2.9 confirms, it performs a lot worse than any of our other forecast models.

Seasonality The baseline model identifies law and order as the only underlying cause of fluctuations in piracy attacks over time. However, Table 2.1 also shows a pronounced seasonal pattern which can be incorporated into the empirical model. Suppose that there is a month-specific shock to the success probability, ξ_t. Now the average number of pirate attacks will depend on the month and equation (2.1) generalizes to

$$\mu_{m\ell} = \xi(\ell) w_m M \tag{2.5}$$

where w_m is the mean "weather" shock to piracy success in month m. This allows us to rewrite the mean number of attacks as an interaction between an indicator for the weak and strong state, $\ell \in \{S, W\}$, and a monthly mean of attacks during times of weak and strong law and order, α_{mW} and α_{mS}.

$$\mu_{m\ell} = I[\ell = W]\alpha_{mW} + I[\ell = S]\alpha_{mS}.$$

Thus, we have a month-dependent mean in the underlying Markov chain which switches between strong and weak law and order. This model allows us to capture Table 2.1 perfectly.

The forecast number of attacks at $t+1$ when that month is m is now a function of the probability of the weak state in t and the mean of attacks during weak and strong law and order states for $t+1$. Thus (2.4) generalizes to:

$$E[a_{mt+1}] = \Omega_m + (\alpha_{mW} - \alpha_{mS})\lambda P_t(\ell = W) \tag{2.6}$$

where Ω_m is again a constant (now specific to month m).[28]

We show in the Appendix Table 2.9 that this model outperforms all other models in its predictive power significantly. It allows us to predict 80 percent of the variation in attacks.[29] The coefficient in column (2) in Table 4 confirms previous estimates. We find that the rise in piracy in 2008 led to an increase in shipping costs by 8.7 percent.

Column (3) of Table 2.4 entertains an alternative measure of expectations that may capture that expectations are a driven by media coverage, instead of past attacks. We obtain data from Google search intensity for the term "Somalia Piracy" for a reduced

[28]We can directly apply the estimation method described in the appendix to a richer parameterization with 28 parameters.

[29]We also ran a out of sample forecast in which we predicted pirate attacks in the period 2011 to 2013 with the models in table 2.9. Again the seasonal EM algorithm outperforms the AR(2) process. Both perform much worse than within sample, however.

form measure of news reports. As the coefficient suggests it does predict shipping costs very well.[30]

Omitted Economic Trends By including time dummy variables (for each month), we are controlling for general developments in the global shipping market. These may be important over this period given that the global financial crisis erupts in 2008 alongside a growth in the capacity in bulk shipping. For this to create a problem for our analysis would require that the routes that we have classified as being affected by piracy are differentially influenced by changes in market conditions in a way that increases bulk shipping costs. The main trend in this period is, however, a switch in bulk trade in Asia *away* from Europe and towards other Asian countries, in particular Australia and the Americas.[31] This would tend to work against our core findings as we would expect it to put downward pressure on prices for bulk charter agreements between Europe and Asia which pass through the piracy affected area. Nonetheless, we look at two further ways of controlling for changes in route-specific economic factors.

Column (4) of Table 2.4 adds GDP growth for the origin and destination of each route to the specification. Due to the coarseness of the destination data in particular (discussed further in Appendix 2.A) we were forced to aggregate to the level of regional GDP for this exercise. Controlling for either annual regional GDP levels (regression not shown), interpolated monthly regional GDP levels (regression not show) or regional GDP growth, as shown in column (5) does not change the main result.

A further possible concern is that trade patterns might change differentially and systematically across time. This concern is particularly important in the light of a recent World Bank (2013) study which tries to identify the effect of piracy from changes in trade. In order to deal with this concern, we gathered monthly trade data from the IMF Direction of Trade Statistics (DOTS) and matched this trade data to our charter contracts. In column (5) we control for the value of trade on the route of the charter during the same month. Controlling for trade in this way has no impact on our piracy cost estimate. The coefficient on trade is positive but insignificant. This suggests that time and dyad fixed effects do a good job in capturing variation in the conditions in shipping markets.

In column (6) of Table 2.4 we further address concerns about unobservable economic trends by incorporating a separate set of region specific time trends for each of the twenty-four regions from which shipping emanates (Eastern Africa, Southwest Asia, etc.).[32] Even with this rather saturated specification, the core finding regarding the effect of piracy attacks is robust, albeit with a somewhat smaller coefficient compared to column (2) of Table 2.3. Our core finding also holds up if we control for

[30]We thank an anonymous referee for suggesting this approach. In Appendix table 2.9, we evaluate the predictive power of news reports over and above passed attacks.

[31]See the detailed discussion in UNCTAD (2009) and UNCTAD (2010).

[32]This entails problems. The time trend for the Middle East, for example, will effectively capture part of the variation induced by piracy as most charters in this area cross the piracy area.

a separate time fixed effect for the region in which each shipping route starts and finishes.

Shipping rates fell considerably when world trade collapsed in Fall 2008. One way to see whether our results are robust to a break in trade patterns around this time is to have separate sets of dyad fixed effects before and after the Lehman Brothers collapse in September 2008. Column (7) in Table 2.4 shows that we still find a significant positive effect of attacks on shipping rates.

Alternative Measures of Vulnerability to Piracy Attacks In order to match the data on piracy attacks to the shipping contracts data, it is necessary to specify criteria according to which some routes are vulnerable to piracy attack. As there is some leeway in the choice of such criteria, we now present some further results which show that our results are robust to alternative ways of doing this. These are shown in Table 2.5.

Columns (1) and (2) study the robustness of our results to the computation of the maritime routes. Ships could be travelling alternative routes in order to avoid the Suez canal fees or the piracy region and we would expect such re-routing to be more of an issue for maritime routes for which there is a feasible alternative route which does not use the Suez Canal and which is not significantly longer compared to a route using the Suez Canal (and thus passing through the piracy region). To examine this, we used our algorithm to compute alternative routes while adding the constraint that vessels cannot travel through the Suez Canal. We then assign treatment based on these alternative routes if they are at most 10 percent (column (1)) or 20 percent (column (2)) longer than the Suez Canal route. The point estimate for the Somalia area becomes slightly higher but is indistinguishable from our main result in Table 2.3 column (2).

In column (3) we use a more narrowly defined piracy region focusing on the key choke point: the Gulf of Aden.[33] The result shows that piracy in the Gulf of Aden still has a significantly positive impact on shipping prices through that area. Again the magnitude of the effect of piracy is very similar to that reported in our core specification.

Column (4) explores variation in exposure to piracy risk by introducing an interaction between our treatment dummy and the share of a trade route that passes through the Somalia area rectangle. We expect piracy attacks to affect the daily charter rate more if a larger share of the charter goes through the piracy-risk area. Column (4) provides a test of this by including the interaction of the share and attacks in addition to attacks. As expected, the interaction term is positive and significant which implies that higher rates are paid on routes that are treated for longer. Conditional on going through Somalia the average trade route is susceptible to piracy for about 20 percent of its length with the maximum being 68 percent. The coefficient on the

[33]For the Gulf of Aden, the bounding box is given by latitude $\in [10.5, 17]$ and longitude $\in [40, 52.2]$. This is drawn as the dark blue area in Figure 2.2.

interaction implies that a route with maximum exposure would become 20 percent more expensive with the rise in piracy.

Column (5) includes only the interaction term between the share of a route through the piracy area with the number of attacks as a measure of treatment. The rationale here is to drop the correlated attacks variable to get a better idea of the magnitudes involved. The coefficient is highly significant. The effect of an increase in number of attacks after May 2008 on shipping rates for the average treatment share is 7.92 percent, which is close to our other estimates.

Columns (6) and (7) present the results from a similar exercise to that conducted in columns (4) and (5). The key difference is that we now allow there to be time-variation in the maritime area that is considered to be affected by piracy. This addresses a potential concern that the choice of the broad Somalia box as our piracy area is somewhat ad-hoc. We generate time variation in the piracy area by computing the convex hull that is spanned by the coordinates of all the piracy attacks that had occurred up to each year inside the rectangular area which we specified for Somalia above (the shaded area in Figure 2.2).[34] We then compute the share of a shipping route that crosses each of these convex hulls. This gives us time-variation in the share of a shipping route that is affected by piracy. Using this, we can conduct the same exercises as we reported in columns (4) and (5) of Table 2.5. The results obtained are very similar suggesting that our initial way of capturing the risk of piracy is robust.[35]

2.2.5 Composition Effects and Re-routing

We now explore the possibility that, as well as affecting costs, piracy attacks also changed the desirability of shipping on routes affected by piracy.

Effects on Shipping Piracy attacks could be a deterrent to shipping goods through areas that are susceptible to piracy attacks. We need to be able to rule this out, because if piracy was positively correlated with the quantity shipped, the observed higher shipping costs may simply reflect increased demand on this particular route. We consider two dimensions in which piracy could affect shipping other than increasing the cost of shipping. First, piracy could directly affect the amount of traffic on piracy routes. Second, piracy could affect the composition of ships going through the piracy areas.

Data on passages through the Suez canal offers a way to analyze the impact of piracy on trade volumes. We obtain data on the quantities of cargo in deadweight tons through the Suez canal for each month of our sample period. The task of identifying a piracy effect in this time series is complicated by the fact that the failure of Lehman Brothers, an event which signalled the onset of the most serious phase of

[34]We are grateful to a referee for suggesting this exercise. The convex hulls for 2005 and 2010 are plotted in figure 2.7.

[35]This is perhaps not too surprising given that the maritime insurance industry considers almost the entire Indian ocean, similar to our Somalia bounding box, to be a war risk area from early 2009 onwards.

the global financial crisis, occurs in September 2008 - only shortly after the upsurge in piracy. As is well known, this led to a significant reduction in world trade.

To disentangle the effect of the economic crisis from the effect of piracy we look for breaks in the time-series of cargo traffic and try to identify in which month, if any, a break took place. Specifically, we use the method described in Bai (2009) to determine the break points in the series for cargo volumes and for piracy attacks in the Somalia region. For the trade volume exercise, we search for the optimal location and number of break-points according to a BIC criterion using the following model:

$$Cargo_t = \beta_0 + \beta_1 t + \epsilon_t$$

for all possible dates t. We find exactly one break-point for the period following November 2008, roughly two months after Lehman Brothers failed. Bai and Perron (2003) propose a method for obtaining a confidence band around an estimated break-point. Applying their approach, we find that with 99% confidence the break occurs in the period October to December 2008. This makes sense given that goods already in transit and on which shipping contracts had been agreed would not have been affected by the Lehman crash. Applying the same approach to piracy attacks, we find that the break in the series is in July 2008. This is different from the break point in the cargo series. That said, the 99% confidence band for the break in the mean level of piracy is a lot wider and ranges from August 2007 to August 2008, the latter still being before Lehman's failure.

This motivates running regressions in which we include a dummy variable for November 2008 onwards to pick up the effect of Lehman Brother's failure when looking for an effect of piracy attacks on the quantity of cargo being shipped through the Suez canal. Thus we run

$$Cargo_t = \lambda_0 + \lambda_1 a_t + \lambda_2 Lehman_t + \lambda_3 t + \eta_t. \tag{2.7}$$

where $Lehman_t$ is a dummy variable that switches from zero to one in November 2008.

The results from running (2.7) with and without the Lehman dummy are in columns (1) and (2) of Table 2.6. Column (1) shows that if we only include the level of piracy attacks, then we get a large and significant effect of piracy attacks on cargo; the effect amounts to a 30 percent reduction at the mean level of monthly piracy attacks after May 2008. Once we include the structural break identified by the method outlined above, this becomes much smaller in size and insignificant as column (2) shows.

These results suggest that piracy did not have a significant effect on the amount of cargo shipped through the Suez canal. That said, the 95% interval of the estimate in column (2) is consistent with a negative effect on trade of up to 3.5% which is in line with the Feyrer (2009) estimates of the effect of transport costs on trade.[36] Feyrer's

[36]The average traffic pre May 2008 was $43,000$ metric tons. The change in the number of attacks was

estimates suggest that an increase of trade costs by 8% would yield a decrease in trade between 1.6% and 4%. As we cannot identify the effect on trade we therefore use Feyrer's estimates in an extension to our core welfare calculations.

We see these results as being very much in line with a recent World Bank report that uses trade value data to identify the welfare effects of piracy. The report attempts to estimate the effect of piracy on trade from gravity equations but finds only marginally significant and inconclusive results.[37]

Effects on Average Ship-Size One possible reaction to piracy would be to use ships that are less susceptible to piracy attack. We look for evidence of a shift in composition by looking at the average DWT of ships in our data over the period and see if this varies in response to the threat of piracy. Thus, we use our data at the route level to calculate the average weight of a ship on route d at t and run the regression:

$$DWT_{dt} = \alpha_d + \alpha_t + \gamma A_{dt} + \psi_{sdt}$$

where (α_d, α_t) are route and month dummies. The effect of piracy is now identified from variation within a route over time using the same treatment assignment as in our core results above.

The result is reported in column (3) of Table 2.6. While there is a negative coefficient on Somali piracy attacks, this coefficient is not significant at conventional levels. Thus, there does not seem to be any evidence of substitution in ship size in response to piracy.

2.3 The Welfare Cost of Piracy

We now discuss what our results imply for the welfare cost of piracy. Our approach is distinct from existing estimates such as One Earth Future Foundation (2010, 2011) since we have estimated the impact of piracy on shipping costs directly rather than using an accounting approach. We also adopt an explicit welfare criterion which recognizes that piracy creates a transfer from consumers of traded goods (who ultimately bear the cost) to pirates. We compare piracy to the cost of making a more efficient transfer via a tax. However, not all costs are necessarily captured by the impact of piracy on shipping costs and we will consider the sensitivity of the estimates to such concerns.

14.33. This implies a point estimate for the decrease in traffic of $\frac{32.89*14.33}{43000} = 1.1\%$. The upper bound is calculated from the 95% interval $1.1 + \frac{1.96*36.9*14.33}{43000} = 3.5\%$.

[37] See World Bank (2013). This is not mentioned in the body of the report. However, the main result and robustness table show that only four out of eleven estimated coefficients have the right sign and are significant. The only coefficient that is significant at the 5 percent level has the wrong sign.

2.3.1 Framework

Piracy leads to a transfer of resources to pirates via ransoms. Resources are used by pirates in securing these ransoms and by ship owners and governments in resisting them. The costs of the ransoms and damage to ships are also borne directly by those who pay them. These costs are pooled across the industry through insurance. Resources are also used in writing insurance costs and in the lengthy process of negotiations with pirates. As with any transfer program, there is a question of who pays in the end. If the market for shipping is competitive then any increased cost will be passed on to consumers of the final goods in the form of higher prices. And full forward shifting is the benchmark that we consider.

Let Δ denote the cost increase per unit of shipping due to piracy. Part of this cost increase is a transfer to pirates, $\tau(\Delta)$, to which we could attach a distributional "welfare" weight. It is somewhat debatable what this weight should be. Ransoms transfer income to a poor country (Somalia) but they go mainly to organized criminals. It is unclear how far these benefits trickle down to the wider Somali population.[38] We feel that it is best to be agnostic about this and base our welfare approach on Coate (2000). Using his reasoning, we should care principally that any transfer made to pirates is accomplished in the most efficient way and hence the welfare loss are the resources spent in the process of delivering the transfer.

We therefore use the following thought experiment. Imagine there were an efficient transfer scheme, t, to transfer money from final consumers to pirates. If we were to keep pirates indifferent but use the efficient transfer, what would be the difference in welfare costs between this hypothetical transfer and the actual costs caused by piracy?[39]

In order to understand this welfare loss we need to first describe demand for final goods as a function of shipping costs. Suppose that there demand for a composite traded good, X, whose transport is susceptible to piracy attacks. Suppose that shipping demand has a fixed coefficient technology so that demand for shipping is νX. The number νX is best thought of as ton days, i.e. as the number of shipped tons multiplied by the average maritime journey time.[40] Suppose that there is a representative consumer with utility $U(X)$ and additive quasi-linear utility.

Shipping costs influence demand through price adjustments. Denote demand for the final good as $\hat{X}(\psi + \phi)$. Where ψ is the cost of production and ϕ is the shipping costs per unit of the final good. Under piracy the shipping cost is

$$\phi(\Delta) = \nu[c + \Delta]$$

[38] Shortland (2011) provides some evidence that piracy revenue trickles into Somali society and has a positive developmental effect.

[39] Of course, a tax would be costly to administer and we are not including this in our thought experiment. But neither are we including the costs to pirates of extracting the resource. We expect this to induce a downward bias in our estimates of the welfare costs.

[40] This view is very much in line with the usual measure of mile tons. For an interesting discussion regarding this see Stopford (2009). We disregard variable shipping speeds which makes the two measures equivalent.

and under the efficient transfer scheme it is

$$\phi(t) = \nu[c+t]$$

where c is the cost of shipping.

If we were able to replace predation with taxation, the required unit tax, t, would be given by

$$t\nu\hat{X}(\psi+\phi(t)) = \tau(\Delta)\nu\hat{X}(\psi+\phi(\Delta)). \quad (2.8)$$

The left hand side of this equation shows total income from the tax. The right hand side shows revenue from predation. Importantly, $\Delta \geq \tau(\Delta)$, the cost incurred by ship owners is potentially larger than what pirates make.

In this simple model the welfare loss caused by piracy is then given by

$$\begin{aligned} L(\Delta) \;=\; & \left\{U\left(\hat{X}(\psi+\phi(t))\right) - \hat{X}(\psi+\phi(t))\left[\psi+\phi(t)\right]\right\} \quad (2.9) \\ & -\left\{U\left(\hat{X}(\psi+\phi(\Delta))\right) - \hat{X}(\psi+\phi(\Delta))\left[\psi+\phi(\Delta)\right]\right\}. \end{aligned}$$

where demand is potentially smaller under higher shipping costs, $\hat{X}(\psi+\phi(\Delta)) \leq \hat{X}(\psi+\phi(t))$, because the price of the final good increases from $\psi+\phi(t)$ to $\psi+\phi(\Delta)$.

2.3.2 Benchmark Estimate

A benchmark (first-order) estimate of (2.9) can be found by ignoring any trade response (i.e. demand response by consumers). Thus $\hat{X}(\psi+\nu[c+\Delta])$ is completely inelastic and $t = \tau(\Delta)$. In this case equation (2.9) takes on the simple form:

$$L^1(\Delta) = [\Delta - \tau(\Delta)] \times \nu\hat{X}. \quad (2.10)$$

Estimates of equation (2.10) for the year 2010 are in column (1) of Table 2.7. Details of all calculations are in Appendix 2.A.4. In Panel A we use the detailed data available from the Suez Canal authority on the total amount of tons shipped through the Gulf of Aden. We translate this number into an amount of DWT×days by using the mean bulk ship speed (from Stopford (2009)) and the average length of the trip in the respective sample.[41] Panel B adds an estimate of the DWT×days that do not travel through the Gulf of Aden but through the Indian Ocean.

To get a feel for the plausible range, we present a low and a high estimate. Our low estimate uses the coefficient from column (1) in Table 2.3 and our high estimate uses column (4) of Table 2.3. Panel B applies these numbers to trade through the Indian Ocean.

We illustrate our calculations of $L^1(\Delta)$ with the low estimate in panel A of Table 2.7. We use the coefficient in column (1) of Table 2.3 and the average rate charter rate

[41] We make the assumption all of this cargo is comparable to ours in terms of its exposure to higher shipping costs, journey length and travels though the Gulf of Aden.

of 0.4726. This yields the following estimate of total piracy costs:

$$\begin{aligned}\Delta \times \nu \hat{X} &= 0.082 * 0.4726 \times 30.3 * 646064000 \\ &= 758 \text{ million USD}\end{aligned}$$

for 2010.[42] The average ship had a cargo capacity of 47,000 DWT which implies a pirate cost of around $55,000$ USD.

Our estimate of $\tau(\Delta) \times \nu \hat{X}$ is the gross ransoms paid less the costs incurred by pirates in generating this. The main problem with calculating total ransom payments is that not all ransom payments are observed. Depending on the assumptions made on the unobserved payments, total ransom amounts vary widely. The One Earth Future Foundation (2011) and Geopolicity (2011) report ransom amounts of up to 240 USD for 2010. A recent report by World Bank (2013) finds much lower numbers of between 70 million USD and 90 million USD for the year 2010. Another World Bank (2013) report finds that labour and capital costs leave a (political) rent of between 70 and 86 percent of revenues. With these estimates of revenues and rents we get a range of 50 million USD up to 205 million USD for $\tau(\Delta) \times \nu \hat{X}$. For now we ignore the margin of uncertainty and pick a value in the middle of this range, 120 million USD.[43]

Together with our estimate of $\Delta \times \nu \hat{X}$ this sums to the number

$$L^1(\Delta) = [758 - 120] \text{ million USD} = 638 \text{ million USD}.$$

Even from this lower-bound estimate it should become clear that the additional costs incurred due to the threat of piracy vastly exceeds what it would cost to offer pirates a tax-financed transfer of comparable magnitude to the revenues that they earn.

Panel B shows, not surprisingly, that the estimated cost is much higher when we calculate the value of shipping for the wider region including trade routes that do not cross the Gulf of Aden. Our estimates of the welfare cost increase by around 70 percent.

One way to understand the welfare loss is to contrast expected ransoms faced by the shipping industry with the increase in shipping costs. In 2010 there were 18,000 vessels travelling through the Suez Canal. In that year, pirates made 50 successful attacks which generated up to 4 Million USD each. This implies an expected loss of up to $11,000$ USD per vessel compared to an increase in shipping costs of $55,000$ USD. Thus, the realized losses due to ransom payments were about five times lower than our most conservative estimate of the welfare loss per vessel. This a fundamental consequence of economic predation combined with private security investments as

[42]Obviously this number is subject to a large margin of error. For example, container traffic is likely to be less affected. Were we to suppose that there was no effect on container ships then the size of the affected deadweight tonnage would be only $279,063,000$ and the cost would be considerably lower. We abstract from this as the value of container goods is likely to be much larger which would increase the cost.

[43]This is also consistent with the calculations at `http://www.goo.gl/5T9nW`.

we discuss further below.

2.3.3 Extended Estimates

There are further reasons to believe that our estimates in column (1) of Table 2.7 are a lower bound on the true cost. We now consider two of these: (i) the possibility of a demand response which reduces trade and (ii) the possibility that only a fraction of the cost of piracy is paid by the charterer.[44]

Allowing for the possibility of a demand response, we show in the Appendix 2.A.4 that the welfare loss due to a decrease in trade can be approximated by a scaling factor on the estimate above, which depends on the elasticity of trade with respect to transport costs, $\hat{\eta}$, and is given by[45]

$$L^2(\Delta) = L^1(\Delta)\left[1 + \frac{1}{2}\frac{\Delta - \tau(\Delta)}{c + \Delta}\hat{\eta}\right]. \tag{2.11}$$

It is clear that $L^2(\Delta) > L^1(\Delta)$ as long as $\hat{\eta} > 0$.

There are several possible estimates of $\hat{\eta}$ that we could use. Recent estimates from Feyrer (2009), who uses the Suez Canal closure from 1967 to 1975 as a shock to distance, suggest that a value of $\hat{\eta}$ between 0.2 and 0.5 is reasonable. This is a little lower than the estimate found in the meta study by Disdier (2008) which is 0.9. However, given the context of the Feyrer (2009) study, we use an estimate of 0.5 in column (2) of Table 2.7. This implies that $L^2(\Delta)$ is larger than $L^1(\Delta)$ by a factor of between 1.017 and 1.03, i.e. the additional welfare loss due to changes in quantity are relatively marginal (consistent with this being a second-order effect in our context). This is confirmed when comparing the new estimates in column (2) of Table 2.7 with column (1).

Column (3) of Table 2.7 allows for the possibility that the increase in shipping rates fails to capture *all* of the additional costs imposed by piracy.[46] To obtain an upper bound on this we check what would happen if costs were split between the ship owner and charterer according to the "general average rule" as it is known in the shipping industry. This shares the costs of protecting the ship in proportion to the value of the vessel and the cargo. Assume then that a share ζ of the piracy costs are borne by the ship-owner. The charter rate increase Δ is the transfer that compensates the owner for piracy costs over and above what the charterer bears. Then if charter rates increase by Δ due to shipping costs the overall cost to the industry is given by

[44] Similarly, if we believe that the market for ship capacity is not competitive, we could see that piracy related expenses may be forwarded with a markup. This is a possibility we do not explicitly consider further.

[45] Note that we calculate an upper bound this way as charter costs are just a part of total (maritime) transport costs.

[46] For example, time charter rates do not cover fuel expenses. If bulk ships speed up or re-route due to piracy then this will not appear in the charter rate leading to an underestimate.

$\frac{\Delta}{2\zeta-1}$.[47] This yields our third measure of welfare cost of:

$$L^3(\Delta) = \left[\frac{\Delta}{2\zeta - 1} - \tau(\Delta)\right] \times \nu \hat{X} \tag{2.12}$$

which is reported in column (3) of Table 2.7. The details on the calibration of ζ can be found in Appendix 2.A.4. This leads to estimates that are somewhat larger than in column (1) of Table 2.7. For example, the low estimate allowing for general averaging is 130 percent higher. The resulting numbers give us a good idea of how much additional costs could be arising on the cargo owner's side in terms of additional fuel costs, insurance and re-routing.

Putting this together, our estimates for the Gulf of Aden and the Indian Ocean are between 1.1 billion USD and 3.7 billion USD. While the range of estimates is quite large, the comparison between these estimates and those of the transfer received by pirates is telling. We used a figure of 120 million USD for the transfer to pirates and the welfare costs would still be substantial even we used the highest estimate of 240 million USD from One Earth Future Foundation (2011) and Geopolicity (2011). And the welfare cost would be higher still using the smaller numbers on transfers to pirates in World Bank (2013). Hence, the results suggest a substantial welfare cost from piracy.

2.3.4 Predation versus Taxation

We can use the analysis above to calculate t - the tax rate on shipping through Aden that would yield the same revenue that is going to pirates. Of course there is no reason to expect that such a tax and transfer system provides a realistic solution to the piracy problem. Identifying those who should receive the transfer would be impossible to identify. However, it does provide another way of conceptualizing the costs involved.

Disregarding the effect on trade we get this tax rate from the following calculation:

$$\begin{aligned} t &= \frac{\tau(\Delta)\,\nu\hat{X}\left(\psi + \nu\left[c+\Delta\right]\right)}{\nu\hat{X}\left(\psi + \nu\left[c+\Delta\right]\right)} \\ &= \frac{120 \text{ million USD}}{0.4726 * 30.3 * 646{,}064{,}000 + 0.4648 * 20.67 * 578{,}000{,}000} \\ &= 0.008. \end{aligned}$$

This implies that a tax rate of just 0.8 percent on chartering would be needed to generate a transfer of comparable magnitude to that generated by piracy. Even if we assume that a rent of 205 million USD was generated by piracy this would still imply a tax rate of only 1.4 percent. This contrasts with our estimates of the increase in shipping costs of between 8 and 12 percent. The predatory activity of the kind un-

[47] To get an intuition for the formula assume that the shipping cost is 100. The owner has additional costs due to piracy of 20 and the charterer pays 10. The charter rate will go up by 10 due to piracy but overall costs due to piracy is 30. And, indeed, $\frac{1}{2\zeta-1} = \frac{1}{2\frac{2}{3}-1} = 3$ in this case.

dertaken by pirates is between 5 and 16 times more costly as a means of transferring resources to pirates than taxation would be.

Somalia is now the focus of international attention although with limited progress. In the context of potential donor interest, it is instructive to consider how many Somali's could be hired for one year using the additional resources that we estimate are expended by the shipping industry in response to the threat of piracy. Using the numbers in panel B of Table 2.7, a conservative estimate of the costs of piracy to the shipping industry is about 1.05 billion USD. We use wage data from the Somali Food Security and Nutrition Analysis Unit (FSNAU) presented in Shortland (2011) to calculate a yearly wage of about 870 USD.[48] This means that the extra spending due to piracy could finance one year of employment for more than 1.2 million laborers at the going market rate in 2010. This does not mean that such a transfer scheme would be realistic or that it would prevent piracy. But it illustrates the scale of losses to the industry relative to the reality of the Somali economy.

2.3.5 Investing in Security

Given the increases in shipping costs that we have found, the question arises of why piracy remains a threat. The question of how security is provided and the optimality of arrangements in place raises a range of issues which go beyond the scope of the paper. However, we briefly discuss some of the issues here and argue that, *prima facie*, there is evidence that there is currently scope for coordinating security.

The obvious course open to ship owners to reduce piracy risk is to make independent investments in defensive measures such as barbed wire, panic rooms and security crews for their ships. We would expect this to be done according to a cost-benefit calculation by each ship owner. Our estimates could be regarded as the expected hijacking costs if no defensive measures were taken.[49] In addition to ransom payments, attempted hijacks generate costs if ships are damaged after the hijacking, especially since pirates have to hold ships for long enough to establish their credibility. The risk of being captured for several months also increases the cost of recruiting crew members who demand a wage premium as compensation. Ransom negotiations for crew and ship are like an inefficient war of attrition which increases the cost of doing business and creates delay over and above the cost of the ransom.[50]

Given that no ship with security teams on board has been hijacked we consider now the costs arising when investments in security are decentralized.[51] From conversations with security firms, we know that they charge about 3000 USD for a security

[48] In 2010 the highest daily wage paid in Somalia was about 100,000 Somali Shillings (SSh). Assuming 261 work days and an exchange rate of about 30,000 SSh/USD this implies a yearly wage of about 870 USD.

[49] In the interpretation of Bowles and Jayadev (2004) the welfare loss we capture is then a direct consequence of the guard labor needed to defend economic inequality.

[50] For an analysis of a closely related ransom bargaining process see Ambrus et al. (2011) who analyze ransom negotiations during a period of piracy in the Mediterranean sea from 1575-1739.

[51] We thank Daron Acemoglu for suggesting that we look at this and Marit Rehavi for suggestions on data.

crew of four per day.[52] The guards typically board the vessel on key points before entering the Indian Ocean. The boarding points are Sri Lanka, the Strait of Hormuz, Madagascar and an anchored vessel in the Red Sea off Djibouti. We compute the average time it takes for a vessel to travel between the boarding points in Sri Lanka, the Strait of Hormuz, Madagascar and the Red Sea. Based on this we compute the total cost of hiring security crews for traffic going through the Suez Canal. We arrive at an estimate of 302 million USD and 486 million USD for 2010. In Table 2.7 we calculated costs which lie between 640 million USD and 2.4 billion USD for traffic through Aden. Taken at face value, it suggests a large loss due to uncoordinated provision of security.

This inefficiency is best explained as due to externalities in protection decisions. There are two plausible externalities at work which have opposite signs. To the extent that being protected increases the chances that unprotected vessels are susceptible to attacks, there is a negative externality from investing in protection.[53] Alone, this might lead to excessive investments. However, investments in protection may reduce the overall level of piracy attacks by reducing pirate intensity. And this will tend to lead to too little investment.[54] Either way, this creates a role for coordinated action.[55] However, such coordination among a myriad of ship-owners will be difficult to achieve and our evidence suggests that externalities leading to free-rider problems may be important.

This apparent inefficiency due to uncoordinated protection notwithstanding, it does appear as if defensive measures in Somalia have been increasingly successful - the number of successful hijacks has declined by more than 70 percent between 2010 and 2012. But the fact that this may be due to higher investments in security implies that the costs of piracy to the industry may not have declined commensurately. These costs of protection continue to be incurred even when there are very few successful

[52] This cost is well in the interval of cost estimates for US security contracts in Iraq. The 2010 United States Government Accountability Office report "Warfighter Support: A Cost Comparison of Using State Department Employees versus Contractors for Security Services in Iraq", for example, gives a range of these costs between 430 USD and 7600 USD for four persons per day.

[53] Bandiera (2003) makes a similar point. In addition, there is anecdotal evidence of an arms race in which pirates are better and better equipped and ship owners move from minor ship modifications to hiring security crews. For a general discussion of these issues see Meza and Gould (1992).

[54] To illustrate this, consider the following simple model. Suppose that there is continuum of ship owners of size one indexed by $n \in [0,1]$ and that each can eliminate the threat of piracy at cost f. Let the threat of any particular ship being attacked when $\hat{n}$ vessels are protected be $p(\hat{n})$. The loss from being attacked is $-\ell$. Now the equilibrium condition determining the fraction of vessels who choose to protect is

$$p(\hat{n})\,\ell = f.$$

For this to be a stable interior solution, we require that $p'(\hat{n}) < 0$. This says that (locally) having more vessels protected, reduces the likelihood of being attacked. The surplus maximizing level of protection maximizes

$$S(\hat{n}) = -\hat{n}f - (1-\hat{n})\,p(\hat{n})\,\ell$$

by choice of $\hat{n}$. And it is straightforward to see that at any stable interior decentralized outcome

$$S'(\hat{n}) = -(1-\hat{n})\,p'(\hat{n})\,\ell > 0$$

i.e., there is too little investment in protection.

[55] See Besley and Ghatak (2010) for development of this argument in relation to property rights enforcement. See Grossman (2002) for a theoretical argument relating to predation.

attacks. Indeed, the possibility of predation can impose welfare costs even where the revenue from predation goes to zero.

There may also be a case for going beyond coordinated private security towards collective provision. Whether this is optimal depends on the technology, coordination among providers and possible scale economies from having vessels that specialize in protection (such as a Navy) as a means of protecting ships. Successful collective provision is likely to occur only if there is an agreed way to share costs. In this regard, it is worthwhile noting that currently member countries of the EU, the US, China, Russia, India, Saudia Arabia, Iran and Japan deploy maritime forces in the area. They patrol an area of sea approximately equal to the size of western Europe.[56] Difficulties in agreeing ways of sharing costs is not a new issue as revealed, for example, in the correspondent report on Chinese piracy in *The London and China Telegraph* from 4th February 1867 noted that

> "Besides we are not the only Power with large interests at stake. French, Americans, and Germans carry on an extensive trade [...] Why should we then incur singly the expense of suppressing piracy if each provided a couple of gunboats the force would suffice for the safety foreign shipping which is all that devolves upon [..] why should the English tax payer alone bear the expense?"

The current reliance of the international community on Naval patrols to combat piracy could succeed in reducing pirate activity further. In the end, the most promising long-term solution would seem to be to restore a functional Somali state which can deny pirates safe haven, thereby dealing with the problem at source.

2.4 Concluding Comments

Piracy is an important source of predation which creates economic disruption. In this paper, we have used estimates of its effect on shipping prices to estimate the welfare cost of Somali piracy.

While what we have studied here is only one specific kind of lawlessness, estimates of the costs of predatory activity in any specific context are rare. We have shown that the cost of piracy is large relative to the size of the transfer to pirates.

The analysis further underlines the difference between organized extraction by the state in the form of taxation and disorganized predation. We estimate that the latter is at least five times more costly. In the language of Olson (1993), pirates are roving bandits while the state is a stationary bandit and hence is in a better place to organize extraction at lower costs. But this requires some commitment power on behalf of a stationary bandit. Absent such commitment in the context of piracy, the shipping industry has started to invest in protection. The resulting reduction in pirate activity, however, is just a change in the way that the costs of piracy manifest

[56]We discuss the additional costs that this might cause in the Appendix 2.A.5.

themselves. Without a return to strong law and order in Somalia, it seems unlikely that the underlying welfare costs will disappear any time soon.

There are a number of insights from our findings which extend beyond the specific context that we study. The results suggest that there can be a substantial cost of predation even if the transfers that are generated are modest. There is a parallel here with the welfare cost of crime more generally. For example, a high perceived risk of burglary can encourage house holders to invest in private security which can lead to significant costs even if, as a consequence, burglars earn low returns from their activity. Thus, gauging the costs of predation and theft requires looking beyond the extent of the crime that actually occurs.

Table 2.1: Seasonality in Attacks in Somalia Region

month	before May 2008	after May 2008
January	1.5	12.5
Februar	2.7	7.5
March	2.9	31.5
April	5.2	34.1
May	3.7	21.0
June	1.8	10.5
July	3.5	4.6
August	2.1	9.4
September	1.3	14.6
October	3.8	18.7
November	2.2	28.0
December	2.4	12.6
average	2.8	17.1
difference (after-before)		14.3

Note: Table shows the mean of attacks in the Somalia area in the periods 2002-2007 and 2008-2009.

Table 2.2: Piracy Attacks and Shipping Costs - Simple Difference in Difference

charter rate per DWT on routes	(1) before May 2008	(2) after May 2008	(3) difference
that do not pass the piracy area	0.486	0.386	0.100
	(0.00306)	(0.00329)	(0.00450)
that pass the piracy area	0.480	0.454	0.026
	(0.00415)	(0.00653)	(0.00781)
difference	0.006	-0.068	**0.074**
	(0.00516)	(0.00731)	(0.00894)

Notes: Charter rates are given in US dollars per deadweight tonnage (DWT).

Table 2.3: Main Results

	(1) Baseline	(2) Adding controls	(3) Post May 08	(4) War Risk	(5) Wind-speed
attacks (Somalia)	8.204**	8.438**	3.862***		
	(3.558)	(3.542)	(1.057)		
war risk area				12.332*	2.183
				(6.924)	(5.556)
calm winds * war risk area					18.254***
					(4.407)
calm winds					-0.036*
					(0.021)
ballast bonus per DWT		-0.001	-1.509***	-0.001	-0.001
		(0.001)	(0.529)	(0.001)	(0.001)
ship age		-0.614***	-0.754***	-0.613***	-0.611***
		(0.098)	(0.191)	(0.098)	(0.098)
handysize	0.622***	0.637***	0.627***	0.639***	0.638***
	(0.031)	(0.031)	(0.030)	(0.030)	(0.030)
handymax	0.400***	0.403***	0.372***	0.404***	0.404***
	(0.031)	(0.036)	(0.031)	(0.036)	(0.036)
panamax	0.149***	0.150***	0.177***	0.152***	0.152***
	(0.018)	(0.012)	(0.014)	(0.012)	(0.012)
capesize	-0.039	-0.051*	-0.082	-0.050*	-0.050*
	(0.037)	(0.029)	(0.067)	(0.030)	(0.030)
Observations	23679	23679	9530	23679	23679
R-squared	.851	.856	.829	.856	.856

Notes: All regressions include dyad fixed effects, ship-size controls and time fixed effects. Robust standard errors adjusted for two-way clustering on the origin and destination country for each voyage are in the parentheses with stars indicating *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The dependent variable is the log of the charter rate in US dollars per dead-weight tonnage (DWT). All attack variables are interactions between a dummy that indicates whether a ship will cross a pirate territory and the number of attacks in that territory. "Handysize" is a dummy that indicates ships with $DWT \leq 35000$. "Handymax" are ships with $35000 < DWT \leq 55000$. "Panamax" are ships with $55000 < DWT \leq 80000$. "Small capesize" are ships with $80000 < DWT \leq 150000$ (omitted). "Capesize" are ships with $DWT > 150000$. “Ballast bonus” is a payment that compensates the ship owner for travelling without cargo on return. “War risk area” is a dummy that indicates whether the area was defined as a war risk area by the Maritime Insurer’s Joint War Committee. The coefficient on the attack variables and of the war-risk and wind-interactions are multiplied by 100 for clearer exposition.

Table 2.4: Robustness Checks - Modelling Expectations and Macroeconomic Controls

	Expectations		Macroeconomic Controls			
	(1) EM	(2) EM (Seasonality)	(3) GDP	(4) Trade	(5) Trends	(6) Lehman Break
forecasted attacks (EM)	9.225*	8.747**				
	(5.539)	(3.709)				
attacks			8.449**	8.385**	5.833**	2.876**
			(3.545)	(3.268)	(2.490)	(1.330)
annual GDP growth origin region			0.024			
			(0.022)			
annual GDP growth destination region			-0.005			
			(0.015)			
monthly trade on dyad				0.008		
				(0.005)		
region time trend	No	No	No	No	Yes	No
post Lehman dyad fixed effect	No	No	No	No	No	Yes
Observations	23679	23679	23679	21469	23679	23282
R-squared	.855	.856	.856	.85	.86	.845

Notes: All regressions include dyad fixed effects, ship-size controls and time fixed effects. Robust standard errors adjusted for two-way clustering on the origin and destination country for each voyage are in the parentheses with stars indicating *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The dependent variable is the log of the daily charter rate per deadweight tonnage (DWT). All attack variables are interactions between a dummy that indicates whether a ship will cross a pirate territory and the number of attacks in that territory. Columns (1) and (2) use alternative models to forecast piracy attacks. We use a simple Markov chain model for column (1) and a Markov chain model that accounts for seasonality in column (2). Ship controls are dummy variables classifying the ship size in terms of DWT and contain the age of the ship and the size of the "Ballast bonus" for a particular voyage. Monthly trade on dyad is the log of the value of monthly trade on a dyad as obtained from the Export time-series from the IMF direction of trade database. For dyads where trade is zero, this is coded as a zero. The coefficient on the attack variables is multiplied by 100 for clearer exposition.

Table 2.5: Robustness Checks - Assignment of Routes, Treatment Areas and Intensity

	Assignment of Routes		Treament Areas				
	(1) Rerouting 10%	(2) Rerouting 20%	(3) Only Aden	(4) Fixed Area	(5) Fixed Area	(6) Varying Area	(7) Varying Area
attacks (Somalia)	9.841***	8.529***		1.911		4.441	
	(3.135)	(3.077)		(4.479)		(3.222)	
attacks (Aden)			7.780*				
			(4.375)				
share of route through piracy area * attacks				29.012*	36.142***	20.586*	40.363***
				(16.157)	(13.557)	(11.272)	(10.797)
share of route through piracy area						0.495	0.412
						(0.849)	(0.816)
Observations	23679	23679	23679	23679	23679	23679	23679
R-squared	.855	.855	.855	.856	.856	.856	.856

Notes: All regressions include dyad fixed effects, ship-size controls and time fixed effects. Robust standard errors adjusted for two-way clustering on the origin and destination country for each voyage are in the parentheses with stars indicating *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The dependent variable is the log of the daily charter rate per deadweight tonnage (DWT). All attack variables are interactions between a dummy that indicates whether a ship will cross a pirate territory and the number of attacks in that territory. Columns (6) and (7) present the results when we use time-varying piracy areas obtained from the year-on-year convex hulls that are spanned by the geo coordinates of the attacks that have occurred up to each year. Columns (4) and (6) test whether the share of a ships journey through maritime areas affected by piracy has an independent effect, while columns (5) and (7) present results when using the continuous measure of the share of route through the piracy area interacted with attacks as a control. The coefficients on the attack variables and war risk dummies are multiplied by 100 for clearer exposition.

Table 2.6: Extended Results - Suez Canal Traffic

	Suez Canal Traffic		Ship Size
	(1)	(2)	(3)
	Basic	Lehman Break	from Chartering Contracts
attacks (Somalia)	-12.015***	-2.296	
	(2.962)	(1.553)	
cargobreak		-0.432***	
		(0.025)	
attacks (Somalia)			-293.468
			(812.198)
linear trend	Yes	Yes	.
route fixed effect	.	.	Yes
month fixed effect	.	.	Yes
Observations	108	108	12753
R-squared	.638	.923	.482

Notes: Robust standard errors are reported. The stars indicate *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The dependent variable is the log of cargo traffic in a particular month through the Suez Canal. The variable "cargobreak" is an indicator that is equal to 1 after the break in cargo trade volumes following the Lehman brothers collapse in November 2008, "attacks" measures the number of attacks in the Somalia area in a given month. The coefficient on the attack variable is multiplied by 100 for clearer exposition.

Table 2.7: The Welfare Cost of Piracy in 2010

Panel A: Aden	(1) L1 (in million USD)	(2) L2 (in million USD)	(3) L3 (in million USD)
low estimate	638	649	1495
high estimate	1017	1045	2302
Panel B: Aden and Indian Ocean	(1) L1 (in million USD)	(2) L2 (in million USD)	(3) L3 (in million USD)
low estimate	1093	1113	2464
high estimate	1700	1749	3757

Notes: Calculations are discussed in section 4 and the appendix F. Column (2) adjusts the welfare loss by taking into account the change in trade. Column (3) adjusts the cost to take into account the share of costs borne by charterers. Panel B uses data on trade to and from the Middle East to calculate the costs for the area including the Indian Ocean.

Table 2.8: Summary Statistics of Main Variables

Variable	Mean	Std. Dev.	Min	Max
trade value (in Mio USD)	3831.42	8101.499	0	42034.211
log(trade value+1)	18.767	5.739	0	24.462
shipage (in years)	9.45	7.31	0	39
deadweight tonnage (dwt)	80092.19	39495.48	5169	300000
rate per day per dwt (in USD)	0.45	0.30	0.01	4.04
ballast bonus per dwt (in USD)	1.03	70.26	0	1.10E+04
distance (in km)	8014	6846	0	2.41E+04
number of attacks in Somalia	7.03	9.06	0	42
number of attacks in Gulf of Aden	4.116	5.493	0	22
average predicted wind speed in m/s (Somalia)	6.34	1.38	4.36	8.81
forecast number of attacks Somalia (Markov Chain)	7.73	5.22	2.79	14.20
forecast number of attacks Somalia (AR(2))	14.59	13.75	1.98	57.26

Table 2.9: Predictive Power of Expectation Models

	Different Expectation Models				Google Searches			
	(1) Lagged Attacks	(2) AR(2)	(3) EM	(4) EM (seasonal)	(5) Somali Piracy	(6) Somali Piracy	(7) Gulf of Aden	(8) Gulf of Aden
forecasted attacks	0.717***	0.710***	0.719***	0.909***		0.621***		0.652***
	(0.099)	(0.097)	(0.107)	(0.072)		(0.166)		(0.148)
Searches ”Somali Piracy”					0.166***	0.032		
					(0.032)	(0.043)		
Searches ”Gulf of Aden”							0.046**	0.008
							(0.018)	(0.015)
Observations	98	98	98	98	83	83	83	83
R-squared	.51	.511	.426	.807	.303	.489	.202	.488

Notes: Results from a regression of piracy in a month on various models of expectations. Robust standard errors reported in parentheses, with stars indicating *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Note that the Google Search data is only available from 2004 onwards.

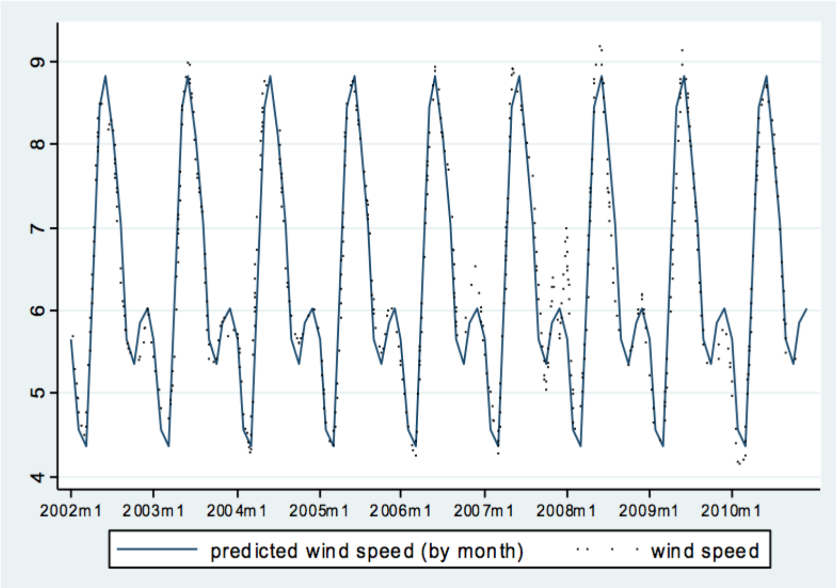

Figure 2.5: Wind Speed in the Somalia Area

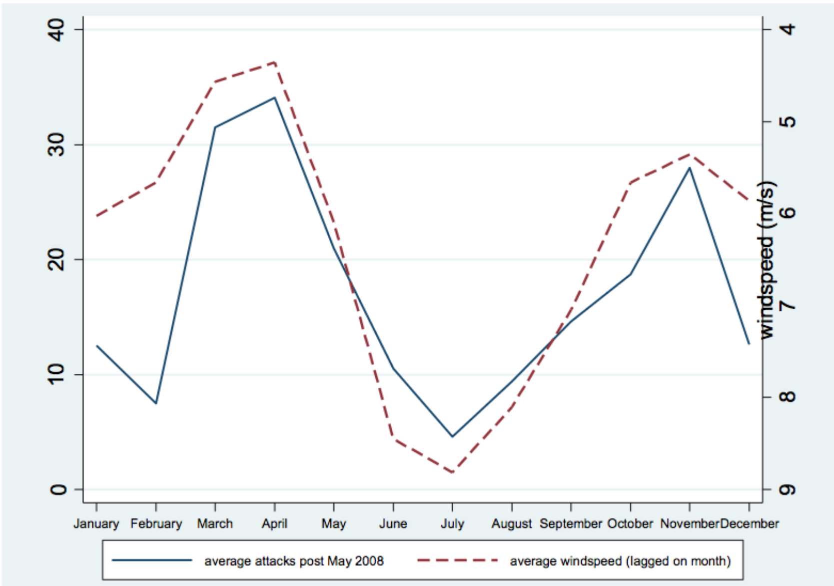

Figure 2.6: Wind Speed and Attacks in the Somalia Area

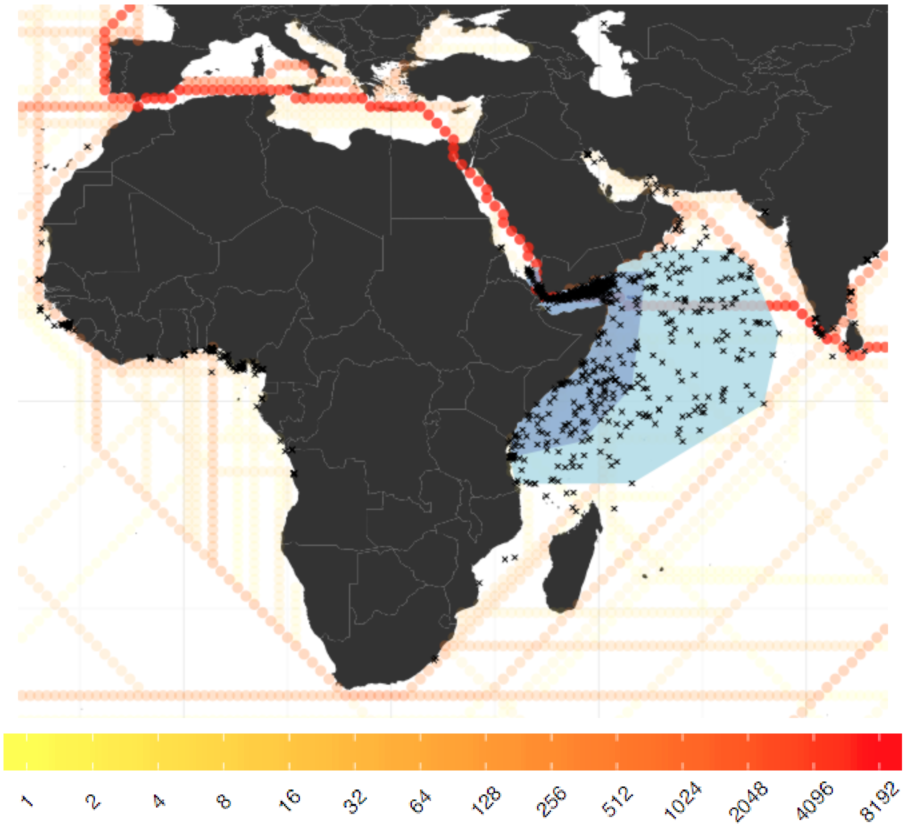

Figure 2.7: Calculated Shipping Lanes and Treatment Areas Based on Convex Hulls
Note: The dark shaded area on the interior is the convex hull spanned by all attacks up to 2005, while the bigger light shaded area is the convex hull spanned by all attacks up to 2010. The location of attacks is indicated by a cross. The circles indicate the shipping lanes, the colouring of which is proportional to the number of observation on each shipping lane according to the continuous colour scheme.

2.A Appendix

This Appendix discusses the data sources and generation of variables.

2.A.1 Data

Chartering Contracts

The data on shipping prices comes from the web-site of N. Cotzias Shipping Consultants which provides monthly reports of the time charter market for the period November 2002 until December 2010.[57] The data is comprised of 33,529 individual fixtures in the dry bulk cargo segment of the market.

It contains details on the vessel that was chartered, the chartering company, the month in which the charter was fixed and the approximate date (day-range / months), when the charter would commence. The details on the vessel give us the current ship name, the year it was built and its deadweight tonnage. The pricing information contains the daily rate in USD, along with a ballast bonus. From these

[57] In early 2011, Cotzias merged with Intermodial (www.intermodal.gr). As of 25th January 2012, the Cotzias data was available on `http://www.goo.gl/g5d0c`.

we construct the daily rate per deadweight ton and the ballast bonus per deadweight ton. On average, about 9% of the charters in our sample include a ballast bonus.

The chartering information provides details about the location of the vessel origin and the vessel destination, i.e. where it will be handed back to the ship owner. Due to the nature of the chartering market, market participants have an active interest in reporting the vessels delivery- and redelivery locations. However, this information comes with varying levels of detail. In particular the redelivery location may either be a port, a country, a maritime region or it may be missing. Further challenges include that sometimes, the port name is spelled wrongly or abbreviations were used. We harmonize the data to country-level pairs. The raw data contains 2,430 distinct delivery- or redelivery locations. We proceeded in two steps:

1. Try an exact match based on a database of port names.[58]. This will give us, in case of an exact match, a port and the country in which this port is located. In case no exact match was found, we used the Google Search Engine to get a spelling suggestion (in case there was a misspelling in the raw data) and try it again with the corrected spelling. Through this, we are able to filter 570 locations, which account for roughly 2/3 of the observations.

2. For the remainder of the delivery- and redelivery locations, we proceed by performing Google searches in a semi-automated way, double checking and validating the results manually.

IMB Piracy Data

The IMB runs the piracy reporting centre which can be contacted 24 hours by vessels under attack. The information received from the ship Masters is immediately relayed to the local law enforcement agencies requesting assistance. In addition, the information received from the ship Masters is broadcast to all vessels in the Ocean region - thus highlighting the threat to a Master en route into the area of risk. The IMB annual reports reproduce the piracy reports received by the piracy reporting centre. They define a piracy attack as

> An act of boarding or attempting to board any ship with the apparent intent to commit theft or any other crime with the apparent intent or capability to use force in the furtherance of that act. (IMB, 2009)

Under this definition, pirate attacks include all actual or attempted attacks on vessels while in port, anchored, berthed or underway. While there is some acknowledged under-reporting, it is the most complete database on maritime piracy that is available. We obtained the annual reports of piracy and robbery incidents from 1999-2010. Each report provides a detailed listing of the piracy incidence, containing the following information:

[58]This database contains the details and locations of 27,625 ports all over the world. They include all major ports, but also smaller ports and docks. It can be accessed on `http://www.goo.gl/s59UE`

- Date (usually to day)
- Name of Ship
- Flag of Ship (sometimes)
- Call sign of ship (not always)
- IMO number of ship (not always)
- Information on location of attack, various levels of detail but mostly a geo-code.
- A narrative of the attack

In total, data on 5,456 incidents is reported. We were not able to use all observations, as quite often for attacks that take place near some ports or just off some islands, the report does not include a geo-coded location. We tried to make use of as many observations as possible by manually geo-coding the missing observations. Furthermore, in early years the data does not give information on whether the vessel was underway or at anchor when it was attacked. This data was manually extracted by analyzing the narrative of the attacks.

Using the maritime areas that we describe in the text, we arrive at a monthly number of piracy attacks in that particular maritime area. This time series is then used throughout the paper.

Wind and Seasonality of Attacks

The connection between wind speed and pirate risk is well-documented. For example, the Office of Naval Intelligence (ONI), a U.S. navy think tank, publishes the Piracy Analysis and Warning Weekly (PAWW) which uses weather data to predict piracy risks in the Somalia area.

We obtained wind data from the National Oceanic and Atmospheric Administration (NOAA), which, among others, provides detailed satellite and observational weather data for the world's oceans. For our purposes we accessed the NOAA Multiple-Satellite Blended Sea Winds database.[59] This particular database has the advantage that it is compiled from several satellites, which limits the number of coverage gaps. Another advantage is, that it provides the data on a fine spatial grid of 0.5° and is available, without gaps from 1987 onwards.

From this database we extracted the monthly mean wind speed pertaining to the geographical grid of our piracy regions. For each month, we have around 8,800 observations of the monthly mean wind speed per 0.5° cell corresponding to our grid. We use this to compute the average wind speed in any month for both the Somalia area.

[59]The data can be accessed via `http://www.goo.gl/DM8O1`.

Figure 2.5 shows the average monthly wind speed for the Somalia area (dotted line) and the predicted wind speed (solid line). The predicted wind speed is calculated from a regression of wind speed on month dummies

$$E\left[wind_t\right] =_{m=1}^{12} month_m\left(t\right) + \epsilon_t.$$

This regression has an R^2 of 0.997. The strong seasonal pattern is also apparent in figure A1 which clearly shows the summer monsoon seasons with increased wind speeds and January and February with very calm winds.

Figure 2.9 shows the connection of the average wind speed prediction (lagged) and mean piracy attacks from Table 2.1. It shows that attacks and lagged wind speed are highly correlated. This is in line with UNOSAT (2010) where the lag reflects the latency period for the pirate militias to redeploy their vessels from the main militia bases along the Puntland coast.

Algorithm for Maritime Routes and Distances

We first determine start and end points for each journey. We use country start and end points rather than specific ports. This is because there is some ambiguity in the port information. This is more severe for some countries. For example, the United States has access to more than one Ocean so that errors could be quite large.

Each country information is interpreted as a specific position. We assigned the most frequently occurring port as our start and finish point for each country. We are then able to automate the way treatment is assigned by computing maritime routes between these points.

The algorithm proceeds as follows.

First, we transform a world map into a coarse 1° grid of the world. The coarseness of the grid allows us to compute optimal routes for the 1,600 routes in a reasonable amount of time on a standard desktop computer. The grid is thus a 360×180 matrix, which we can think of as a graph. Each cell in the matrix represents a node of the graph. We assume that vessels can travel into any of the 8 neighboring cells. The transformation into a grid takes into account that moving along a diagonal corresponds to a larger distance (i.e. higher costs) than moving along straight line vertices.

Second, we then assigned to each cell a cost of crossing using the map on which the grid was defined. We normalize this cost of crossing to be 1 for sea- or oceans and passing a very large number for landmass. We had to manually close the North-West passage and, due to the coarseness of the grid, we had to open up the Suez canal, the Malacca Straits and the Panama canal.

Third, the start- and end-locations, given as GPS coordinates, are then mapped into a particular cell in this graph. We can use simple shortest-path algorithms to compute an optimal path from any two points on the grid. The shortest-path implementation we used is a Dijkstra algorithm implemented in the *R* package Gdistance.[60]

[60]The R package is available from `http://www.goo.gl/BCj6G`. The procedure and the code used is

The algorithm delivers three outputs: a shortest path as a sequence of GPS coordinates, its distance and a cost measure. We use the actual path for the intention to treat assignment that we describe in the text.

2.A.2 Predicting Pirate Attacks

This Appendix discusses table 2.9 which reports the predictive power of five different ways to model equation (2.1). Our baseline specification, in effect, supposes that the best estimate of piracy en route is the level of piracy attacks in the current month, i.e. $E[a_{t+1}] = a_t$. The result is reported in column (1) of Table 2.9.

As an alternative, we also fitted an AR(2) process to the pattern of attacks in the piracy region which we report in column (2) of Table 2.9. The R squared of this model is only marginally higher than in the baseline model.

In section 2.2.4 we also discuss two Markov Chain models which have a more intuitive appeal in the context of the distinct shift in pirate activity after May 2008. The first model uses a Markov Chain to model just the shift from one mean number of attacks to another. We report the fit to the actual attacks in column (3) of Table 2.9. This model performs slightly worse than the baseline.

The second model distinguishes twelve different means, one for each month, in each regime. As can be seen in column (4) this, season specific, Markov Chain model produces an extremely good fit to the realized number of attacks.

Finally, we gathered data on google searches on "Somali piracy", a proxy for news stories, which we use to predict attacks. Results are presented in column (5) of Table 2.9. This variable performs worse than any of the models using the attacks data which suggests that news stories lag attacks instead of leading them. Column (6) shows that news do not add additional predictive power beyond attacks. In columns (7) and (8) we run the same analysis for the search term "Gulf of Aden".

Markov Chain Forecasts

Basics Assume that attacks in region r at time t are given by the following "switching" model:

$$a_t = \mu_\ell (1 - \delta(\ell_t)) + \mu_W \delta(\ell_t) + \varepsilon_t \text{ with } \varepsilon_t \sim N(0, \sigma^2_{\ell_t}) \tag{2.13}$$

where $\delta(S) = 0$ and $\delta(W) = 1$. Thus, μ_S is the mean number of attacks in the inactive state and μ_W is the number of attacks when pirates are active. This allows for the possibility that $\mu_S > 0$. The transition matrix between states is given by:

$$\begin{array}{ccc} & \ell_{t-1} = W & \ell_{t-1} = S \\ \ell_t = W & p & 1-q \\ \ell_t = S & 1-p & q \end{array}$$

available from `http://www.goo.gl/irRxgv`.

at date t, follows the process:

$$\ell_t = 1 - q + \lambda \ell_{t-1} + v_t \text{ where } \lambda = q + p - 1$$

where v_t is an error term with a state-contingent distribution of

$$v_t \mid (\ell_{t-1} = W) = \begin{cases} 1-p & \text{with probability} \quad p \\ -p & \text{with probability} \quad 1-p \end{cases}$$

and

$$v_t \mid (\ell_{t-1} = S) = \begin{cases} -(1-q) & \text{with probability} \quad q \\ q & \text{with probability} \quad 1-q. \end{cases}$$

The model has a vector of six region-specific parameters

$$\theta \equiv \left\{ \mu_W, \mu_S, \sigma_W^2, \sigma_S^2, p, q \right\}$$

which is a complete description of the parameters governing the process of piracy. Most of our use of the model will turn around just four parameters from this vector: μ_W, μ_S, p and q

The history of attacks is used to estimate the probability $P(\ell_t = W \mid H_t, \theta)$ given the attack history H_t and the parameter vector θ. (Details are provided below.) This probability can then be used to form expectations about the level of future attacks, i.e. a_{t+1}. It is easy to show that given equation (2.13) the estimate of attacks in the next month is

$$\begin{aligned} E(a_{t+1} : H_t) &= \mu_W (1-q) + \mu_S q \\ &\quad + (\mu_W - \mu_S) \lambda P(s_t = W \mid H_t, \theta) \end{aligned} \tag{2.14}$$

where $\lambda \equiv p + q - 1$. The first two terms in equation (2.14) are time-invariant functions of the regional parameters θ. One can interpret them as the expected level of attacks in times of inactivity, i.e. at $P(s_t = W \mid H_t, \theta) = 0$. The second term shows that the expected violence in the next period only depends on the estimated probability of conflict in t, the differences in attacks between active and inactive months and the persistence, λ.

Estimation A good starting point for the calculation of the probability of being in conflict, $P(\ell_t = W \mid H_t, \theta)$, is Bayesian updating in period t. In period t, the extrapolation of last period $P(\ell_t = W \mid H_{t-1}, \theta)$ is updated with attacks in t according to the standard formula:

$$P(\ell_t = W \mid H_t, \theta) = \frac{f(a_t \mid \ell_t = W, H_{t-1}, \theta) P(\ell_t = W \mid H_{t-1}, \theta)}{\sum_{j=S}^{W} f(a_t \mid \ell_t = j, H_{t-1}, \theta) P(\ell_t = W \mid H_{t-1}, \theta)}.$$

The immediate insight from this formula is that the probability can only be calculated with an estimate of θ_r because the conditional densities are given by

$$f\left(a_t \mid \ell_t = j, H_{t-1}, \theta\right) = \frac{1}{\sqrt{2\pi\sigma_j^2}} \exp\left(-\frac{\left(a_t - \mu_j\right)^2}{2\sigma_j^2}\right)$$

and therefore depend on parameters in θ.

The probability $P\left(\ell_t = W \mid H_t, \theta\right)$ can be calculated if the past estimate $P\left(\ell_{t-1} = W \mid H_{t-1}, \theta\right)$ is known. To see that this dependency of $P\left(\ell_t = W \mid H_t, \theta\right)$ on $P\left(\ell_{t-1} = W \mid H_{t-1}, \theta\right)$ note that

$$P\left(\ell_t = W \mid H_t, \theta\right) = \sum_{j=0}^{1} P\left(\ell_t = W, \ell_{t-1} = j \mid H_{t-1}, \theta\right).$$

and

$$P\left(\ell_t = W, \ell_{t-1} = j \mid H_{t-1}, \theta\right) = P\left(\ell_t = 1 \mid \ell_{t-1} = j\right) P\left(\ell_{t-1} = W \mid H_{t-1}, \theta\right)$$

where $P\left(\ell_t = W \mid \ell_{t-1} = j\right)$ is nothing else than the estimated p and $1-q$ contained in θ. Hence, one needs $P\left(\ell_{t-1} = W \mid H_{t-1}, \theta\right)$ to calculate $P\left(\ell_t = W \mid H_t, \theta\right)$.

This reliance of $P\left(\ell_t = W \mid H_t, \theta\right)$ on $P\left(\ell_{t-1} = W \mid H_{t-1}, \theta\right)$ implies that previous probabilities of conflict have to be calculated first. The filter therefore takes a starting value $P\left(\ell_0 = 1 \mid H_0, \theta\right)$ and calculates

$$P\left(\ell_1 = 1 \mid H_1, \theta\right), P\left(\ell_2 = 1 \mid H_2, \theta\right) ... P\left(\ell_T = 1 \mid H_T, \theta\right)$$

by iteratively updating the probability of conflict with the monthly attacks data a_t. To some degree this is what the charter parties of a shipment through the Somalia area would have done, too.

However, this simple filter relies on the availability of the vector θ. The problem is that θ cannot be calculated without knowing the states $\ell_1, \ell_2 ... \ell_T$ which are unobserved. Hence, the estimation method needs to determine when regime shifts occurred and at the same time estimate the parameters of the model. One way of estimating the parameters of the violence process is the Expectation Maximization (EM) Algorithm described in Hamilton (1990) which generates an estimate of θ by iteration.

In each iteration the algorithm makes use of the "smoothed" probability of conflict which is based on the entire attack time series data

$$P\left(\ell_t = 1 \mid a_T, a_{T-1}, ..., a_1, \theta\right).$$

Nothing in the process changes if we assume a distinct value of μ_{jm} that is a function of the month in addition to the state. The EM algorithm simply fits 12 means instead of 1 mean per state and calculates probabilities $P\left(\ell_t = 1 \mid a_T, a_{T-1}, ..., a_1, \theta_m\right)$ as described above.

2.A.3 Cost Factors

Damage to Vessels

Direct damage is typically due to attempts to board a vessel. This could be damage due to small arms fire or rocket propelled grenades. Damages to the cargo are typically small, at least in bulk shipping which we focus on, while damage to the hull is more common.[61] As a consequence, the risk to hulls has now been unbundled from the Hull and Machinery (H&M) insurance and put into special War Risk Insurance. The War Risk Insurance is typically an annual police, but additional premiums are charged if vessels travel through high risk areas. These premiums are passed on to the charterers. In May 2008 the Joint War Committee, an advisory body set up by the maritime underwriters based in London, declared the Gulf of Aden to be an area of high risk for which these additional premiums apply. The high risk area has since then expanded considerably and now covers the whole large rectangle in Figure 2.2.[62] Cargo insurances do not typically charge additional premiums for specific sea areas.[63] Since hull damage is covered by insurance we expect such costs to be passed on to ship charterers.

Loss of Hire and Delay

The distribution of costs coming from loss of hire depends on the individual chartering agreements. These determine to what extent a charterer has to pay the daily chartering rate for the time that a ship is being held by pirates. According to an industry norm the charterer is responsible for the first 90 days following seizure.[64] With an estimated rolling average of 205 days under seizure at the end of 2010 this implies a relatively even share of costs.[65] The risk of not being operational after release (due to damage to ship during captivity) is with the ship owner. This risk is substantial as immobility of several months without maintenance is bound to incapacitate a ship.

Ransom Payments

Ransom payments and the costs of negotiators typically reach several million dollars and are, in principle, shared between the owner of the vessel, a chartering party and the owner of the cargo or special insurances that these parties purchased.[66] However, this applies only on journeys with cargo on board. In addition, the crew falls into the ship owners obligations if brought off the ship.[67] Both the ship owner's H&M insurance and the war risk insurance will cover part of this ransom. Kidnap and

[61] Hastings (2009) stresses that cargo is not stolen during captivity in the case of Somalia because the infrastructure for transporting it off is lacking.

[62] For details see `http://www.goo.gl/MOg7S`.

[63] See Marsh's Global Marine Practice available at `http://www.goo.gl/vhXoJ`.

[64] This norm is the "BIMCO Piracy Clause 2009". BIMCO is the largest international shipping association representing ship-owners.

[65] For a summary see MARSH (2011).

[66] See `http://www.goo.gl/jSO3f`.

[67] For a discussion see MARSH (2011) and `http://www.goo.gl/vhXoJ`, accessed on 10.04.2012.

Ransom (K&R) insurance policies, introduced in 2008, provide additional cover for the payment of ransoms. It is unclear what proportion of ships are insured by these policies.[68] However, the fact that these are designed for shipowners is indicative that these bear the main burden of ransom payments.

Even if ransoms are not paid, ship-owners need to pay a significant wage risk bonus to crew when travelling through pirate territory. According to the International Maritime Employers' Council (IMEC) seafarers are entitled to a compensation amounting to 100% of the basic wage on each day a vessel stays in a high risk area.

Security

The maritime industry's Best Practices manual lists a long list of changes to ship and crew stretching from barbed wire, high pressure fire hoses and citadels to additional security teams, that can help prevent a successful pirate attack/hijack.[69] All these expenses will be borne by the ship owner. The notion of an "arms race" between better equipped pirates and ever more sophisticated defence mechanisms by ship owners suggests that there might be costs on the side of ship owners that exceed the expected sum of ransom payments. According to *The Economist* newspaper, some 40% of ships carried security crews by 2012.[70] Conversations with industry experts suggest that the price per security crew of four is fixed and does not generally vary with the type of ship under consideration. The quoted price we work with in the paper is 3000 USD per day for a crew of four.

Re-routing, Speeding-up

The cost of re-routing around the Cape of Good Hope, especially among very large vessels, has been highlighted as a major element of the costs of piracy in early publications on the issue.[71] In the public debate this notion was often supported by a drastic decrease in Suez canal traffic in 2008. However, Suez canal traffic data can be misleading in this regard as world bulk trade collapsed only a few months before the increase in pirate activity. In addition, it should be kept in mind that large Capesize Bulk Carriers were never able to cross the Suez canal and would go around the Cape regardless of pirate activity. Indeed, more recent evidence using satellite imaging suggests that re-routing around the Cape is likely to be a minor issue.[72] Rerouting costs are in principle fully recoverable from the charterer since contracts are written for daily ship hire. A different issue are additional fuel costs which are borne by the charterer under the time-charter.

[68] Though some industry experts claim that as of 2009, the proportion of ships covered by such policies was less than 10 %, see `http://www.goo.gl/Uh3zX`.

[69] These are updated regularly. The version referred to here is BMP4 (2011) "Best Management Practices for Protection against Somalia Based Piracy".

[70] *Laws and guns*, The Economist, April 14th 2012.

[71] See, for example, One Earth Future (2010) and Bendall (2011).

[72] See One Earth Future (2011).

Additional Wage Costs

There are large welfare costs borne by the captured individuals in hijacking incidents. With captivity lasting on average 11 months and a high level of physical violence the hijacking risk looms large for individuals. In addition, according to the Ocean's Beyond Piracy think tank 3,863 seafarers were fired upon by pirates in 2012. It is difficult to measure this human cost in monetary terms. Still, one way to capture it is the wage compensation that seafarers receive when shipping through piracy areas. After negotiations in the International Bargaining Forum (IBF) it was agreed that workers should be entitled to a 100 percent basic wage bonus to compensate for travels through war risk areas. As this cost is directly borne by the shipowner it will also increase shipping rates in a competitive market.

The bottom line from this discussion is that looking at contract prices in shipping should pick up a good deal of the increased costs imposed by piracy. However, we would expect this to be a lower bound on the overall cost to the shipping industry since some of the direct costs paid by charterers may not be captured. This issue is taken into account in our welfare calculations.

2.A.4 Welfare Cost Calculations

Basic Estimate

The first column in Table 2.7 reports:

$$L^1(\Delta) = [\Delta - \tau(\Delta)]\,\nu\hat{X}(\psi + \nu[c + \Delta]).$$

In this Appendix, we first present the calculations for column (1) in Panel (A) and (B). We then discuss the calculations of column (2) and (3).

Total Cargo shipped through the Suez Canal is around $646,064,000$ tons per year.[73] According to data from Stopford (2009) bulk ships travel at around 26km per hour (14 knots) and the average distance that charters travel which pass through the Gulf of Aden is $16,400$ km with a typical charter length of 26.3 days. To this we add 4 days on charter for loading and unloading. This does not include waiting time in Suez and neglects the possibility of re-routing.

Our estimates in Panel B in Table 2.7 add the costs imposed by piracy on maritime traffic through the broader Somali area to this cost. In order to calculate this we use the same estimates as before and estimate the number of tonnage travelling through this area (but not the Gulf of Aden) we use COMTRADE data on commodity trade between the Middle East and Africa/Asia (excluding India).[74] The data suggests that

[73] See `http://www.goo.gl/J31Gl`.

[74] For this we define two groups of countries and calculate total tons of trade between the two groups. A) Middle Eastern countries: Bahrain, Iran, Iraq, Kuwait, Oman, Pakistan, Qatar, Saudi Arabia, United Arab Emirates; and B) Africa/Asia: Angola, Australia, Bangladesh, Cambodia, China, Hong Kong SAR, Macao SAR, Dem. People's Rep. of Korea, Fmr Dem. Rep. of Vietnam, Fmr Rep. of Vietnam, Indonesia, Japan, Kenya, Madagascar, Malaysia, Mozambique, Myanmar, Nepal, New Zealand, Philippines, Rep.

about $578,000,000$ tons were shipped through the area in 2010. Most of this is oil exports from the Middle East. As before, we use our data to calculate the average charter length (20.67 days) and the average charter rate (0.4646 USD/DWT days).
Low estimate:
Gulf of Aden:

$$\begin{aligned} 0.082 * 0.4726 * 30.3 * 646064000 &= 758 \text{ million USD} \\ & \quad -120 \text{ million USD} \\ &= 638 \text{ million USD} \end{aligned}$$

Gulf of Aden+Indian Ocean:

$$\begin{aligned} 0.082 * 0.4726 * 30.3 * 646064000 &= 758 \text{ million USD} \\ 0.082 * 0.4648 * 20.67 * 578000000 &= 455 \text{ million USD} \\ & \quad -120 \text{ million USD} \\ &= 1,093 \text{ billion USD.} \end{aligned}$$

High estimate:
Our high estimate uses the estimate on the dummy on war area risk from Column (4) Table 2.5 to derive the costs of piracy. That estimate suggests that piracy leads to an increase of charter rates by 12.3%.
Gulf of Aden:

$$\begin{aligned} 0.123 * 0.4726 * 30.3 * 646064000 &= 1,137 \text{ million USD} \\ & \quad -120 \text{ million USD} \\ &= 1017 \text{ million USD} \end{aligned}$$

Gulf of Aden+Indian Ocean: we use the same coefficient but apply it to the Indian Ocean Trade Thus:

$$\begin{aligned} 0.123 * 0.4726 * 30.3 * 646064000 &= 1,137 \text{ million USD} \\ 0.123 * 0.4648 * 20.67 * 578000000 &= 683 \text{ million USD} \\ & \quad -120 \text{ million USD} \\ &= 1,700 \text{ billion USD.} \end{aligned}$$

Column (2) in Table 2.7 applies the additional factor derived in equation (2.11). Details are in the following section.

of Korea, Singapore, South Africa, Sri Lanka, Thailand, United Rep. of Tanzania, Viet Nam.

Quantity Effects

Formula for $L^2(\Delta)$ The general formula for the welfare loss can be written

$$\begin{aligned} V(\psi+\nu[c+t]) - V(\psi+\nu[c+\Delta]) &= Q(t) \\ &\simeq Q(\Delta) + Q'(\Delta)[t-\Delta] + \frac{1}{2}Q''(\Delta)[t-\Delta]^2 . \end{aligned}$$

Note that

$$V(\psi+\nu[c+t]) = U\left(\hat{X}(\psi+\nu[c+t])\right) - \hat{X}(\psi+\nu[c+t])[\psi+\nu[c+t]] .$$

When we derive the partial derivative using

$$\frac{\partial U\left(\hat{X}(\psi+\nu[c+t])\right)}{\partial \hat{X}(\psi+\nu[c+t])} = \psi+\nu[c+t]$$

we find that

$$Q'(t) = -v\hat{X}(\psi+\nu[c+t]) .$$

Now observe that:

$$\begin{aligned} Q(\Delta) &= 0 \\ Q'(\Delta) &= -\nu\hat{X}(\psi+\nu[c+\Delta]) \\ Q''(\Delta) &= -\nu^2\hat{X}'(\psi+\nu[c+\Delta]) \end{aligned}$$

We assume that the demand function has a constant price elasticity η so that we can write

$$\hat{X}(\psi+\nu[c+t]) = (\psi+\nu[c+t])^{-\eta} .$$

and inserting all this we get an approximation of the welfare loss

$$\begin{aligned} & Q(\Delta) + Q'(\Delta)[t-\Delta] + \frac{1}{2}Q''(\Delta)[t-\Delta]^2 \\ = \ & \nu\hat{X}(\psi+\nu[c+\Delta])[\Delta-t] - \frac{1}{2}\nu^2\hat{X}'(\psi+\nu[c+\Delta])[t-\Delta]^2 \\ = \ & \nu\hat{X}(\psi+\nu[c+\Delta])[\Delta-t]\left[1+\frac{1}{2}\eta\frac{\nu(\Delta-t)}{\psi+\nu[c+t]}\right] \\ = \ & \nu\hat{X}(\psi+\nu[c+\Delta])[\Delta-t]\left[1+\frac{1}{2}\hat{\eta}\frac{\Delta-t}{c+\Delta}\right] \\ \geq \ & \nu\hat{X}(\psi+\nu[c+\Delta])[\Delta-\tau(\Delta)]\left[1+\frac{1}{2}\frac{\Delta-\tau(\Delta)}{c+\Delta}\hat{\eta}\right] \end{aligned}$$

where we have replaced the trade elasticity with regard to price η (which we do not have) with the trade elasticity with regard to transport costs, $\hat{\eta}$ (available from the trade literature). Observe that the trade elasticity with respect to transport costs, $\hat{\eta}$, in terms of our model is

$$\hat{\eta} = \frac{\partial \log X}{\partial \log \phi} = \eta\frac{\phi}{\psi+\phi}$$

so that, using the definition of ϕ above, we get

$$\eta = \hat{\eta}\frac{\psi + \nu\left[c+\Delta\right]}{\nu\left[c+\Delta\right]}.$$

The last approximation uses the fact that $\tau(\Delta) \leq t$. So this gives a lower bound on the welfare loss and depends on observables. Comparing this to equation (2.10) we have that

$$L^2(\Delta) \simeq L^1(\Delta)\left[1 + \frac{1}{2}\frac{\Delta - \tau(\Delta)}{c+\Delta}\hat{\eta}\right].$$

Implementation In the low estimate the relative increase in transport costs due to piracy is

$$\frac{\Delta}{c+\Delta} = 0.082$$

while in the high estimate it is

$$\frac{\Delta}{c+\Delta} = 0.123.$$

We use four different estimates for $1 - \frac{\tau(\Delta)}{\Delta}$. The low Gulf of Aden estimate is

$$1 - \frac{\tau(\Delta)}{\Delta} = 1 - \frac{120 \text{ million USD}}{758 \text{ million USD}} = 0.84$$

the other estimates are calculated analogously.

There are several possible numbers we could use for $\hat{\eta}$. Latest results from Feyrer (2009) who uses the Suez Canal closure as a shock to distance and calculates the effects on trade from distance costs suggests that an estimate between 0.2 and 0.5 for $\hat{\eta}$ is realistic. The estimate found in a meta study in Disdier (2008) is 0.9. Given the similarity of the Feyrer study we use the estimate of 0.5 in column 2. This leads to an adjustment of

$$\begin{aligned} L^2(\Delta) &= L^1(\Delta) \times \left[1 + \frac{1}{2}\left(1 - \frac{\tau(\Delta)}{\Delta}\right)\frac{\Delta}{c+\Delta}\hat{\eta}\right] \\ &= L^1(\Delta) \times 1.017 \end{aligned}$$

for the low estimate in the Gulf of Aden. This is applied to the whole welfare loss caused by price increases. For the low estimate in the Gulf of Aden this is

$$(758 \text{ million USD} - 120 \text{ million USD}) \times 1.017 = 649 \text{ million USD}.$$

Insurance Averaging

The general average insurance rules imply that the cost of piracy is borne by both cargo owners as well as by the ship owners. It is the ship owners, who in turn pass on this cost to the chartering parties in form of higher chartering rates. This is what we estimate in our main specification. However due to the general average principle, this effect is underestimated, since the ship owner's insurer pays only a share of the piracy cost in cases in which the ship is laden. In this Appendix we describe at how we arrive at the scaling factor $\zeta > 1$ used in the welfare estimates shown in the main text (Table 2.7).

The first step is to estimate the market value of the vessels in our dataset. Second, we estimate the values of the cargo that these ships transport. The ratio of the values is indicative for general average rules. In a third step, we estimate the share of ballast journeys, in order to correct for the fact that, during these journeys, the ship owner bears the entire cost of piracy.

From weekly market reports of the ship brokerage firm Intermodial[75], we obtained recorded sales of dry bulk vessels on the second hand market for 2010. In total, there were 402 recorded transactions. For a subset of 379 of these transactions, we know the age of the ship, the vessel's deadweight tonnage and the value of the transaction. Using these data on transactions, we can estimate the value of the ships 2010 in our dataset for the year. These estimates use two common controls in both data-sets: the age of ship and its tonnage to carry out this matching. Clearly, there are many more controls that correlate with the price that a vessel achieves on the market. However, we abstract from these due to data limitations. Either way, our estimated values are likely constitute a lower bound on a ship's value due to the standard adverse selection problem.

Using the 379 recorded sales, we estimate a regression of the form:

$$\text{ShipPrice}_l = \beta_0 + \beta_1 Age_l + \beta_2 DWT_l + \epsilon_l$$

Using the estimated coefficients, we generate fitted values for our main sample for the ships in 2010. The estimated values for vessels travelling through the Suez Canal in our sample are as follows:

Quartile	Value (USD)
Lower Quartile	26,791,260
Median	32,637,280
Upper Quartile	37,281,280

This compares well with industry-wide figures published by ship brokerage firms. For 2010, Intermodal for example reports that a five year old Panamax vessel with 75,000 tons deadweight was estimated to be worth 39 Million USD. In our dataset, the median ship on the Aden route is 7 years old, i.e. slightly older and with 73,726 tons deadweight slightly smaller. This makes us confident that the fitted ship values are indeed reasonably realistic for 2010.

We estimate the value of the cargo carried by the dry bulk ships in our sample using Suez Canal traffic statistics. These provide a very crude disaggregation into the different types and quantities of goods carried through the Suez Canal. We try to link this disaggregation with average commodity price data for the year 2010 obtained

[75] These reports can be accessed on `http://www.goo.gl/RmUZU`.

from the IMF and the World Bank. Any matching to these average commodity values is quite crude since the Suez authorities, for example, do not decompose such broad categories as cereals, ores and metals, coal and coke or oil seeds.[76] With this caveat, we match to our data using four main commodity prices: coal, iron ore, soybean and wheat. Using the traffic statistics on these four broad commodities, we compute the value of the average ton of these commodities passing through the Suez canal.

Using this, we estimate the value of the average ton of dry bulk carried through the Suez Canal. Using the median ship in our dataset, this allows us to estimate the value of cargo. We compute lower- and upper-bound values for these estimates using plain commodity prices for coal and wheat. This yields the following range of estimates:

Cargo type	Price (USD) per Ton	Cargo Value (USD)
(Low value) Coal cargo	106.03	7,451,675.41
Average Suez dry bulk cargo	165.97	11,663,908.20
(High value) Wheat	223.67	15,719,087.90

Using these estimates, we can compute the ratio of the cargo to ship value. However, using this share as a scaling factor ζ, without correcting for the share of ballast (i.e. without cargo) journeys, we are likely to underestimate the general average share paid by the ship owner. Using Suez canal traffic data, we find that, in 2010, 25.7% of the dry bulk carrier transits were ballast journeys. Hence, the general average share of the ship owner is:

$$\zeta = (1 - b)(1 - \text{cargo}/\text{ship}) + b$$

where b is the share of the journey in ballast.

Using this, we arrive at the following general average shares for our median ship value:

Cargo type	Cargo-to-ship value	ζ
Average Suez dry bulk cargo	0.35738	0.7346

The value of ζ from this table is used in the Table 2.7 to estimate the welfare loss.

This implies that $L^1(\Delta)$ can underestimate the welfare cost by a factor of up to 2.13. Combined with our high estimate this would imply an increase in chartering cost by 27%. However, for reasons laid out in section 2.1.3 this is likely to be an upper

[76] These four commodities make up at least 48.3 % of all commodities in the Suez traffic that can broadly be classified as (dry) bulk cargo.

bound. The low estimate for Aden, for example, can then be calculated as

$$
\begin{aligned}
& 758 \text{ million USD } \times 2.13 \\
& -120 \text{ million USD} \\
= \; & 1,495 \text{ billion USD.}
\end{aligned}
$$

This is the figure reported in column (3), Table 2.7.

2.A.5 Cost of Military Operations

While somewhat sketchy, our estimates in Table 2.7 can be augmented to include the costs of naval operations which try to limit pirate activities. The costs of Atalanta for the European Union in 2009 was 11 million USD.[77] To this we need to add the costs of the EU member countries. The only available estimates indicate that additional operational costs for the German military involvement (1 vessel, 300 personal) in 2010 was around 60 million USD.[78] Since the overall size of the Atalanta mission is between 4 and 7 vessels this indicates total costs of about 340 million USD for the Atalanta mission.

In addition to Atalanta there are two more operations which are, at least partially, occupied with preventing piracy attacks: NATO's Ocean Shield and the Combined Force 151. Causality from piracy to the presence of some of the military forces in the Arabian sea is harder to establish. For example, the Combined Force 151 includes two US aircraft carriers stationed there.

[77] See `http://www.goo.gl/hrqPA`, accessed on 10.04.2012.

[78] Deutscher Bundestag Drucksache 17/179. *Fortsetzung der Beteiligung bewaffneter deutscher Streitkräfte an der EU-geführten Operation Atalanta zur Bekämpfung der Piraterie vor der Küste Somalias.*

Chapter 3

Group Lending Without Joint Liability

While joint liability lending by microfinance institutions (MFIs) continues to attract attention as a key vehicle of lending to the poor, recently some MFIs have moved away from explicit joint liability towards individual lending. The most prominent such institutions are Grameen Bank of Bangladesh and BancoSol of Bolivia.[1] However, interestingly, Grameen and others have chosen to retain the regular group meetings that traditionally went hand-in-hand with joint liability lending.

Now it should be pointed out that in the absence of good panel data on lending methods it cannot be conclusively said that there has been a significant overall decline in joint liability among MFIs worldwide just on the basis of various anecdotes about a handful of high-profile MFIs. Indeed, existing evidence suggests that joint liability continues to be widely used. For example, de Quidt et al. (2012) use a sample of 715 MFIs from the MIX Market (Microfinance Information Exchange) database for 2009, and estimate that 54% of loans are made under "solidarity group" lending as opposed to "individual" lending.[2]

Nevertheless, these phenomena raise the question of the costs and benefits of using joint liability, and the choice between group loans with and without (explicit) joint liability. Besley and Coate (1995) is one of the first papers to point out both benefits and costs of joint liability: joint liability can increase repayment rates by inducing borrowers to repay on behalf of their unsuccessful partners but there are also states of the world where an individual borrower may default because of this burden, even if she was willing to pay back her own loan. Using a limited enforcement or "ex-

[1] For a discussion of the reasons for the shift in Grameen Bank's lending strategy, see Muhammad Yunus's article "Grameen Bank II: Lessons Learnt Over Quarter of A Century," at `http://www.grameen.com/index.php?option=com_content&task=view&id=30&Itemid=0`, accessed 18 December 2012.

[2] An earlier study Cull et al. (2009) puts this number at 51% using 2002/04 data involving 315 institutions. The year 2009 is one for which the largest cross-section of lending methodologies is available. Solidarity group loans defined by MIX as those for which "some aspect of loan consideration depends on the group, including credit analysis, liability, guarantee, collateral, and loan size and conditions." Individual loans are "made to individuals, and any guarantee or collateral required comes from that individual." We excluded 154 "village banks" for which lending methodology is unclear. See `http://www.mixmarket.org/about/faqs/glossary`.

post moral hazard" framework introduced by Besley and Coate (1995) in the group lending context, in this paper we study two issues raised by this apparent shift.

First, we analyze how by leveraging the borrowers social capital, individual liability lending (henceforth, IL) can mimic or even improve on the repayment performance and borrower welfare of explicit joint liability (EJ). When this occurs, we term it "implicit joint liability" (IJ). For this argument to work, there is no need for group lending per se - borrowers can, in theory, sustain this without any explicit effort on the part of the lender. Second, to understand better the logic of group lending, we introduce a purely operational argument for its use under IL, namely, it simply reduce the lender's transactions costs, shifting the burden to the borrowers. This is valuable because lower interest rates relax the borrowers' repayment incentive constraints, increasing repayment and welfare. We then show how this related to first issue: group lending may contribute to the creation of social capital, and therefore, may induce IJ.[3]

Next we carry out some simple simulation exercises using empirically estimated parameters. The goal is to complement the theoretical analysis and to get a quantitative sense of the welfare effects as well as the relevant parameter thresholds that determine which lending method is preferred. Our key findings are as follows. First, in low social capital environments, EJ does quite well compared to IJ. For example, when the standard deviation of project returns of 0.5, for social capital worth 10% of the loan size, the welfare attainable under IJ is 32.4% lower compared to the welfare under EJ. However, with social capital worth 50% of the loan size, the welfare attainable under EJ is 5% lower to the one attainable under IJ. Second, we find that the interest rate, repayment rate and borrower welfare are all rather insensitive to social capital under EJ, whereas in the case of IJ, they are all highly sensitive. This is what we would expect, since the only sanction available under IJ is coming through social capital. Third, when project returns are high variance, the welfare gains from higher social capital are quite large under IJ, which is not the case under EJ. To illustrate consider the case where project returns have a standard deviation of 0.5. If borrowers share social capital worth 10% of the loan size, borrower welfare under IJ is 35.9% lower than that of borrowers who share social capital worth 50% of the loan size.

Our analysis is motivated by two influential recent empirical studies. Giné and Karlan (2011) found that removing the joint liability clause, but retaining the group meetings, of a random subset of borrowing groups of Green Bank in the Philippines had no meaningful effect on repayment rates. In our model, this outcome arises when the newly individually liable groups have sufficient social capital to continue to assist one another with repayments, as under EJ. Secondly, Feigenberg et al. (2011) randomly varied the meeting frequency of individually liable borrowing groups of the Village Welfare Society in India. They found that groups who met more frequently

[3]It could even be that without the group, borrowers would be less able to interact. Indeed, in some conservative societies, social norms may prevent women from attending social gatherings (for instance under the Purdah customs in some parts of India and the Middle East). Then externally mandated borrowing groups can be a valuable vehicle for social interaction. See, for instance Sanyal (2009), Anderson et al. (2002), Kabeer (2005).

had subsequently higher repayment rates. In particular, they present evidence suggesting that this is due to improved informal insurance among these groups due to higher social capital. Both Giné and Karlan (2011) and Feigenberg et al. (2011) find evidence for intra-group transfers to help a borrower repay her loan even without explicit joint liability.[4] We argue that more frequent group meetings give borrowers a stronger incentive to build social capital, and that this is then leveraged to generate IJ. Grameen Bank states that Grameen II is designed to "lean on solidarity groups: small informal groups consisting of co-opted members coming from the same background and trusting each other."[5] The emphasis on trust suggests that the group continues to play an important role in Grameen's lending methodology beyond simply moderating the lender's transaction costs.

The main conclusions of our analysis is that it is premature to write off EJ as a valuable contractual tool and group lending without (explicit) joint liability may still harness some of the benefits of joint liability via implicit joint liability. Thus far we have one high quality randomized study of contractual form (Giné and Karlan (2011)) in which EJ seems not to play an important role. However in our theoretical analysis there are always parameter regions over which EJ is the most efficient of the simple contracts we analyze. A recent randomized control trial by Attanasio et al. (2011) finds stronger consumption and business creation impacts under EJ (albeit no significant difference in repayment rates - note that in their context mandatory group meetings are not used under either IL or EJ). Carpena et al. (2010) analyze an episode in which a lender switched from IL to EJ and found a significant improvement in repayment performance. For the same reasons, Banerjee (2012) stresses the need for more empirical work in the vein of Giné and Karlan (2011) before concluding that EJ is no longer relevant.

It is instructive to briefly look at the types of contracts currently used by MFIs. As mentioned, from the MIX dataset, 54% of borrowers were borrowing under what are classified as solidarity group loans. Although the solidarity group loans might not correspond exactly to pure EJ, this is the best measure we have. Our concept of IJ is most relevant to the "individual" category; the MIX Market notes that "loans based on consideration of the sole borrower, but disbursed through and recollected from group mechanisms, are still considered individual loans." A notable example is the Indian MFI Bandhan, which is one of the top MFIs in India, and is listed as having 3.6m outstanding loans in 2011, all classified as "individual". Bandhan does not use joint liability but disburses the majority of its loans through borrowing groups. Unfortunately, we do not have data on the method of disbursement of the

[4] In table IX of Giné and Karlan (2011) we see that conversion to individual liability caused a decrease, significant at 10%, in side-loans between borrowers, although no significant effect on borrowers "voluntarily [helping others] repay loans". Note that one challenge of interpreting these results in the light of our analysis is that group composition changed in Giné and Karlan (2011)'s experiment, while our model analyzes contract choice for a given level of social capital. Converted centers tended to take in members that were less well-known by existing members, presumably because individual liability made doing so less risky.

[5] See `http://www.grameen-info.org/index.php?option=com_content&task=view&id=33&Itemid=107`, accessed 28th September, 2012.

full sample of loans classified as individual, but it seems likely that many institutions are indeed using groups to disburse individual loans. This paper highlights how this may improve welfare through two channels: first of all, borrowers with sufficient social capital can mutually insure one another and secondly, attending costly group meetings may give borrowers incentives to invest in social capital.

Much of the existing theoretical work has sought to show how explicit joint liability improves repayment rates (see Ghatak and Guinnane (1999) for a review). In the model of Besley and Coate (1995), joint liability gives borrowers an incentive to repay on behalf of their partner when the partner is unable to repay her own loan. If borrowers can threaten social sanctions against one another, this effect is strengthened further. However, there are two problems with EJ. Firstly, since repaying on behalf of a partner will be costly, incentive compatibility requires the lender to use large sanctions and/or charge lower interest rates, relative to individual liability.[6] Secondly, when a borrower is unsuccessful, sometimes EJ induces the successful partner to bail them out, but sometimes it has a perverse effect, inducing them to default completely, while under IL they would have repaid. Rai and Sjöström (2004) and Bhole and Ogden (2010) approach these issues from a mechanism design perspective - designing cross-reporting mechanisms or stochastic dynamic incentives that minimize the sanctions used by the lender. Baland et al. (2010) provide an alternative explanation of the apparent trend away from what we call EJ towards IL, based on loan size. They find that the largest loan offered under IL cannot be supported under joint liability and that the benefits of the latter are increasing in borrower wealth. We do not focus on this angle but briefly touch on the issue of loan size in section 3.1. Allen (2012) shows how partial EJ, whereby borrowers are liable only for a fraction of their partner's repayment, can improve repayment performance by optimally trading off risk-sharing with the perverse effect on strategic default. In contrast, we focus on how simple group lending with no joint liability can achieve some of these effects, as side-contracting by the borrowers can substitute for the lender's enforcement mechanism.

Our model is also related to Rai and Sjöström (2010). In that paper, borrowers are assumed to have sufficient social capital to support incentive-compatible loan guarantees through a side-contract between borrowers, provided they have sufficient information to enforce such side contracts. The role of groups is to provide publicly observable repayment so as to enable efficient side-contracting. In contrast, in our setting, repayment behavior is common knowledge among the borrowers, and it is the amount of social capital that is key. Groups play a role that depends on meeting costs introduced in the next two sections. Secondly, in our model, borrowers are better off when they guarantee one another as their probability of contract renewal is higher. In Rai and Sjöström (2010) this is not the case as the lender is simply assumed to use a punishment that simply imposes a utility cost on the borrowers in case of

[6]This issue is the focus of the analysis in Rai and Sjöström (2010). Because of this, de Quidt et al. (2012) show that with a for-profit monopolist lender borrowers are better off under EJ than IL lending, because the lender must typically charge lower interest rates under EJ.

default. In fact, the optimal contract delivers the same borrower welfare whether they guarantee one another or not.

Other than the above mentioned papers, our paper is also broadly related to the theoretical literature in microfinance that have emerged in the light of the Grameen Bank of Bangladesh abandoning explicit joint liability and switching to the Grameen II model, focusing on aspects other than joint liability, such as sequential lending (e.g., Chowdhury (2005)), frequent repayment (Jain and Mansuri (2003), Fischer and Ghatak (2010)), exploring more general mechanisms than joint liability (e.g., Laffont and Rey (2003)), and exploring market and general equilibrium (Ahlin and Jiang (2008); **?** and de Quidt et al. (2012)).

The paper is structured as follows: in section 3.1 we present the basic model where in principle lending may take place with or without group meetings. We introduce our concept of implicit joint liability and show when it will occur and be welfare improving. Section 3.2 formalizes a key transaction cost in group and individual lending - the time spent attending repayment meetings. Section 3.3 then shows how meeting costs can give borrowers incentives to invest in social capital, and shows when this is welfare improving. Section 3.4 presents results of a simulation of the core model, while section 3.5 summarises the results and concludes.

3.1 Model

We model a lending environment characterized by costly state verification and limited liability. Borrowers are risk neutral, have zero outside option, no capital and limited liability. They have access to a stochastic production technology that requires 1 unit of capital per period with expected output $\bar{R}$, and therefore must borrow 1 per period to invest (we assume no savings for simplicity). There are three possible output realizations, $R \in \{R_h, R_m, 0\}$, $R_h \geq R_m > 0$ which occur with positive probabilities p_h, p_m and $1 - p_h - p_m$ respectively. We define the following:

$$\begin{aligned} p &\equiv p_h + p_m \\ \triangle &\equiv p_h - p_m \\ \bar{R} &\equiv p_h R_h + p_m R_m. \end{aligned}$$

We will refer to p as the probability of "success", and $\bar{R}$ as expected output.

We assume that output is not observable to the lender and hence the only relevant state variable from his perspective is whether or not a loan is repaid. Since output is non-contractible, the lender uses dynamic repayment incentives, as in Bolton and Scharfstein (1990). We assume that if a borrower's loan contract is terminated following a default, she can never borrow again. Under individual liability (IL), a borrower's contract is renewed if she repays and terminated otherwise. Under explicit joint liability (EJ), both contracts are renewed if and only if both loans are repaid.

Now we introduce the notion of social capital used in the paper. We assume that pairs of individuals in the village share some pair-specific social capital worth S in discounted lifetime utility, that either can credibly threaten to destroy. In other words, a friendship yields lifetime utility S to each person. If the social capital is destroyed it is lost forever. We assume that each individual has a very large number of friends, each worth S. Thus each friendship that breaks up represents a loss of size S.[7]

We assume a single lender with opportunity cost of funds equal to $\rho > 1$. In the first period, the lender enters the community, observes S and commits to a contract to all potential borrowers. The contract specifies a gross interest rate, r and EJ or IL. We assume the lender to be a non-profit who offers the borrower welfare maximizing contract, subject to a zero-profit constraint.[8]

In this section we ignore the role of groups altogether - being in a group or not has no effect on the information or cost structure faced by borrowers and lenders. Although borrower output is unobservable to the lender, we assume it is observable to other borrowers. As a result, they are able to write informal side contracts to guarantee one another's repayments, conditional on the output realizations. For simplicity, in the theoretical analysis we assume such arrangements are formed between pairs of borrowers.[9]

EJ borrowers will naturally side contract with their partner, with whom they are already bound by the EJ clause. Specifically, we assume that once the loan contract has been fixed, pairs of borrowers can agree a "repayment rule" which specifies each member's repayment in each possible state $Y \in \{R_h, R_m, 0\} \times \{R_h, R_m, 0\}$. Then in each period, they observe the state and make their repayments in a simultaneous-move "repayment game". Deviations from the agreed repayment rule are punished by a social sanction: destruction of S. The repayment rule, social sanction and liability structure of the borrowing contract thus determine the payoffs of the repayment game and beliefs about the other borrower's strategy. To summarize, once the lender has entered and committed to the contract, the timings each period are:

1. Borrowers form pairs, and agree on a repayment rule.

2. Loans are disbursed, borrowers observe the state and simultaneously make repayments (the repayment game).

3. Conditional on repayments, contracts are renewed or terminated and social

[7]One way to conceptualize S is as the net present value of lifetime payoffs in a repeated "social game" played alongside the borrowing relationship, similar to the multi-market contact literature, such as Spagnolo (1999), who models agents interacting simultaneously in a social and business context, using one to support cooperation in the other. As an illustration, suppose the borrowers play the following "coordination" stage-game each period: if both play A, both receive s. If one plays A and the other, B, both receive -1. If both play B, both receive 0. Clearly, both (A, A) and (B, B) are Nash equilibria in the stage-game. If players expect to play (A, A) forever, their expected payoff is $S \equiv \frac{s}{1-\delta}$. However, switching to (B, B) forever as a social sanction is always a credible threat, and can be used to support the repayment rule.

[8]We abstract from other organizational issues related to non-profits, see e.g. Glaeser and Shleifer (2001).

[9]This could be for example because there are two types of investment project available and returns within a project type are perfectly correlated, such that side-contracting with another borrower who has the same project type yields no benefit. In the simulations we extend the analysis to larger groups.

sanctions carried out.

4. If an IL borrower's partner was terminated but she repaid, she rematches with a new partner.

We restrict attention to repayment rules that are stationary (depending only on the state) and symmetric (do not depend on the identity of the borrower). This enables us to focus on the stationary value function of a representative borrower. Stationarity also rules out repayment rules that depend on repayment histories, such as reciprocal arrangements. In addition, we assume that the borrowers choose the repayment rule to maximize joint welfare. Welfare maximization implies that social sanctions are never used on the equilibrium path, since joint surplus would be increased by an alternative repayment rule that did not punish this specific deviation.

Given repayment probability π, the lender's profits are:

$$\Pi = \pi r - \rho$$

and therefore the zero-profit interest rate is:

$$\hat{r} \equiv \frac{\rho}{\pi}. \tag{3.1}$$

By symmetry, each borrower i pays $\pi r = \rho$ per period in expectation.

There are two interesting cases that arise endogenously and determine the feasibility of borrowers guaranteeing one another's loans. In Case A $R_m \geq 2r$ and hence a successful borrower can always afford to repay both loans. In Case B we have $R_h \geq 2r > R_m \geq r$, thus it is not feasible for a borrower with output R_m to repay both loans. Case B will turn out to be the more interesting case for our analysis, since in this case there is a cost to using joint liability lending. Specifically there are states of the world (when one borrower has zero output and the other has R_m) in which under joint liability both borrowers will default, since it is not feasible to repay both loans and they will therefore be punished whether or not the successful partner repays her loan. Meanwhile under individual liability, the successful partner is able to repay her loan and will not be punished if she does so.

Consider Case A. If borrowers agree to guarantee one another's loans, they will repay in every state except $(0, 0)$, so the repayment probability is $\pi = 1 - (1 - p)^2 = p(2 - p)$, in which case $\hat{r} = \frac{\rho}{p(2-p)}$. Therefore Case A applies if $R_m \geq \frac{2\rho}{p(2-p)}$, i.e. when the successful partner can afford to repay both loans even if her income is only R_m. If this condition does not hold, then it will not be feasible for the successful borrower to help her partner in this state of the world, and therefore Case B applies.

Definition 1 *Case A applies when $R_m \geq \frac{2\rho}{p(2-p)}$. Case B applies when $R_m < \frac{2\rho}{p(2-p)}$.*

Suppose that borrowers only repay when both are successful, i.e. when both have at least R_m, which occurs with probability p^2. If this is the equilibrium repayment rate, then $\hat{r} = \frac{\rho}{p^2}$. We make a simple parameter assumption that ensures that this will be the highest possible equilibrium interest rate (lowest possible repayment rate), by ensuring that even with income R_m, borrowers can afford to repay $\frac{\rho}{p^2}$.

Assumption 1 $R_m \geq \frac{\rho}{p^2}$.

We also assume that R_h is sufficiently large that a borrower with R_h could afford to repay both loans even at interest rate $\hat{r} = \frac{\rho}{p^2}$. Since this is the highest possible equilibrium interest rate, this implies that R_h is always sufficiently large for a borrower to repay both loans.

Assumption 2 $R_h \geq 2\frac{\rho}{p^2}$.

To summarize, together these assumptions guarantee that $R_m \geq r$ and $R_h \geq 2r$ on the equilibrium path.

We can now write down the value function V for the representative borrower, which represents the utility from access to credit. Suppose that borrower i's loan is repaid with some probability π. Since the repayment rule is assumed to maximize joint welfare, it follows that borrowers's loans are only repaid when repayment leads to the loan contracts being renewed, and therefore the representative borrower's contract is also renewed with probability π. Since the lender charges zero profit interest rate $\hat{r} = \frac{\rho}{\pi}$, the borrower repays $\pi\hat{r} = \rho$ in expectation. Hence, her welfare is:

$$\begin{aligned} V &= \bar{R} - \rho + \delta\pi V \\ &= \frac{\bar{R} - \rho}{1 - \delta\pi}. \end{aligned} \tag{3.2}$$

For any borrower to be willing to repay her loan, it must be that the value of access to future loans exceeds the interest rate, or $\delta V \geq r$. If this condition does not hold, all borrowers will default immediately. We refer to this condition as Incentive Condition 1 (IC1), and it must hold under any equilibrium contract.

Provided IC1 is satisfied, borrower welfare is maximized by achieving the highest repayment rate possible. To see this, suppose the lender charges some interest rate r. Then $V = \frac{\bar{R} - \pi r}{1 - \delta\pi}$. It can be verified that this is increasing in π if and only if IC1 holds. Therefore, in the subsequent discussion the ranking of welfare will be equivalent to the ranking in terms of the repayment probability.

Using (3.2) and $\hat{r} = \frac{\rho}{\pi}$ we can derive the equilibrium IC1 explicitly:

$$\rho \leq \delta\pi\bar{R}. \tag{IC1}$$

By Assumption 1, the lowest possible equilibrium repayment probability π is equal to p^2. For the theoretical analysis we make the following parameter assumption that ensures IC1 is satisfied in equilibrium:

Assumption 3 $\delta p^2 \bar{R} > \rho$.

Now that the model is set up we analyze the choice of contract type.

3.1.1 Individual Liability

Suppose first of all that the borrower does not reach a repayment guarantee arrangement with a partner. Since IC1 is satisfied, the borrower will repay her own loan whenever she is successful, so her repayment probability is p. Her utility V is then equal to $\frac{\bar{R}-\rho}{1-\delta p}$.

Now we consider when pairs of IL borrowers will agree a repayment guarantee arrangement. If this occurs, we term it implicit joint liability (IJ).

Since IC1 holds, the borrowers want to agree a repayment rule that maximizes their repayment probability. There are many possible such rules that can achieve the same repayment rate, so for simplicity we focus on the most intuitive one, whereby borrowers agree to repay their own loan whenever they are successful, and also repay their unsuccessful partner's loan if possible.[10]

We already know that repayment of the borrower's own loan is incentive compatible by IC1. For it to be incentive compatible for her to repay on behalf of her partner as well, it must be that social sanction outweighs the cost of the extra repayment, i.e. $r \leq \delta S$. This gives us a constraint which we term IJ Incentive Constraint 2, or IJ IC2. For equilibrium interest rate $\hat{r} = \frac{\rho}{\pi^{IJ}}$ IJ IC2 reduces to:

$$\rho \leq \delta \pi^{IJ} S. \tag{IJ IC2}$$

There is a threshold value of S, $\hat{S}^{IJ}$, such that IJ IC2 holds for $S \geq \hat{S}^{IJ}$:

$$\hat{S}^{IJ}_k \equiv \frac{\rho}{\delta \pi^{IJ}_k}, k \in \{A, B\},$$

where k denotes the relevant case. When $S \geq \hat{S}^{IJ}$, it is feasible and incentive compatible for borrowers to guarantee one another's loans, and therefore they will do so as this increases the repayment probability and thus joint welfare. Therefore IJ applies for $S \geq \hat{S}^{IJ}$.

Next we work out the equilibrium repayment probabilities and interest rates in cases A and B respectively. Assume $S \geq \hat{S}^{IJ}$. In Case A, a successful borrower can always afford to repay both loans, so both loans are repaid with probability $\pi^{IJ}_A \equiv 1-(1-p)^2 = p(2-p)$. In Case B, both loans are repaid whenever both are successful, and in states $(R_h, 0), (0, R_h)$. In state $(R_m, 0)$, borrower 1 cannot afford to repay borrower 2's loan, so she repays her own loan, while borrower 2 defaults and is replaced in the next period with a new partner. Therefore $\pi^{IJ}_B \equiv p^2 + 2p_h(1-p) + p_m(1-p) = p + p_h(1-p)$. Notice that both π^{IJ}_A and π^{IJ}_B are greater than p.

The lender observes whether Case A or Case B applies, and the value of S in the community, and offers an individual liability contract at the appropriate zero profit interest rate. Equilibrium borrower welfare under individual liability is equal to:

[10] An example of an alternative, less intuitive rule that can sometimes achieve the same repayment rate but cannot do better is where borrowers agree to repay their partner's loan, and then repay their own as well if they can afford to do so.

$$V_k^{IL}(S) = \begin{cases} \frac{\bar{R}-\rho}{1-\delta p} & S < \hat{S}_k^{IJ} \\ \frac{\bar{R}-\rho}{1-\delta\pi_k^{IJ}} & S \geq \hat{S}_k^{IJ} \end{cases}, k \in \{A, B\}.$$

It is straightforward to see that as S switches from less than $\hat{S}_k^{IJ}$ to greater than or equal to it, $V_k^{IL}(S)$ goes up as $\pi_k^{IJ} > p$.

3.1.2 Explicit Joint Liability

Now we analyze EJ contracts. Recall that under EJ, a pair of borrowers are offered a contract such that unless both loans are repaid, both partners lose access to credit in the future. The advantage of this contractual form is that it gives additional incentives to the borrowers to guarantee one another's loans. However, the disadvantage is that when borrower i is successful and j is unsuccessful, there may be states in which borrower i would repay were she under individual liability, but she will default under joint liability because she is either unwilling or unable to repay both loans.

The borrowers will agree a repayment rule, just as under IJ. Since this will be chosen to maximize joint welfare, it will only ever involve either both loans being repaid or both defaulting, due to the joint liability clause that gives no incentive to repay only one loan. Subject to this, because IC1 holds, joint welfare is maximized by ensuring both loans are repaid as frequently as possible.

IC1 implies that when both borrowers are successful, they will both be willing to repay their own loans. We therefore need to consider i's incentive to repay both loans when j is unsuccessful. Borrower i will be willing to make this loan guarantee payment provided the threat of termination of her contract, plus the social sanction for failing to do so, exceeds the cost of repaying two loans. Formally, this requires $2r \leq \delta(V^{EJ} + S)$. We refer to this condition as EJ IC2. Rearranging, and substituting for $\hat{r} = \frac{\rho}{\pi^{EJ}}$, we obtain:

$$\rho \leq \frac{\delta\pi^{EJ}[\bar{R} + (1 - \delta\pi^{EJ})S]}{2 - \delta\pi^{EJ}}. \qquad \text{(EJ IC2)}$$

We can derive a threshold, $\hat{S}^{EJ}$, such that EJ IC2 is satisfied for $S \geq \hat{S}^{EJ}$:

$$\hat{S}_k^{EJ} \equiv \max\left\{0, \frac{\rho}{\delta\pi_k^{EJ}} - \frac{\delta\pi_k^{EJ}\bar{R} - \rho}{\delta\pi_k^{EJ}(1 - \delta\pi_k^{EJ})}\right\}, k \in \{A, B\}$$

where as before, k denotes the relevant Case.

Note that $\hat{S}^{EJ}$ can be equal to zero. This corresponds to the basic case in Besley and Coate (1995) where borrowers can be induced to guarantee one another even without any social capital. This relies on the lender's use of joint liability to give borrowers incentives to help one another, and is not possible under individual liability.

Provided $S \geq \hat{S}^{EJ}$, borrowers are willing to guarantee one another's repayments. The repayment rule will then specify that i repays on j's behalf whenever i can afford to and j is unsuccessful. If $S < \hat{S}^{EJ}$, borrowers will not guarantee one another. They will therefore only repay when both are successful.

We now derive the equilibrium repayment probability under each Case. Firstly, if $S < \hat{S}^{EJ}$, borrowers repay only when both are successful, so $\pi^{EJ} = p^2$ in either Case.

Now suppose $S \geq \hat{S}^{EJ}$. In Case A, both loans can be repaid whenever at least one borrower earns at least R_m. Thus the repayment probability is $\pi_A^{EJ} = p(2-p)$. In Case B, R_m is not sufficient to repay both loans. Therefore both loans are repaid in all states except $(0,0), (R_m, 0), (0, R_m)$. In these three states both borrowers default. The repayment probability is therefore $\pi_B^{EJ} = p^2 + 2p_h(1-p) = p + \triangle(1-p)$.

Borrower welfare is:

$$V_k^{EJ}(S) = \begin{cases} \frac{\bar{R}-\rho}{1-\delta p^2} & S < \hat{S}_k^{EJ} \\ \frac{\bar{R}-\rho}{1-\delta \pi_k^{EJ}} & S \geq \hat{S}_k^{EJ} \end{cases}, k \in \{A, B\}.$$

Note that $\hat{S}_A^{EJ} \leq \hat{S}_B^{EJ}$. This is because the interest rate is lower in Case A, and V is higher (due to the higher renewal probability), so the threat of termination is more potent.

Now that we have derived the equilibrium contracts assuming either IL or EJ, we turn to analyzing the lender's choice of contractual form in equilibrium, which will depend crucially on the borrowers' ability to guarantee one another's loans.

Let us define $V(S) \equiv \max\{V^{EJ}(S), V^{IL}(S)\}$ as the maximum borrower welfare from access to credit. Observe that the repayment probability and borrower welfare from access to credit, $V(S)$, are stepwise increasing in S.

3.1.3 Comparing contracts

In this section we compare borrower welfare under each contractual form. We have seen that EJ has the advantage that it may be able to induce borrowers to guarantee one another even when they have no social capital. However, in Case B it has a perverse effect: in some states of the world borrowers will default when they would have repaid under IL.

This is most acute when $p_m > p_h$. Then $\pi_B^{EJ} = p + \triangle(p_h - p_m) < p$. Therefore in Case B, EJ actually performs worse than IL for all levels of social capital - the perverse effect dominates. Thus for Case B, EJ would never be offered.

We have already derived thresholds for S, $\hat{S}^{IJ}$ and $\hat{S}^{EJ}$, above which borrowers will guarantee one another's loans under individual and joint liability respectively. The lender's choice of contract will depend on the borrowers ability to do so, so first we derive a lemma that orders these thresholds in Case A and Case B.

Lemma 1

1. $\hat{S}_A^{IJ} > \hat{S}_A^{EJ}$.

2. *Suppose* $p_h \geq p_m$. *Then* $\hat{S}_B^{IJ} > \hat{S}_B^{EJ}$.

Proof. See appendix. ■

Lemma 1 shows that supporting a loan guarantee arrangement requires more social capital under IL than under EJ. The reason for this is that the lender's sanction under EJ is a substitute for social capital in providing incentives to borrowers to guarantee one another.[11]

The lender is a non-profit who offers the borrower welfare-maximizing contract. Therefore he offers IL if $V^{EJ}(S) \leq V^{IL}(S)$ and EJ otherwise. This will depend on the Case (A or B), the sign of $\triangle$, and S. We summarize the key result of this section as:

Proposition 1 *The contracts offered in equilibrium are as follows:*

Case A: *IL is offered at $\hat{r} = \frac{\rho}{p}$ for $S < \hat{S}_A^{EJ}$, otherwise EJ is offered at $r = \frac{\rho}{\pi_A^{EJ}}$.*

Case B, $\triangle > 0$***:*** *IL is offered at $\hat{r} = \frac{\rho}{p}$ for $S < \hat{S}_A^{EJ}$, EJ is offered at $\hat{r} = \frac{\rho}{\pi_B^{EJ}}$ for $S \in [\hat{S}_B^{EJ}, \hat{S}_B^{IJ})$, IL is offered at $\hat{r} = \frac{\rho}{\pi_B^{IJ}}$ for $S \geq \hat{S}_B^{IJ}$.*

Case B, $\triangle \leq 0$***:*** *IL is offered at $\hat{r} = \frac{\rho}{p}$ for $S < \hat{S}_B^{IJ}$, IL is offered at $\hat{r} = \frac{\rho}{\pi_B^{IJ}}$ otherwise.*

Whenever EJ is offered borrowers guarantee one another's repayments. Whenever IL is offered and $S \geq \hat{S}^{IJ}$ borrowers guarantee one another's repayments.

Proof. See appendix. ■

The result is summarized in Table 3.1, which gives the equilibrium contract and repayment probability π in alternate rows. Borrower welfare is not shown, but is easily computed as $V = \frac{pR-\rho}{1-\delta\pi}$, is strictly increasing in π.

<table>
<tr><th></th><th>Case A</th><th>Case B, $\triangle > 0$</th><th>Case B, $\triangle \leq 0$</th></tr>
<tr><td>$S < \hat{S}^{EJ}$</td><td>IL (no IJ)
p</td><td>IL (no IJ)
p</td><td rowspan="2">IL (no IJ)
p</td></tr>
<tr><td>$S \in [\hat{S}^{EJ}, \hat{S}^{IJ})$</td><td>EJ
$p(2-p)$</td><td>EJ
$p + \triangle(1-p)$</td></tr>
<tr><td>$S \geq \hat{S}^{IJ}$</td><td>EJ
$p(2-p)$</td><td>IL (with IJ)
$p + p_h(1-p)$</td><td>IL (with IJ)
$p + p_h(1-p)$</td></tr>
</table>

Table 3.1: Equilibrium contracts and repayment probabilities

This table shows that there are clear trade-offs in the contractual choice. In Case A, IJ has no advantage over EJ because in both cases borrowers repay both loans whenever successful. Therefore IL is offered for low S, and EJ for high S. In Case B when $\triangle \leq 0$, we have already remarked that EJ is always dominated by IL. Therefore basic IL is offered for low S, and when S is high enough, borrowers will begin to

[11] A slight complication arises in the proof because in Case B the repayment probability is higher and therefore the interest payment is lower under IJ. As a result, the size of the guarantee payment that must be incentive compatible is actually smaller under IJ, but the net effect is still that borrowers are more willing to guarantee one another under EJ.

guarantee one another, leading to an increase in the repayment rate and a fall in the equilibrium interest rate.

The most interesting case is Case B for $\triangle > 0$. Here there is a clear progression as S increases. For low S, borrowers cannot guarantee one another under either contract, so basic IL is offered. For intermediate S, EJ can sustain a loan guarantee arrangement but IL cannot, so EJ is offered. Finally for high S, borrowers are able to guarantee one another under IL as well. Since this avoids the perverse effect of EJ, the lender switches back to IL lending.

3.1.4 A remark on loan size

For simplicity, our core model assumes loans of a fixed size. However we can allow for variable loan size as a simple extension. To keep things simple, we assume that borrowers require a loan of size L. The relation between loan size and output is linear, that is, with a loan of size L, output is LR_h with probability p_h, LR_m with probability p_m, and 0 otherwise. Therefore we can simply scale $\bar{R}$ and r by L, so borrower welfare is now equal to LV. However, borrowers' social capital is derived from relationships external to the production function and therefore is assumed not to depend on L. Thus for a given amount of social capital S, borrowers are less willing to guarantee one another's loans as the loan size increases.[12] Thus we have the following observation:

Observation 1 *$\hat{S}^{EJ}(L)$ and $\hat{S}^{IJ}(L)$ are increasing in loan size, L. For a given S borrowers are less likely to guarantee one another's repayments as loan sizes increase. The repayment probability is thus decreasing in L.*

Note that the region $L(\hat{S}^{IJ} - \hat{S}^{EJ})$ is increasing in L. In particular, as L increases, the region $[0, \hat{S}^{EJ})$ expands. Over this region, borrowers are receiving "basic" IL, and not guaranteeing one another. Thus this result suggests a simple intuition for the stylized fact that IL loans tend to be larger. When loan sizes are small, the borrowers' social capital can be tapped to smooth out occasional small imbalances in income. As loan sizes and incomes increase, this becomes less feasible. As borrowers become unwilling to guarantee one another's loans, EJ becomes unattractive as it induces the borrowers to default unless both are successful.[13]

3.1.5 Discussion

Borrowers form partnerships that optimally leverage their social capital to maximize their joint repayment probability. Thus when social capital is sufficiently high to generate implicit joint liability, IL lending can dominate EJ: borrower i no longer

[12] Formally, the IJ IC2 is $Lr \leq \delta S$ and the EJ IC2 is $Lr \leq \delta(LV + S)$. Both are tighter as L increases. Replacing $\bar{R}$ with $L\bar{R}$, we observe that $\hat{S}^{EJ}(L) = L\hat{S}^{EJ}$ and $\hat{S}^{IJ}(L) = L\hat{S}^{IJ}$.

[13] Baland et al. (2010) obtain a result that gives the same negative correlation between the use of IL and loan size. Our above result is different in a nuanced way. In their model the poorest borrowers need the largest loan. Hence, their model generates a positive correlation between loan size and poverty.

defaults in state $(R_m, 0)$. This does not however mean there is no role for EJ. In particular, for intermediate levels of social capital, EJ can dominate IL - social capital is high enough for repayment guarantees under EJ but not under IL. We analyze borrower welfare under EJ and IL/IJ quantitatively in the simulations.

The results of Giné and Karlan (2011) are consistent with our Case A. Here, IL and EJ lending can achieve the same repayment probability, provided S is sufficiently high. This does not imply that those same borrowers would repay as frequently if they were not able to side-contract. Giné and Karlan (2011) additionally find that borrowers with weak social ties are more likely to default after switching to IL lending - this is consistent with these borrowers having $\hat{S}^{EJ} \leq S < \hat{S}^{IJ}$, so they are unable to support implicit joint liability.

So far, we have ignored the use of groups for disbursal and repayment of loans. However, it is frequently argued (see e.g. Armendáriz de Aghion and Morduch (2010)) that group meetings generate costs that differ from those under individual repayment. In the next section we show that this may induce the lender to prefer one or the other. We then proceed to show that by interacting with the benefits from social capital, group meetings may induce the creation of social capital. This is consistent with the results of a field experiment by Feigenberg et al. (2011).

3.2 Meeting Costs

In this section we lay out a simple model of loan repayment meeting costs. This immediately suggests a motivation for the use of groups. Holding group repayment meetings shifts the burden of meeting costs from the lender to the borrowers. This enables the lender to reduce the interest rate, which in turn makes it easier for borrowers to guarantee one another. Then in the next section we explore how the use of groups might create social capital, and thus generate implicit joint liability.

Since we want to focus on the interplay between meeting costs and social capital under individual liability, we assume that Case B applies and $\triangle \leq 0$. Therefore we can ignore EJ and drop the A, B notation.

A common justification for the use of group meetings by lenders is that it minimizes transaction costs. Meeting with several borrowers simultaneously is less time-consuming than meeting with each individually. However, group meetings might be costly for the borrowers, as they take longer and are less convenient than individual meetings. We term IL lending to groups ILG and IL lending to individuals ILI.

We assume that loan repayment meetings have two components, each of which takes a fixed amount of time. For simplicity, we assume that the value of time is the same for borrowers and loan officers[14] Also, for simplicity, we assume that the cost of borrower time is non-monetary so that borrowers are able to attend the meeting even if they have no income. However, more time spent in meetings by the loan officer increases monetary lending costs, for example because more staff must be hired.

[14]This may not be too unrealistic. For example, the large Indian MFI, Bandhan, deliberately hires loan officers from the communities that they lend to.

Each meeting incurs a fixed and variable cost. The fixed cost includes travel to the meeting location (which we assume to be the same for borrower and loan officer for simplicity), setting up the meeting, any discussions or advice sessions that take place at the meeting, reminding borrowers of the MFI's policies, and so on. This costs each borrower and the loan officer an amount of time worth γ_f irrespective of the number of borrowers in the group. Secondly there is a variable cost that depends on the number of borrowers at the meeting. This time cost is worth γ_v per borrower in the meeting. This covers tasks that must be carried out once for each borrower: collecting and recording repayments and attendance, reporting back on productive activities, rounding up missing borrowers, and so on. As with the fixed cost, each borrower and the loan officer incurs the variable cost. We assume that for group loans, each borrower also has to incur the cost having to sit through the one-to-one discussion between the loan officer and the other borrower, i.e., in a two group setting, the total variable cost per borrower is $2\gamma_v$ whereas under individual lending, it is γ_v.

Therefore, in a meeting with one borrower, the total cost incurred by the loan officer is $\gamma_f + \gamma_v$, and the total cost incurred by the borrower is the same, bringing the aggregate total time cost of the meeting to $2\gamma_f + 2\gamma_v$. In a meeting with two borrowers the loan officer incurs a cost of $\gamma_f + 2\gamma_v$, and similarly for the borrowers. Thus the aggregate cost in this case is $3\gamma_f + 6\gamma_v$. The lender's cost of lending per loan under ILI is $\rho + \gamma_f + \gamma_v$. Under ILG it is $\rho + \frac{\gamma_f}{2} + \gamma_v$. Therefore the corresponding zero-profit interest rates are $\hat{r}^{ILI} \equiv \frac{\rho+\gamma_f+\gamma_v}{\pi}$ and $\hat{r}^{ILG} \equiv \frac{\rho+\frac{\gamma_f}{2}+\gamma_v}{\pi}$.

Accounting for these costs, per-period expected utility for borrowers under ILI is $\bar{R} - \rho - 2(\gamma_f + \gamma_v)$. Under ILG, the per-period utility is $\bar{R} - \rho - \frac{3}{2}(\gamma_f + 2\gamma_v)$.[15]

Of course, the first thing to check is whether one lending method is less costly than the other in the absence of any loan guarantee arrangement between borrowers. This is covered by the following observation:

Observation 2 *Suppose $S = 0$. The lender uses ILG if and only if $\gamma_v < \frac{\gamma_f}{2}$.*[16]

The intuition is straightforward. When $\frac{\gamma_f}{\gamma_v}$ is large, i.e., fixed costs are important relative variable costs (e.g., when a large part of repayment meetings is repetitious) it is economical to hold group meetings. However, the more time is spent on individual concerns, the more costly it is to the borrowers to have to attend repayment meetings in groups because they have to sit through all the bilateral exchanges between another borrower and the loan officer. Microfinance loans are typically highly standardized and so $\frac{\gamma_f}{\gamma_v}$ will be relatively large, which is consistent with the common usage of group lending methods in microfinance.

Now consider borrowers' incentives to guarantee one another's loans. First we observe that for a given γ_v, γ_f, half of the aggregate meeting cost per borrower is

[15] We need to adapt Assumptions 1, 2 and 3 to reflect the additional costs. We assume $R_m \geq \frac{\rho+\frac{1}{2}(\gamma_f+2\gamma_v)}{p^2}$, $R_h \geq 2\frac{\rho+\frac{1}{2}(\gamma_f+2\gamma_v)}{p^2}$, $\delta p^2 \bar{R} - \max\left\{(1+\delta p^2)(\gamma_f+\gamma_v), \left(\frac{1}{2}+\delta p^2\right)(\gamma_f+2\gamma_v)\right\} \geq \rho$.

[16] Proof: $S = 0$ implies IJ is not possible so $\pi = p$ under ILI and ILG. The result then follows from comparison of per-period borrower welfare.

borne by the lender under ILI, while only a third is borne by the lender under ILG. The lender passes on all costs through the interest rate, so inspecting the value functions suggests that it is innocuous upon whom the cost of meetings falls. In fact this is not the case. Consider once again IJ IC2: $r \leq \delta S$. The only benefit a borrower receives from bailing out her partner is the avoidance of a social sanction, while the cost depends on the interest payment she must make. Therefore a lending arrangement in which the lender bears a greater share of the costs, and thus must charge a higher interest rate, tightens IJ IC2. This gives us the next proposition, which is straightforward:

Proposition 2 *Borrowers are more likely to engage in IJ under group lending than individual lending:* $\hat{S}^{IJG} < \hat{S}^{IJI}$.[17]

The implication of this result is that there is a trade-off between minimizing total meeting costs, and minimizing those costs borne by the lender. It may actually not be optimal to minimize total costs as shown by the following corollary, the proof of which is straightforward and given in the appendix. This arises from the fact that in an environment where the borrowers' participation constraints are not binding, the lender does not put weight on the disutility costs of meetings (individual or group) to the borrowers.

Corollary 1 *Suppose* $S \in [\hat{S}^{IJG}, \hat{S}^{IJI})$. *Borrower welfare under ILG may be higher than under ILI, even if* $\gamma_v > \frac{\gamma_f}{2}$.

We have now set the stage to analyze the interaction between meeting costs and social capital.

3.3 Social capital creation

In this section we show how group lending can actually generate social capital that is then used to sustain IJ. This analysis is motivated by the findings of Feigenberg et al. (2011). In their experiment, borrowers who were randomly assigned to higher frequency repayment meetings went on to achieve higher repayment rates. The authors attribute this to social capital being created by frequent meetings, social capital which can then support mutual insurance.

We show two main results. Firstly, group lending may create social capital where individual lending does not. The reason is simply that forcing the borrowers to spend time together in group meetings gives them an added incentive to invest in getting to know one another, as this makes the time spent in group meetings less costly. The knock-on effect is then that individual liability in groups may outperform individual liability with individual meetings because the groups are creating social capital that is then being used to support IJ.

[17]Proof: Borrowers are willing to guarantee their partner's repayments provided $r \leq \delta S$. Plugging in for the interest rates under ILG and ILI, we obtain $\hat{S}^{IJG} = \frac{\rho + \frac{1}{2}(\gamma_f + 2\gamma_v)}{\delta \pi^{IJ}} < \frac{\rho + \gamma_f + \gamma_v}{\delta \pi^{IJ}} = \hat{S}^{IJI}$.

Secondly, we turn to a comparative static more closely related to the Feigenberg et al. (2011) finding. Our simple framework does not easily allow us to model varying meeting frequency, so instead we study the effect on social capital creation of increasing the meeting costs (γ_f or γ_v). We find that an increase in the amount of time spent in group meetings can induce borrowers to switch to creating social capital, and can in fact be welfare-increasing.

Suppose that initially borrowers do not have any social capital, because creating social capital is too costly. For example, borrowers must invest time and effort in getting to know and understand one another, extend trust that might not be reciprocated, and so forth. Assume that social capital can take two values only, 0 and $S > 0$ and for a pair to generate social capital worth S in utility terms, each must make a discrete non-monetary investment that costs them η. To make the analysis interesting, we assume that in the absence of microfinance, they prefer not to do so, namely, $\eta > S$.

Once we introduce group lending, social capital generates an indirect benefit, by enabling the formation of a guarantee arrangement. This may or may not be sufficient to induce them to make the investment - that would depend on how $\eta - S$ compares with the insurance gains from

.[18]

Suppose the lender offers ILI and S is sufficiently large to sustain IJ. If the borrowers prefer to invest in social capital, each time their partner defaults they must invest in social capital with their new partner. We obtain the following result:

Lemma 2 *Borrowers will not invest in social capital under ILI if:*

$$\eta - S > G_1. \tag{3.3}$$

where

$$G_1 \equiv \frac{p_h(1-p)\left[\delta\left(\bar{R} - \frac{\rho}{\pi^{IJ}}\right) - \frac{1+\delta\pi^{IJ}}{\pi^{IJ}}(\gamma_f + \gamma_v)\right]}{(1-\delta p)(1-\delta(p+\triangle(1-p)))}.$$

The proof is given in the appendix. The greater the welfare gain from insurance, the higher is G_1 so the more likely the borrowers will invest in social capital. If (3.3) holds, the only equilibrium under ILI is one in which the borrowers do not invest in social capital, and therefore are not able to guarantee one another's loans.

Now assume that under ILG, the per-meeting cost to borrowers is decreasing in S. Attending group meetings is a chore unless the other group members are friends, in which case it can be a social occasion. By forcing the borrowers to meet together, the lender might give them an incentive to create social capital, benefiting them.

For simplicity, we assume that the cost to the borrowers of the time spent in group meetings is $(1 - \lambda(S))(\gamma_f + 2\gamma_v)$. In particular, $\lambda(0) = 0$ and $\lambda(S) = \lambda > 0$. The larger is λ, the smaller the disutility of group meetings, and when $\lambda > 1$, borrowers

[18]Note that each time a borrower's partner defaults and is replaced, she must invest in social capital with the new partner in order to continue with IJ.

actually derive positive utility from group meetings that is increasing in the length of the meeting. We can now check when social capital will be created in groups.

Lemma 3 *Borrowers invest in social capital under ILG if:*

$$\eta - S \leq G_2. \tag{3.4}$$

where

$$G_2 \equiv \frac{p_h(1-p)\left[\delta\left(\bar{R} - \frac{\rho}{\pi^{IJ}}\right) - \frac{1+2\delta\pi^{IJ}}{2\pi^{IJ}}(\gamma_f + 2\gamma_v)\right] + \lambda(1-\delta p)(\gamma_f + 2\gamma_v)}{(1-\delta p)(1-\delta(p + \triangle(1-p)))}.$$

The proof is given in the appendix. The greater the welfare gain from insurance, the higher is G_2, but in addition, G_2 is increasing in λ, which represents the reduction in the cost of attending group meetings when the borrowers have social capital. The larger is G_2, the more likely borrowers are to invest in social capital.

Lemmas 2 and 3 suggest that there may exist an interval, $(G_1, G_2]$ for $\eta - S$ over which groups create social capital but individual borrowers do not. The condition for this to be the case is derived in the next proposition, which follows from straightforward comparison of (3.3) and (3.4):

Proposition 3 *If the following condition holds:*

$$\lambda > \frac{p_h(1-p)(\delta\pi^{IJ}\gamma_v - \frac{\gamma_f}{2})}{4\pi^{IJ}(1-\delta p)(\gamma_f + 2\gamma_v)} \tag{3.5}$$

then there exists a non-empty interval for $\eta - S$ over which both (3.3) *and* (3.4) *are satisfied. If $\eta - S$ lies in this interval, groups create social capital, and individual lending does not.*

This is a key result, as it shows that when creating social capital sufficiently offsets the cost to borrowers of attending group meetings, borrowing groups may create social capital and guarantee one another's loans, while individual borrowers may not do so. We can see that the threshold for λ in (3.5) is negative if $\frac{\gamma_f}{2} > \gamma_v > \delta\pi^{IJ}\gamma_v$ and so the condition (3.5) is always satisfied if group lending has a cost advantage to the lender. What can be checked is, even if this is not the case, and $\delta\pi^{IJ}\gamma_v - \frac{\gamma_f}{2} > 0$ the critical threshold for λ is always strictly less than 1 and therefore, there always exists a λ high enough (but strictly less than 1) such that the condition (3.5) would hold. However it does not yet establish that the use of groups is necessarily welfare-improving. In other words, observing that groups are bonding and creating social capital does not tell the observer that group lending is the welfare-maximizing lending methodology. All it tells us is that investment is preferred to no investment under ILG, and no investment is preferred to investment under ILI. The welfare ranking of these two will depend on the meeting costs, η and S. The following proposition addresses the welfare question.

Proposition 4 *Suppose condition* (3.5) *is satisfied and* $\eta - S \in (G_1, G_2]$. *Borrower welfare under ILG is higher than that under ILI if:*

$$\eta - S \leq G_3 \tag{3.6}$$

where

$$G_3 \equiv \frac{\delta p_h(1-p)\,(\bar{R}-\rho) + 2(1-\delta\pi^{IJ})(\gamma_f+\gamma_v) - \frac{1}{2}(1-\delta p)(\gamma_f + 2\gamma_v)(3-2\lambda)}{(1-\delta p)(1-\delta(p+\triangle(1-p)))}.$$

The proof is given in the appendix. G_3 is higher the larger is the meeting cost under ILI relative to under ILG. It is also increasing in λ, representing the reduction in the cost of attending group meetings when the borrowers have social capital. Note that (3.6) is always satisfied for sufficiently large λ.

The expressions G_1, G_2 and G_3 are somewhat unwieldy. The following proposition establishes a sufficient condition under which $G_1 < G_2 < G_3$, i.e. there is guaranteed to exist an interval for $\eta - S$ over which groups invest in social capital and individuals do not, and over which borrower welfare is higher under group than individual lending:

Proposition 5 *Suppose total meeting costs per borrower are weakly lower under ILG than ILI, i.e.* $\gamma_v \leq \frac{\gamma_f}{2}$. *Then* $G_1 < G_2 < G_3$, *i.e.:*

1. *There always exists an interval for* $\eta - S$ *over which groups create social capital and individuals do not.*

2. *Borrower welfare is weakly higher under ILG than ILI for all values of* $\eta - S$.

The proof is given in the appendix. The condition $\gamma_v \leq \frac{\gamma_f}{2}$ implies that ILG has a (weak) cost advantage over ILI, as was discussed in Observation 2. In addition, when $G_1 < \eta - S \leq G_2$, groups invest in social capital while individuals do not, and this gives ILG a further advantage.

3.3.1 Meeting frequency and social capital creation

Now we take this basic framework and carry out one particular comparative-static exercise, motivated by the findings of Feigenberg et al. (2011). They find that groups that were randomly assigned to meet more frequently have better long-run repayment performance, which they attribute to higher social capital and informal insurance within the group. It is not possible to model repayment frequency in our simple setup, but nevertheless our model is able to capture some of this intuition.

We model an increase in meeting frequency as an increase in meeting costs, represented by an increase in either γ_f or γ_v. The more time spent in group meetings, the greater the benefit from social interaction within those meetings, captured by λ. Intuitively, it may not be too costly to attend meetings once a month with a stranger, but the more frequent those meetings are, the greater the incentive the borrowers have to build social capital.

However, more frequent meetings require more of the loan officer's time as well, leading to higher lending costs and a higher interest rate. This reduces the borrowers' incentive to invest in S, since the higher meeting costs reduce the value of maintaining access to credit.

The net effect on borrowers willingness to invest in S is positive if λ is sufficiently large, as shown by the following proposition.

Proposition 6 *Increases in γ_f or γ_v make borrowers under group lending more willing to invest in social capital if and only if the following condition holds:*

$$\lambda > \frac{p_h(1-p)(1+2\delta\pi^{IJ})}{2\pi^{IJ}(1-\delta p)}. \tag{3.7}$$

The proof is immediate from inspection of (3.4). This proposition implies an interesting corollary: an increase in meeting costs can actually be welfare-improving, by inducing borrowers to invest in social capital and thus engage in implicit joint liability.

Corollary 2 *Suppose* (3.7) *holds. Then there exists a threshold at which increases in the costs γ_f or γ_v cause group borrowers to switch to creating social capital, and this is welfare-improving.*

The proof is given in the appendix. The reason for this result is that in the neighborhood of (3.4) binding, the no-investment equilibrium is inefficient. A marginal increase in the meeting cost can be enough to give the borrowers sufficient incentive to switch to the investment equilibrium, generating a strict welfare increase.

It is worth explaining here why it is that there may not be an investment equilibrium even when utility is strictly higher under the investment than the no-investment equilibrium. In fact the reasoning is straightforward: the welfare cost of switching from investment to no-investment may be high. This is because of two things: the repayment rate is lower in the no-investment equilibrium, and the interest rate is higher. However, a borrower considering whether to deviate under the investment equilibrium does not consider the effect on the interest rate, since this only changes in equilibrium. Hence the cost of deviating from a hypothetical investment equilibrium is lower than the cost of switching from investment to no-investment.

Proposition 3 derives a condition on λ under which groups are better able to create social capital than individual borrowers. Proposition 6 simply focuses on group lending and asks when higher meeting costs actually lead to more social capital creation. As meeting costs increase, two things occur. Firstly, the lender must charge a higher interest rate, which reduces borrower welfare and tightens IJ IC2. Secondly, the cost to borrowers of being in a group with a stranger increase: by creating social capital the cost to borrowers of time spent in meetings decreases by $\lambda(\gamma_f + 2\gamma_v)$. If

λ is sufficiently large, the second effect dominates and higher meeting costs increase the borrowers' incentive to invest in S.

Feigenberg et al. (2011) show that the improvement in repayment performance associated with higher meeting frequency approximately offset the increase in the lender's cost. This implies that among contracts with group meetings the total surplus was increasing in meeting frequency in their experiment. In our model, all surplus accrues to the borrower, so condition (3.7) is necessary for there to exist a region over which total surplus is increasing in the meeting frequency.

If the lender holds the interest rate fixed, as in Feigenberg et al. (2011), borrowers will be more willing to create social capital for a given increase in the meeting frequency (the extra cost is not passed on through a higher interest rate). However, a parallel condition must then hold for the increase in repayment frequency to offset the lender's costs.

3.4 Simulation

In this section, we simulate a simple extension of the model calibrated to empirically estimated parameters. This enables us to illustrate the costs and benefits of explicit joint liability and explore under which environments it will be dominated by individual liability lending that induces implicit joint liability.

We find that in low social capital environments, EJ does quite well compared to IJ. For example, when the standard deviation of project returns of 0.5, for social capital worth 10% of the loan size, the welfare attainable under IJ is 32.4% lower compared to the welfare under EJ. However, with social capital worth 50% of the loan size, the welfare attainable under EJ is 5% lower to the one attainable under IJ. We find that for social capital worth around 25% of the loan size, EJ and IJ perform approximately equally well in terms of borrower welfare. For lower values of S, EJ dominates, and for higher values of S, IJ dominates. This analysis thus gives us insights into the extent of the perverse effect of JL. With high S under IJ, the borrowers now have enough social capital to help one another when they can afford to do so, but are not penalized in states of the world where only some of the group can repay. We also find that the interest rate, repayment rate and borrower welfare are highly insensitive to social capital under EJ, whereas IJ is highly sensitive to social capital, since the only sanction available is coming through the social capital. For example, when the standard deviation of project returns is 0.5, the EJ net interest rate is 11.3%, while the IJ net interest rate ranges between 10.4% and 21.4% for levels of S valued at 10% to 50% of the loan size respectively. The difference in the interest rate translates correspondingly into borrower welfare. If borrowers share social capital worth 10% of the loan size, the attainable IJ welfare is $V^{IJ} = 2.29$, which is 35.9% lower compared to the IJ welfare of $V^{IJ} = 3.57$ attained by borrowers who share social capital worth 50% of the loan size. We also find that these welfare and interest rate differentials between low and high levels of social capital S are increasing in the variance of project returns.

From theory we know the basic trade off between EJ, Il and IJ and how that changes with social capital. What this analysis adds is to give a quantitative magnitude to the relevant thresholds and also suggests some policy implications. In low social capital environments, despite its well known costs (Besley and Coate (1995)) EJ is an effective device to induce repayment incentives and moreover, if the extent of social capital is not known *ex ante* it is a robust instrument. It also suggests a high payoff from encouraging investing in social capital given the welfare implications of higher S on borrower welfare in IJ.

3.4.1 Approach

We approach the simulations in a very similar way to de Quidt et al. (2012). Firstly, while it is theoretically convenient to model groups of size two, these require empirically implausibly high returns to investment for the borrowers to be able to repay on one another's behalf, so instead we extend the model to groups of size 5, the group size originally used by Grameen Bank and others. For simplicity, we carry over our concept of social capital unaltered to the larger groups. Previously a borrower who did not help her partner when the repayment rule stipulated she should was sanctioned by her partner. Now we simply assume she is sanctioned by the whole group, losing social capital worth S.

We express all units in multiples of the loan size and a loan term of 12 months. For example, if $S = 0.15$ this means the borrowers have social capital worth 15% of the loan size. We obtain our parameter values from the estimates in de Quidt et al. (2012). $\bar{R}$, the expected return to borrowers' investments, is set to 1.6, i.e. a 60% annual return, based on De Mel et al. (2008)'s preferred estimates of the rate of return to capital among microenterprises in Sri Lanka. The lender's cost of capital, ρ, is set to 1.098, which was estimated using lender cost data from the MixMarket database of financial information from MFIs around the world. Lastly, we set δ equal to 0.864. This is the midpoint between the value implied by the return on US treasury bills and a lower bound implied by the model in de Quidt et al. (2012).

The two key ingredients that drive the trade-off between explicit and implicit joint liability are the level of social capital and the shape of the borrowers' return distribution function. We do not have data on social capital, so instead we estimate the equilibrium interest rate, repayment rate and welfare for a range of values for S. This enables us to say, for example, how much social capital is required for implicit and joint liability to perform as well or better than explicit joint liability.

It is more difficult to explore how the shape of the returns distribution affects the trade-off between EJ and IJ. In the theoretical analysis it was convenient to illustrate the key intuition using a simple categorical distribution with three output values and associated discrete probabilities. With larger groups, this distribution function is less useful. It no longer gives a simple and intuitive set of states of the world in which EJ does and does not perform well (with a group of size n, there are 3^n possible states of the world). More problematic is that the distribution has four parameters (p_m, R_m, p_h, R_h), only one of which can be tied down by our calibrated value of $\bar{R}$. As a result, it is very difficult to perform meaningful comparative statics - there are too many degrees of freedom.[19]

Therefore, for the main simulations we use the most obvious two-parameter distribution function, the Normal distribution.[20] Fixing the mean at $\bar{R}$, we can vary the shape of the distribution by changing the standard deviation. The range for σ was chosen to obtain the highest and lowest possible repayment rates at which the lender is able to break even. For the benchmark simulations, we assume the borrowers' returns are uncorrelated, but we also allow for positive and negative correlations in an extension.

To simulate the model, for each contract we work out a welfare-maximizing repayment rule for the borrowing group, i.e. one that maximizes the repayment rate, subject to the borrowers' incentive constraints. Solving analytically for the equilibrium repayment probability (which then gives us the interest rate and borrower welfare) is complex, so instead we simulate a large number of hypothetical borrowing groups and use these to compute the equilibrium repayment probability. We describe the simulation approach in detail in appendix 3.B.

3.4.2 Results

The main results for uncorrelated borrower incomes are presented in Figure 3.1. The standard deviation σ of individual borrower returns is varied on the horizontal axis of each figure.

For the distribution and parameter values used, it turns out that individual liability is in fact marginally loss-making for all σ, so we just present results for implicit joint liability and explicit joint liability for values $S \in \{0.1, 0.3, 0.5\}$.

The figures show that increasing the variance of returns is bad for repayment and thus welfare under both contracts. This is unsurprising: higher variance income processes are more difficult to insure (the required transfers between members tend to be larger), so states in which members cannot or will not help one another out become more common. Increasing S partially mitigates this effect since it increases the size of incentive-compatible transfers between borrowers.

Our simulated repayment rates vary between around 85% to close to 100% as the variance of borrower income decreases. These high repayment rates follow from the fact that the calibrated mean return $\bar{R}$ is higher than the lender's cost of funds, ρ, so perfect repayment is attainable for sufficiently low variance. However, these values are fairly typical for microfinance repayment rates. For example, in de Quidt et al. (2012) we conservatively estimate a repayment rate in the MIX Market dataset

[19] We perform one exercise in the appendix, where we vary $p_h - p_m$ while holding p, R_h, R_m constant. The confound here is that the mean return also varies as we vary p_h and p_m.

[20] One complication arises, namely the possibility of negative income realizations. For simplicity, we allow these to occur, but we assume that only borrowers with positive incomes can assist others with repayment.

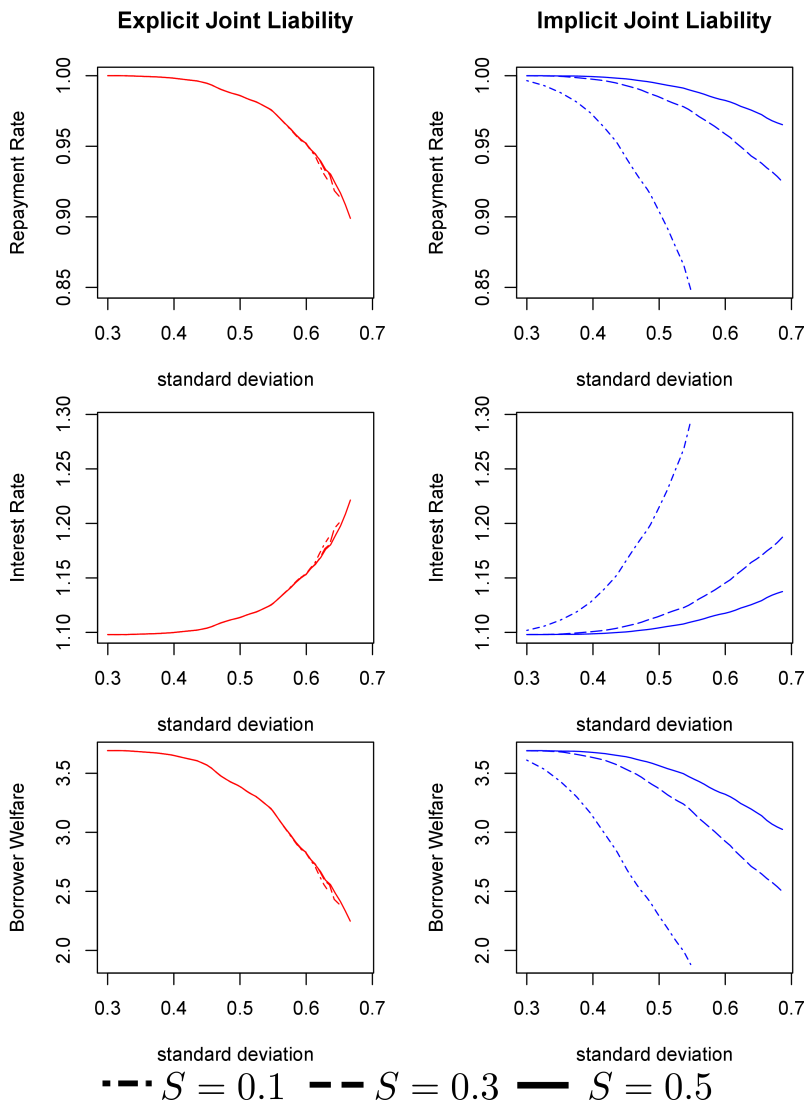

Figure 3.1: Simulation results for uncorrelated borrower returns. Explicit joint liability results are in the left column and implicit joint liability in the right column. Each figure plots the relevant object (repayment rate, interest rate and borrower welfare) for three levels of social capital, $S = 0.1, 0.3, 0.5$. The standard deviation of the individual borrower's income is varied on the horizontal axis of each figure.

of around 0.92. Using the simulated repayment rate, we can obtain the zero-profit interest rate and borrower welfare. The net interest rate varies between 10% and 30% per year (again, these are not unreasonable values for the microfinance context), while borrower welfare varies between around 1.8 and 3.7 multiples of the loan size.

One of the most striking lessons we learn from the graphs is that the interest rate, repayment rate and borrower welfare are highly insensitive to social capital under explicit joint liability. The reason is that social capital is only shifting the borrowers from default to repayment in states of the world where they can afford to help one another and where the joint liability penalty is not already sufficient. The probability that such a state occurs is lower, the bigger the sample of borrowers. Meanwhile, implicit joint liability is highly sensitive to social capital, since the only sanction available is coming through the social capital. For example, at $\sigma = 0.5$, the IJ repayment rate is 91% for $S = 0.1$, 98% for $S = 0.25$, and close to 100% for $S = 0.5$, while the EJ repayment rate is fixed at 98% throughout.[21]

In order to more easily compare EJ and IJ, in Figure 3.2 we overlay the welfare curves for EJ and IJ. The simulation exercise emphasizes much of the core intuition from the model. When S is low, explicit joint liability tends to dominate since the joint liability clause gives the borrowers an additional incentive to help one another. When S is high, implicit joint liability dominates, due to the perverse effect of JL - the borrowers now have enough social capital to help one another when they can afford to do so, but are not penalized in states of the world where only some of the group can repay.

To give a numerical example of the magnitudes of the welfare gains from EJ and IJ as a function of S, consider the case of a standard deviation of project returns of 0.5. Here for social capital worth 10% of the loan size for example, the welfare attainable under IJ, $V^{IJ} = 2.29$ is 32.4% lower compared to the welfare under EJ $V^{EJ} = 3.39$. This highlights the clear welfare gains that are possible under EJ in environments with low S. These gains disappear however for higher levels of S. With social capital worth 50% of the loan size, the welfare attainable under EJ $V^{EJ} = 3.39$ is in fact 5 % lower to the one attainable under EJ $V^{IJ} = 3.56$. The higher levels of social capital make it incentive compatible to help each other out, when they are able to, while not being punished when not the whole group is able to repay.

The graph also highlights that the EJ and IJ contracts are almost completely overlapping for intermediate values of $S = 0.3$ of the loan size, suggesting that in environments with intermediate levels of social capital both contracts can perform equally well.

While these results illustrate the problems with strict EJ,[22] we also interpret them as showing why EJ should not be prematurely dismissed as an important contractual

[21] Note that in de Quidt et al. (2012) we find that the interest rate and borrower welfare are sensitive to social capital when the lender is a monopolist, since higher social capital relaxes IC2, and therefore enables the lender to increase the interest rate. The non-profit lender, as modeled in this paper, does not do this.

[22] Problems that have also received attention in Besley and Coate (1995), Rai and Sjöström (2004), Bhole and Ogden (2010), Rai and Sjöström (2010) and Allen (2012).

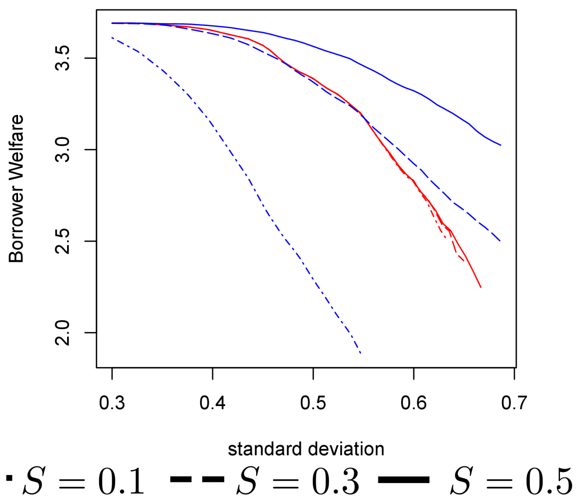

Figure 3.2: Simulation results for uncorrelated borrower returns. Explicit joint liability results are in red and implicit joint liability in blue. The figure plots borrower welfare for three levels of social capital, $S = 0.1, 0.3, 0.5$. The standard deviation of the individual borrower's income is varied on the horizontal axis.

tool (as also recently argued by Banerjee (2012)). Many of the candidates for alternative mechanisms discussed in the literature are complex and potentially difficult to implement, so we have focused on two extremely simple mechanisms that we feel are empirically relevant. What we find is that implicit joint liability can perform very well, provided borrowers have enough social capital: borrowers have to be willing to impose sanctions on one another worth at least 25% of their loan size. Meanwhile EJ functions well in our simulations even for low levels of social capital. This illustrates how important the lending environment, and in particular borrowers' social ties are for determining the preferred contract in our framework.

3.4.3 Correlated returns

As an extension, we now present simulation results when borrowers' returns are correlated. A number of recent papers have analyzed how correlated returns affect repayment behavior under joint liability lending.[23]. As a simple extension, we consider how our EJ and IJ borrowers are affected by introducing positively or negatively correlated returns into the model. We simulate the borrowing group's per-period income vector $[Y_1, ..., Y_n]$ as a multivariate Normal distribution. We fix the standard deviation at 0.5, the midpoint of the range considered in the previous section, and vary the pairwise correlation between group members from -0.25 to 0.45.[24] We graph the results in Figure 3.4.[25]

[23] For example, Laffont (2003), Ahlin and Townsend (2007) and Allen (2012)

[24] For correlation smaller than -0.25 we essentially have 100% repayment everywhere, and for greater than 0.45 there is typically no lending equilibrium.

[25] Note that the graphs are less smooth than those in Figure 3.1. This is because for the benchmark simulations we are able to reuse the same underlying random draws for each set of output realizations, simply by rescaling as the standard deviation changes. This is not possible when considering variously correlated returns, so we need to generate a new sample of borrower output realizations for each value

The main conclusion from this analysis is that for a given level of social capital, EJ is sufficiently more sensitive to the strength of correlation between borrower incomes. EJ requires all loans to be repaid. When borrower incomes are only weakly correlated, there will typically only be a small number of failures in a group, which are relatively easy for the other members to assist with. With a strongly positive correlation this is no longer the case, it becomes more common to have large numbers of failures. In this environment IJ is an advantage because the borrowers are not penalized when their partners default. This becomes evident when comparing the gradient of the IJ curves relative to the EJ curves as the correlation increases.

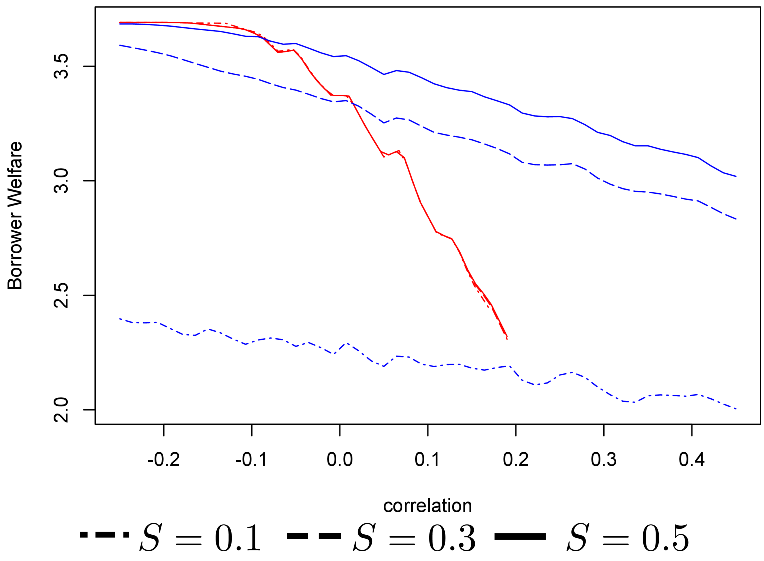

Figure 3.3: Simulation results for correlated borrower returns. Explicit joint liability results are in red and implicit joint liability in blue. The figure plots borrower welfare for three levels of social capital, $S = 0.1, 0.3, 0.5$. The correlation between pairs of borrower's returns is varied on the horizontal axis.

3.5 Conclusion

Anecdotal evidence suggests that there has been a move away from explicit joint liability towards individual liability by some prominent institutions. Most of these institutions have retained the use of groups to facilitate credit disbursal. The key question now is whether groups do more than just facilitate the lender's operations. The interest in this question has been strengthened by two recent field experiments by Giné and Karlan (2011) and Feigenberg et al. (2011).

The first of these, Giné and Karlan (2011), found that removing the joint liability clause, but retaining the group meetings, of a random subset of borrowing groups of Green Bank in the Philippines had no meaningful effect on repayment rates, although borrowers with weak social ties to other borrowers were more likely to drop out.

In this paper we have shown that this outcome may result when the newly in-

of the correlation coefficient, and this naturally introduces some extra noise.

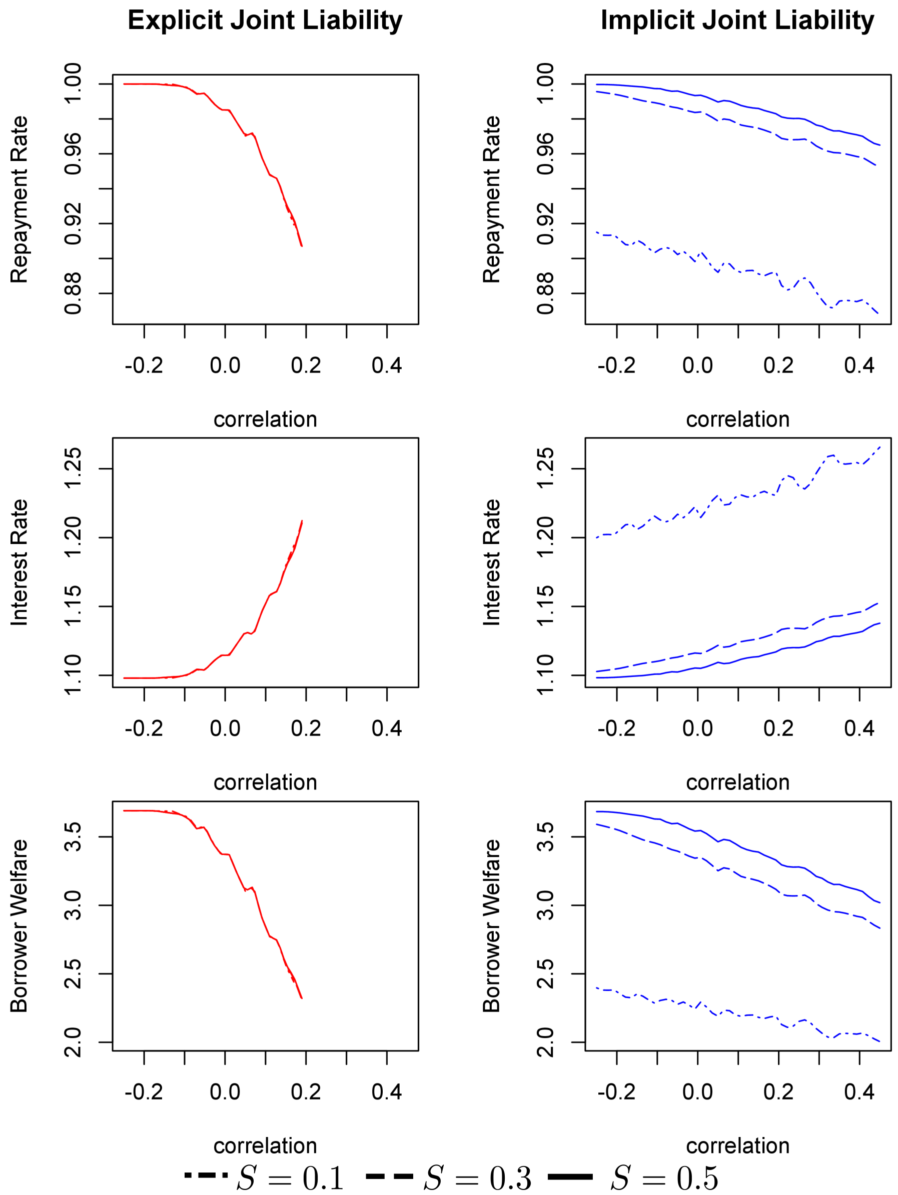

Figure 3.4: Simulation results for uncorrelated borrower returns. Explicit joint liability results are in the left column and implicit joint liability in the right column. Each figure plots the relevant object (repayment rate, interest rate and borrower welfare) for three levels of social capital, $S = 0.1, 0.3, 0.5$. The correlation between pairs of borrower's returns is varied on the horizontal axis of each figure.

dividually liable groups have sufficient social capital to continue to guarantee one another's repayments, as under EJ, which we call implicit joint liability (IJ). We show that this may even lead to higher repayment rates and borrower welfare. However this first result does not depend upon the use of groups, provided borrowers are able to side contract on loan repayments outside of repayment meetings.

We next show that when individual and group repayment meetings are costly, mutual insurance or IJ are easier to sustain under group lending, because IJ depends crucially on the interest rate, which in turn depends on the share of total meeting costs borne by the lender. Group meeting reduces the lender's share of meeting costs, enhancing the advantages of IJ.

The second experimental paper highlighting the role of groups is Feigenberg et al. (2011). They find that varying meeting frequency for a subset of individually liable borrowing groups seemed to have persistent positive effects on repayment rates. They suggest that this is due to improved informal insurance among these groups due to higher social capital.

We analyze situations under which microcredit might induce borrowers to create social capital, which in turn enables them to sustain IJ. We derive conditions under which group lending is more likely than individual lending to create social capital, and show when this is indeed welfare increasing. Finally, relating to one of the key findings of Feigenberg et al. (2011), we derive conditions under which more frequent meetings, modeled here as an increase in the amount of time borrowers and loan officers must spend in loan repayment meetings, increases borrowers' incentive to invest in social capital. This provides a theoretical foundation for Feigenberg et al. (2011)'s observation. We also carry out a simulation exercises to assess the quantitative magnitudes of the effects of alternative forms of lending, as well as some of the relevant thresholds of social capital.

3.A Mathematical appendix

Proof of Lemma 1

Proof. Comparing the expressions for $\hat{S}_A^{EJ}$ and $\hat{S}_A^{IJ}$, it is immediate that $\hat{S}_A^{EJ} < \hat{S}_A^{IJ}$ since $\pi_A^{EJ} = \pi_A^{IJ}$ and $\delta\pi_A^{EJ}\bar{R} - \rho > 0$ by Assumption 3.

Now consider Case B. It is obvious that if $\hat{S}^{EJ} = 0$, $\hat{S}^{IJ} > \hat{S}^{EJ}$, since $\hat{S}^{IJ} > 0$. Suppose therefore that $\hat{S}^{EJ} > 0$. It is straightforward to check that Assumptions 1, 2 and 3 imply that $\delta p \geq \frac{1}{2}$. Given this, and $p_h \geq p_m$, it follows that $\pi^{IJ} \geq p \geq \frac{1}{2}$ and $\pi^{EJ} \geq p \geq \frac{1}{2}$. Also using the fact that π_B^{EJ} can be written as $p^2 + 2p_h(1-p)$. We have:

$$\begin{aligned}
\hat{S}^{IJ} - \hat{S}^{EJ} &= \frac{\delta\pi_B^{EJ}\bar{R} - \rho}{\delta\pi_B^{EJ}(1-\delta\pi_B^{EJ})} + \frac{\rho}{\delta\pi_B^{IJ}} - \frac{\rho}{\delta\pi_B^{EJ}} \\
&= \frac{\pi_B^{IJ}(\delta\pi_B^{EJ}\bar{R} - \rho) - p_m(1-p)(1-\delta\pi_B^{EJ})\rho}{\delta\pi_B^{IJ}\pi_B^{EJ}(1-\delta\pi_B^{EJ})} \\
&\geq \frac{(\delta\pi_B^{EJ}\bar{R} - \rho) - p_m(1-p)\rho}{2\delta\pi_B^{IJ}\pi_B^{EJ}(1-\delta\pi_B^{EJ})} \\
&= \frac{\delta p^2\bar{R} - \rho + p_h(1-p)(2\delta\bar{R} - \rho) + (p_h - p_m)(1-p)\rho}{2\delta\pi_B^{IJ}\pi_B^{EJ}(1-\delta\pi_B^{EJ})} \\
&> 0
\end{aligned}$$

which follows from $2\delta\bar{R} - \rho > 0$ by Assumption 3. ■

Proof of Proposition 1

To compare IL and EJ, we consider first Case A, then Case B with $p_h > p_m$, and lastly Case B with $p_h \leq p_m$.

In Case A, borrower repayment guarantees under IL offer no advantage over EJ, so provided $S \geq \hat{S}_A^{EJ}$, EJ is the borrower welfare-maximizing contract (with indifference for $S \geq \hat{S}^{IJ}$). For $S < \hat{S}_A^{EJ}$, borrower will not mutually guarantee under EJ and also default unless their partner is successful, so IL is preferred to EJ:

$$V_A^{EJ}(S) - V_A^{IL}(S) = \begin{cases} -\frac{\delta p(1-p)(\bar{R}-\rho)}{(1-\delta p)(1-\delta p^2)} & S < \hat{S}_A^{EJ} \\ \frac{\delta p(1-p)(\bar{R}-\rho)}{(1-\delta p)(1-\delta p(2-p))} & S \in [\hat{S}_A^{EJ}, \hat{S}_A^{IJ}) \\ 0 & S \geq \hat{S}_A^{IJ} \end{cases}$$

In Case B, with $p_h > p_m$, EJ dominates IL when borrowers guarantee one another under EJ but not under IL, for $S \in [\hat{S}_B^{EJ}, \hat{S}_B^{IJ})$, so EJ is preferred in this region. However, once IJ is possible, for $S \geq \hat{S}_B^{IJ}$, it dominates EJ. This is because borrower 1 repays her own loan in state $(R_m, 0)$, while she would default under EJ. We have:

$$V_B^{EJ}(S) - V_B^{IL}(S) = \begin{cases} -\frac{\delta p(1-p)(\bar{R}-\rho)}{(1-\delta p)(1-\delta p^2)} & S < \hat{S}_B^{EJ} \\ \frac{\delta\triangle(1-p)(\bar{R}-\rho)}{(1-\delta p)(1-\delta(p+\triangle(1-p)))} & S \in [\hat{S}_B^{EJ}, \hat{S}_B^{IJ}) \\ -\frac{\delta p_m(1-p)(\bar{R}-\rho)}{(1-\delta(p+p_h(1-p)))(1-\delta(p+\triangle(1-p)))} & S \geq \hat{S}_B^{IJ} \end{cases}$$

Lastly, in Case B with $p_h \leq p_m$, EJ is always dominated by IL. This is because under EJ the highest possible repayment probability is $p + \triangle(1-p)$, which is weakly smaller than p, the lowest possible repayment probability under IL. Therefore we do not need to know the ordering of $\hat{S}_B^{EJ}$ and $\hat{S}_B^{IJ}$ for this case - EJ will never be used.

Proof of Corollary 1

Suppose total meeting costs are higher under ILG: $\frac{3}{2}(\gamma_f + 2\gamma_v) > 2(\gamma_f + \gamma_v)$ or $2\gamma_v > \gamma_f$. Suppose also that $S \in [\hat{S}^{IJG}, \hat{S}^{IJI})$. Then group lending sustains IJ but individual lending does not. Welfare is higher under group lending if:

$$\frac{\bar{R} - \rho - \frac{3}{2}(\gamma_f + 2\gamma_v)}{1 - \delta(p + p_h(1-p))} > \frac{\bar{R} - \rho - 2(\gamma_f + \gamma_v)}{1 - \delta p}$$

Taking the limit as $\gamma_f \to 2\gamma_v$, it is clear that this condition holds strictly, while $\hat{S}^{IJG} > \hat{S}^{IJI}$ continues to hold, thus the corollary follows for a non-trivial interval of costs by a standard open set argument.

Proof of Lemma 2

First, note that $\frac{\partial^2 V}{\partial r \partial \pi} < 0$. Therefore, the benefit of increasing π is higher when interest rates are low.

We want to find conditions under which ILI borrowers will not invest in social capital in equilibrium. To show this, we hypothesize a (low interest rate) equilibrium in which ILI borrowers do invest, and show that there exists a profitable deviation. Then, we know that in a (high interest rate) equilibrium in which borrowers do not invest, they will not wish to deviate to investing; this follows from $\frac{\partial^2 V}{\partial r \partial \pi} < 0$ as noted above.

Consider then a hypothetical equilibrium in which the borrowers do invest in social capital and repay with probability $\pi^{IJ} \equiv p + p_h(1-p)$. They are charged $\hat{r} = \frac{\rho+\gamma_f+\gamma_v}{\pi^{IJ}}$.

At the beginning of the first period, the borrower and her partner pay cost η and create social capital. Then, each period with probability $p + \triangle(1-p)$, both loans are repaid and both contracts renewed. With probability $p_m(1-p)$, only borrower i's loan is repaid. As a result, at the beginning of the next period, she must again pay cost η to create social capital with her new partner.[26]

[26]Since no social capital is destroyed on the equilibrium path, the S created with the original partner

Consider an ILI borrower in the first period, or one whose partner has just defaulted. We know that IC1 is satisfied, since by repaying her loan she can guarantee herself at least $\delta(\bar{R} - (\gamma_f + \gamma_v)) - \frac{\rho+\gamma_f+\gamma_v}{\pi^{IJ}}$ if she agrees with the new partner to simply take a loan and default immediately. This expression is positive by the modified Assumption 3 in footnote 15. Then we note that if it is an equilibrium for the borrower to invest in social capital, it must be that she does even better than this, and therefore IC1 must hold.

As we are considering an equilibrium in which she invests in social capital, we use an "IJI" superscript to denote the fact that IJ is taking place. If she invests in social capital with the new partner, she earns utility U, defined as follows:

$$U_1^{IJI} = S - \eta + W_1^{IJI}$$

where

$$W_1^{IJI} = \big(\bar{R} - \rho - 2(\gamma_f + \gamma_v)\big) + \delta(p + \triangle(1-p))W_1^{IJI} + \delta p_m(1-p)U_1^{IJI}.$$

The first term in W is the per-period utility under ILI. The second term represents the continuation payoff when both borrowers repay and have their contracts renewed. This occurs with probability $p + \triangle(1-p)$. In this case she earns W_1^{IJI} next period. The third term represents the continuation payoff if she repays but her partner defaults, which occurs with probability $p_m(1-p)$. In this case she matches with a new partner and therefore earns U_1^{IJI} next period.

Substituting for W, we can write U as:

$$\begin{aligned} U_1^{IJI} &= S - \eta + \frac{\big(\bar{R} - \rho - 2(\gamma_f + \gamma_v)\big) + \delta p_m(1-p)(S-\eta)}{1 - \delta\pi^{IJ}} \\ &= \frac{\big(\bar{R} - \rho - 2(\gamma_f + \gamma_v)\big) + (1 - \delta(p + \triangle(1-p)))(S-\eta)}{1 - \delta\pi^{IJ}}. \end{aligned}$$

Now we check for a one-shot deviation. In this context, a deviation is to defer investing in social capital by one period, i.e. to undergo one period without social capital (and therefore with repayment probability p), then invest in social capital next period. She prefers to deviate if:

$$U_1^{IJI} < \left(\bar{R} - p\frac{\rho + \gamma_f + \gamma_v}{\pi^{IJ}} - (\gamma_f + \gamma_v)\right) + \delta p U_1^{IJI}. \tag{3.8}$$

The first term on the right hand side represents the per-period utility of a borrower under ILI without social capital, paying an interest rate of $\hat{r} = \frac{\rho+\gamma_f+\gamma_v}{\pi^{IJ}}$ (intuitively, since the lender does not know she has deviated, the interest rate is not adjusted). With probability p her loan is repaid, and in the next period she invests in S, thus receiving continuation value U_1^{IJI}. Substituting for U_1^{IJI} and rearranging yields condition (3.3).

is not lost but cannot be leveraged in the credit contract.

Proof of Lemma 3

Hypothesize an equilibrium in which borrowers invest in social capital. We know that IC1 is satisfied, since by repaying her loan she can guarantee herself at least $\delta(\bar{R} - (\gamma_f + 2\gamma_v)) - \frac{\rho + \frac{1}{2}(\gamma_f + 2\gamma_v)}{\pi^{IJ}}$ if she agrees with the new partner to simply take a loan and default immediately. This expression is positive by the modified Assumption 3 in footnote 15.

We need to check that no borrower prefers to deviate by deferring their investment by one period, exactly as in Lemma 2. We define the value functions analogously to those in the proof of Lemma 2:

$$U_1^{IJG} = S - \eta + W_1^{IJG}$$
$$W_1^{IJG} = \left(\bar{R} - \rho - \frac{1}{2}(\gamma_f + 2\gamma_v)(3 - 2\lambda)\right) + \delta(p + \triangle(1-p))W_1^{IJG} + \delta p_m(1-p)U_1^{IJG}.$$

Where the possession of social capital reduces the borrowers' cost of group meetings by $\lambda(\gamma_f + 2\gamma_v)$. The appropriate substitutions yield:

$$U_1^{IJG} = \frac{\bar{R} - \rho - \frac{1}{2}(\gamma_f + 2\gamma_v)(3 - 2\lambda) + (1 - \delta(p + \triangle(1-p)))(S - \eta)}{1 - \delta\pi^{IJ}}.$$

There will be no deviation if $U_1^{IJG} \geq \left(\bar{R} - p\frac{\rho + \frac{1}{2}(\gamma_f + 2\gamma_v)}{\pi^{IJ}} - (\gamma_f + 2\gamma_v)\right) + \delta p U_1^{IJG}$. Simplifying yields condition (3.4).

Proof of Proposition 4

Total borrower welfare under ILI (where borrowers do not invest in social capital) is:

$$\begin{aligned} V^{ILI} &= \bar{R} - \rho - 2(\gamma_f + \gamma_v) + \delta p V^{ILI} \\ &= \frac{\bar{R} - \rho - 2(\gamma_f + \gamma_v)}{1 - \delta p} \end{aligned}$$

and when groups are used (and the borrowers do invest in social capital) it is:

$$U_1^{IJG} = \frac{\bar{R} - \rho - \frac{1}{2}(\gamma_f + 2\gamma_v)(3 - 2\lambda) + (1 - \delta(p + \triangle(1-p)))(S - \eta)}{1 - \delta\pi^{IJ}}.$$

as was derived in the proof of Lemma 3. The result then follows from comparison of these value functions.

Proof of Proposition 5

First, observe that if $\gamma_v \leq \frac{\gamma_f}{2}$, condition (3.5) is satisfied for all $\lambda \geq 0$, hence $G_1 < G_2$. From the proof of Lemma 3, $\eta - S \leq G_2$ if and only if $U_1^{IJG} \geq \frac{\bar{R} - p\frac{\rho + \frac{1}{2}(\gamma_f + 2\gamma_v)}{\pi^{IJ}} - (\gamma_f + 2\gamma_v)}{1 - \delta p}$. Call the RHS of this condition B. From the proof of Proposition 4, $\eta - S < G_3$ if and only if $U_1^{IJG} > V^{ILI}$. Finally, note that $B - V^{ILI} = \frac{p_h(1-p)(\rho + \gamma_f) + p\left(\frac{\gamma_f}{2} - \gamma_v\right)}{\pi(1 - \delta p)}$, which is strictly positive if $\gamma_v < \frac{\gamma_f}{2}$. Thus, $\eta - S \leq G_2$ implies $\eta - S < G_3$, or $G_2 < G_3$.

Claim 1 follows immediately from $G_1 < G_2 < G_3$. Claim 2, that borrower welfare is *always* higher under ILG, can be broken into three parts. Firstly, if $\eta - S \leq G_1$, both groups and individuals invest in social capital. Then, the cost advantage of ILG ($\gamma_v \leq \frac{\gamma_f}{2}$) implies that welfare is higher under ILG. Secondly, if $\eta - S > G_2$, neither groups nor individuals invest in S, and again the cost advantage leads to ILG dominating. Lastly, if $G_1 < \eta - S \leq G_2$, groups invest and individuals do not, and thus ILG dominates by Proposition 4.

Proof of Corollary 2

Suppose condition (3.4) binds, such that a small decrease in γ_f causes borrowers to stop investing in social capital. We want to show that this leads to a discontinuous decrease in welfare.

Before the change, welfare is:

$$U_1^{IJG} = \frac{\bar{R} - \rho - \frac{1}{2}(\gamma_f + 2\gamma_v)(3 - 2\lambda) + (1 - \delta(p + \triangle(1-p)))(S - \eta)}{1 - \delta\pi^{IJ}}.$$

after the change (in the limit as the increase in γ_f approaches zero), it is:

$$V^{ILG} = \frac{\bar{R} - \rho - \frac{3}{2}(\gamma_f + 2\gamma_v)}{1 - \delta p}$$

since the borrowers can no longer sustain IJ, so the new equilibrium is one in which they repay with probability p and the interest rate is $\frac{\rho + \frac{1}{2}(\gamma_f + 2\gamma_v)}{p}$. From condition (3.4) binding we know that:

$$\eta - S = \frac{p_h(1-p)\left[\delta\left(\bar{R} - \frac{\rho}{\pi^{IJ}}\right) - \frac{1+2\delta\pi^{IJ}}{2\pi^{IJ}}(\gamma_f + 2\gamma_v)\right] + \lambda(1-\delta p)(\gamma_f + 2\gamma_v)}{(1-\delta p)(1 - \delta(p + \triangle(1-p)))}. \quad (3.9)$$

For U_1^{IJG} to be strictly larger than V^{ILG} we require:

$$\frac{\bar{R} - \rho - \frac{1}{2}(\gamma_f + 2\gamma_v)(3 - 2\lambda) + (1 - \delta(p + \triangle(1-p)))(S - \eta)}{1 - \delta\pi^{IJ}} > \frac{\bar{R} - \rho - \frac{3}{2}(\gamma_f + 2\gamma_v)}{1 - \delta p}$$

which reduces to

$$\frac{\delta p_h(1-p)\left(\bar{R} - \rho - \frac{3}{2}(\gamma_f + 2\gamma_v)\right) + \lambda(1-\delta p)(\gamma_f + 2\gamma_v)}{(1-\delta p)(1 - \delta(p + \triangle(1-p)))} > \eta - S.$$

Substituting for $\eta - S$ from (3.9) and simplifying, we obtain:

$$2\delta\rho(1 - \pi^{IJ}) + (1 - \delta\pi^{IJ})(\gamma_f + 2\gamma_v) > 0$$

which is satisfied.

More generally, this demonstrates that the no-investment equilibrium is inefficient in the neighborhood of $\eta - S = G_2$. A marginal increase in the meeting cost that gives the borrowers greater incentive to invest in social capital can lead to a strict increase in borrower welfare.

3.B Simulation approach

This Appendix outlines the algorithm used to simulate the core model. The simulation was implemented in *R*. The intuition of the simulation procedure is very straightforward. We use a random sample of N groups with n members each. A group merely constitutes a vector of income realizations. These incomes are drawn from some distribution function F. We assume that F is a Normal distribution with $\mu = \bar{R} = 1.6$, however we allow the standard deviation σ to vary.

Given these income realizations, we compute the repayment rate that would arise under each contract for a given interest rate r. This process gives us a repayment probability function $\pi(r)$ under either contract.

Given this repayment probability function, we can then compute the break-even repayment rate and thus the break-even interest rate under each contract, along with borrower welfare. This then allows us to make comparisons between the two contractual forms.

We now describe in detail how the group-level repayment rate is computed, as this is different under each contract type due to the different incentive constraints.

We denote an income realization of a group i with n borrowers is represented by an n-vector, $Y_i = (y_1, ..., y_n)$, where y_j is group member j's income draw.

We want to find a repayment rule analogous to the one outlined in the theory that allows for larger groups and the continuous output distribution. The most obvious way to do this is to construct for each Y_i a "group bailout fund" that can be used for transfers between group members to assist with repayments. Since the incentive constraints differ between EJ and IJ, the construction of the group fund also differs and is described below.

Group Lending without Joint Liability

The relevant incentive constraint under group lending without joint liability implies that the maximum amount a group member j is willing to contribute to the group fund is $c_{ij} = max(y_{ij}, \delta S)$. All the transfers are put into a common pool C_j. This pool is then used to ensure the maximum possible number of repayments. The borrowers are sorted in ascending order of the amount of transfer they require to repay their own loan. and transfers made from the fund until it is exhausted.[27]. If m group members

[27]This in fact implies that in some cases the worse off borrowers will be bailing out the better off borrowers. In particular, it may be that an unlucky borrower gives her whole income to a partner to repay their loan, but defaults on her own loan. This is because the worse off borrowers require a larger transfer, which is thus less likely to be incentive compatible. This mechanism achieves the maximum possible repayment rate and therefore maximizes ex-ante expected utility.

This does not imply that a borrower with $y_j > r$ would ever default (i.e. be forced to choose between losing δV and δS. The reason is that all borrowers "above" her in the bail out chain also have $y > r$, so repay, then we obtain a group level repayment rate $\pi_i = \frac{m}{n}$. As this procedure is repeated for a sample of N groups, we can then estimate the overall repayment probability as the simple average.

The procedure in pseudo-code:

Group Lending without JL

1. Generate a $N \times n$ matrix of income realizations from F.
2. For each possible value of the interest rate r:
   (a) For each Y_i: compute the maximum level of contributions that each group member is willing to make to the common pool as $c_{ij} = max(y_{ij}, \delta S)$. This pot amounts to $C_{ij} = \sum_n c_{ij}$
   (b) Compute the redistributions required by members to ensure repayment as $t_{ij} = max(0, r - y_{ij} - c_{ij})$.
   (c) Order the required transfer in ascending order and redistribute the pot C_{ij} until it is exhausted.
   (d) Compute the group level repayment rate $\pi_i(r)$.
3. Given all the π_i, compute $\pi(r) = \frac{\sum_i \pi_i}{N}$.

Group Lending with Joint Liability

The simulation of this contract is more involved, since the relevant incentive constraint is $c_{ij} \leq \delta(V + S)$. This implies that in order to construct the repayment rate π, a number for the continuation value V is needed. V however, is itself a function of π.

The method proceeds as follows, for each possible value of r. First, we construct a set of possible candidates for $\pi(r)$, denoted $\tilde{\pi}$ [28], we calculate the associated $V(\tilde{\pi})$. Given these candidate $\tilde{V}$'s, the group fund C_{ij} is computed as follows. Each member is willing to contribute at most $c_{ij} = max(y_{ij}, \delta(\tilde{V} + S))$ toward repayment of the group's loan obligations. Explicit joint liability implies that the group will only repay when $C_j = \sum_n c_{ij} \geq nr$. Thus a group's repayment rate is $\pi_i = \mathcal{I}[C_j \geq nr] \in \{0, 1\}$. Taking the average we obtain the simulated repayment rate given $\hat{\pi}(V(\tilde{\pi}))$. In other words, taking as given a value for $V(\tilde{\pi})$, the implied repayment rate $\hat{pi}$ is computed. Then, the true π (and thus the true V) is found by solving for the fixed point $\pi = \hat{\pi}(V(\hat{\pi}))$. By iterating over r, we obtain the schedule $\pi(r)$ and the associated $V(\pi(r))$.

The procedure in pseudo code:

are making net positive contributions to the fund, which therefore has a positive "balance" when her turn comes

[28] These candidate π's exploit the monotonicity of the $\pi(r)$ schedule. The upper bound is given by the previous iteration for a higher r, while the lower bound is globally defined as $\frac{\rho}{(\delta \bar{R})}$

Group Lending with JL

1. Generate a $N \times n$ matrix of income realizations from F.

2. For each interest rate r:

   (a) Construct a set of candidates for $\tilde{\pi}(r)$.

   (b) For each $\tilde{\pi}(r)$:

   - For each Y_i: compute the maximum level of contributions that each group member is willing to make to the common pool as $c_{ij} = max(y_{ij}, \delta(S + V(\tilde{\pi})))$. This pot amounts to $C_{ij} = \sum_n c_{ij}$
   - The group defaults if $C_j = \sum_n c_{ij} < nr$
   - Compute the group level repayment rate $\hat{\pi}_i(\tilde{\pi})$.

3. Given all the $\hat{\pi}_i(V(\tilde{p}))$, compute $\hat{\pi}(V(\tilde{\pi}))$ as the average and find the fixed point π such that $\pi = \hat{\pi}(V(\pi))$.

3.C Simulation Results for Piecewise Returns

As discussed in the main text, there is no straightforward approach to simulate the model with the piecewise returns distribution. The problem is one of too many degrees of freedom. A sensible approach would be to vary the difference between the parameters p_h and p_m, as we saw in the main draft that for $p_h < p_m$, group lending with joint liability performs particularly bad. We can vary this difference, but still hold the sum $p_h + p_m = \bar{p}$ fixed, where $\bar{p} = 0.921$, as in de Quidt et al. (2012).[29]

We still have three parameters to tie down. Namely R_m, R_h and the mean return. There is no straightforward approach to tie down either of these parameters when varying the difference between p_h and p_m. This appendix will show the results from one pragmatic way. First, we tie down $R_m = \rho/p^2$. This condition is motivated by assumption 1 for the two player model. It implies that the medium return is high enough to repay a individual liability loan. Given this and the value of $\bar{R} = 1.6$, we compute R_h imposing the constraint that $p_h = p_m$. This thus gives us the value for R_h, when the difference between p_h and p_m is zero. Given these fixed values, we then simply vary the difference between p_h and p_m, holding everything else constant. This exercise thus maps somewhat into the table of the two-player model, where the model suggest that there is only an IL equilibrium for low S and only IJ equilibria for sufficiently high S. There is no EJ equilibrium in this case however. For $\Delta > 0$, the simple model would predict EJ lending for some range of parameter values. In the two-player model thus, the Δ is key. For groups with larger size, we would not expect this simple result to go through as now there are a lot more states of the

[29]Please refer to this paper for details on how this value was estimated using cross-sectional data from the MIX Market database

world. However, when plotting the simulation results as a function of the difference between p_h and p_m in figure 3.5, we do see that EJ performs better the larger $p_h - p_m$. However, this may simply be due to the fact that for higher p_h relative to p_m, the mean return in this case is changing as well.

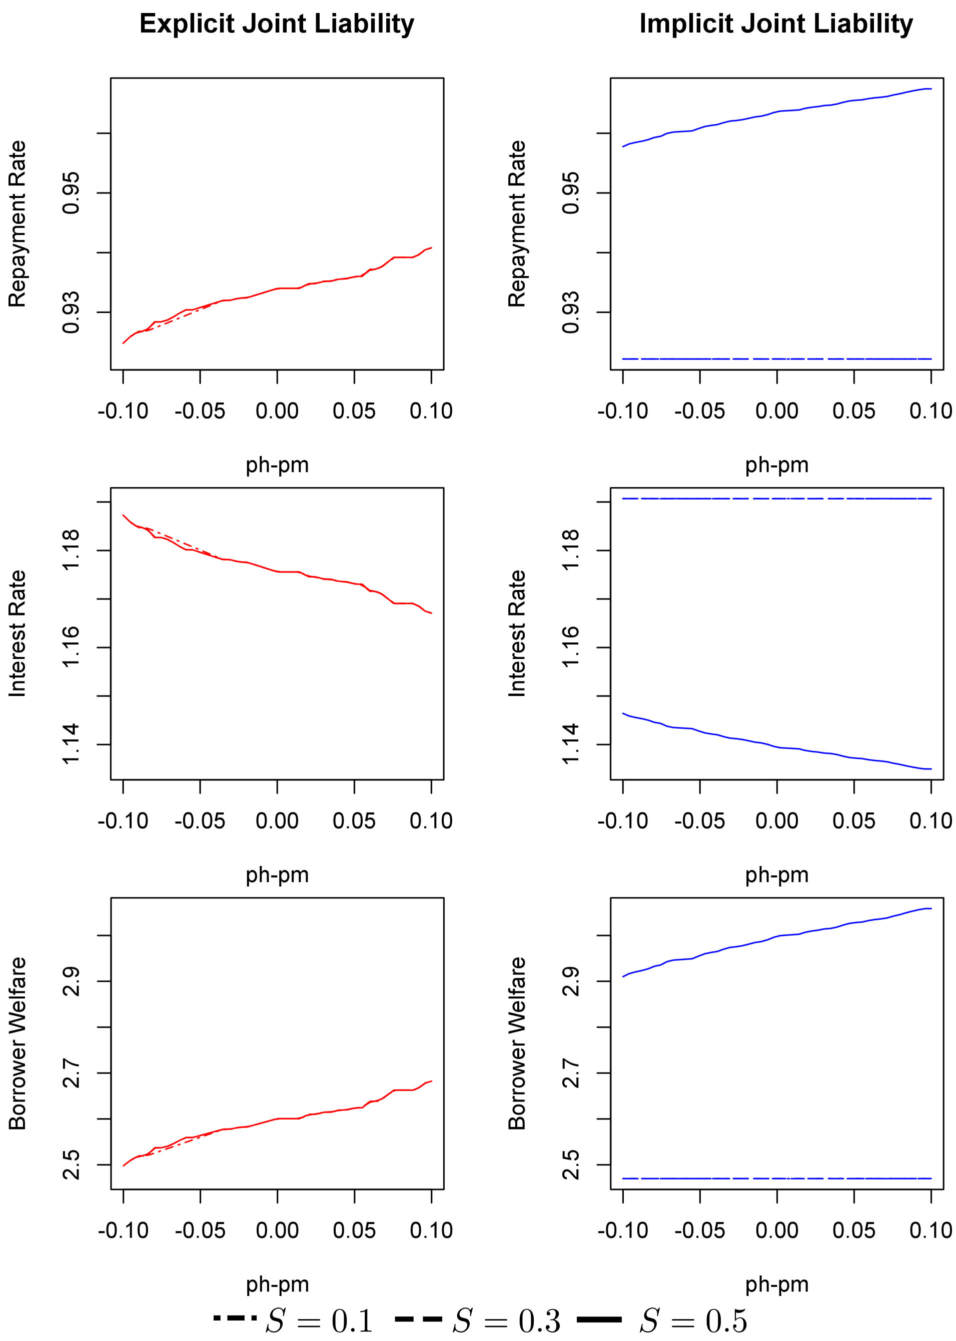

Figure 3.5: Simulation results for piecewise borrower returns distribution. Curves for explicit joint liability are drawn in red, and implicit joint liability in blue. Each figure plots the relevant object (repayment rate, interest rate and borrower welfare) for three levels of social capital, $S = 0.1, 0.3, 0.5$. The difference between p_h and p_m of individual borrower returns is varied on the horizontal axis of each figure.

Bibliography

Abadie, A. (2005, January). Semiparametric Difference-in-difference Estimators. *The Review of Economic Studies 72*(1), 1–19.

Abadie, A. and J. Gardeazabal (2008, January). Terrorism and the world economy. *European Economic Review 52*(1), 1–27.

Acemoglu, D., S. Johnson, and J. Robinson (2001). The Colonial Origins of Comparative Development: An Empirical Investigation. *American Economic Review 91*(5), 1369–1401.

Aggarwal, S. (2014). Do Rural Roads Create Pathways Out of Poverty? Evidence from India. *mimeo*.

Ahlin, C. and N. Jiang (2008, January). Can micro-credit bring development? *Journal of Development Economics 86*(1), 1–21.

Ahlin, C. and R. M. Townsend (2007, January). Using Repayment Data to Test Across Models of Joint Liability Lending. *The Economic Journal 117*(517), F11—-F51.

Akresh, R. and D. D. Walque (2008). Armed conflict and schooling: Evidence from the 1994 Rwandan genocide. *mimeo* (April).

Allen, T. (2012). Optimal (Partial) Group Liability in Microfinance Lending. *mimeo*.

Ambrus, A., E. Chaney, and I. Salitskiy (2011). Pirates of the Mediterranean: An Empirical Investigation of Bargaining with Transaction Costs. *mimeo*.

Anderson, C., L. Locker, and R. Nugent (2002, January). Microcredit, Social Capital, and Common Pool Resources. *World Development 30*(1), 95–105.

Anderson, J. and O. Bandiera (2006). Traders, Cops and Robbers. *Journal of International Economics 70*(1), 197–215.

Anderson, J. E. and E. van Wincoop (2003, March). Gravity with Gravitas: A Solution to the Border Puzzle. *American Economic Review 93*(1), 170–192.

Annan, J. and C. Blattman (2014). Can employment reduce lawlessness and rebellion? Experimental evidence from an agricultural Intervention in a fragile state. *mimeo*.

Armendáriz de Aghion, B. and J. Morduch (2010, May). *The Economics of Microfinance, Second Edition*. MIT Press Books. Cambridge, MA: The MIT Press.

Attanasio, O., B. Augsburg, R. De Haas, E. Fitzsimons, and H. Harmgart (2011). Group Lending or Individual Lending? Evidence from a Randomized Field Experiment in Mongolia. *EBRD Working Paper No. 136*.

Azam, M. (2011). The Impact of Indian Job Guarantee Scheme on Labor Market Outcomes: Evidence from a Natural Experiment. *IZA Discussion Paper 6548*.

Bai, J. (2009, February). Estimating Multiple Breaks One at a Time. *Econometric Theory 13*(03), 315.

Bai, J. and P. Perron (2003, January). Computation and analysis of multiple structural change models. *Journal of Applied Econometrics 18*(1), 1–22.

Baland, J.-M., R. Somanathan, and Z. Wahhaj (2010, January). Repayment incentives and the distribution of gains from group lending. *CDEDSE Working Paper* (192).

Bandiera, O. (2003). Land reform, the market for protection, and the origins of the Sicilian mafia: theory and evidence. *Journal of Law, Economics, and Organization* (1881).

Banerjee, A. V. (2012, September). Microcredit Under the Microscope: What Have We Learned in the Past Two Decades, and What Do We Need to Know? *Annual Review of Economics 5*(1).

Banerjee, K. and P. Saha (2010). The NREGA, the Maoists and the developmental woes of the Indian state. *Economic and Political Weekly xlv*(28), 42–48.

Bazzi, S. and C. Blattman (2014, October). Economic Shocks and Conflict: Evidence from Commodity Prices. *American Economic Journal: Macroeconomics 6*(4), 1–38.

Behar, A. and A. Venables (2011). Transport Costs and International Trade. In A. de Palma, R. Lindsey, E. Quinet, R. Vickerman, and E. Elgar (Eds.), *Handbook of Transport Economics*, Number January.

Bensassi, S. and I. Martinez-Zazosa (2010). How Costly is Modern Maritime Piracy for the International Community? *MPRA Paper* (27134).

Berg, E., S. Bhattacharyya, R. Durgam, and M. Ramachandra (2012). Can Rural Public Works Affect Agricultural Wages? Evidence from India. *CSAE Working Paper*.

Besley, T. and S. Coate (1992). Workfare versus welfare: Incentive arguments for work requirements in poverty-alleviation programs. *American Economic Review 82*(1), 249–261.

Besley, T. and S. Coate (1995). Group lending, repayment incentives and social collateral. *Journal of Development Economics 46*, 1–18.

Besley, T. and M. Ghatak (2010). Property Rights and Economic Development. In D. Rodrik and M. Rosenzweig (Eds.), *Handbook of Development Economics*, Volume V, Chapter 68.

Besley, T. and H. Mueller (2012, April). Estimating the Peace Dividend: The Impact of Violence on House Prices in Northern Ireland. *American Economic Review 102*(2), 810–833.

Besley, T., H. Mueller, and T. Fetzer (2014). The Welfare Cost of Lawlessness: Evidence from Somali Piracy. *forthcoming, Journal of the European Economics Association*.

Besley, T. and T. Persson (2008). The Incidence of Civil War: Theory and Evidence. *NBER Working Paper*.

Bhole, B. and S. Ogden (2010). Group lending and individual lending with strategic default. *Journal of Development Economics 91*(2), 348–363.

Blakeslee, D. and R. Fishman (2014). Weather Shocks, Crime, and Agriculture: Evidence from India. *mimeo*.

Blattman, C. and J. Annan (2010, November). The Consequences of Child Soldiering. *Review of Economics and Statistics 92*(4), 882–898.

Blattman, C., N. Fiala, and S. Martinez (2014, December). Generating Skilled Self-Employment in Developing Countries: Experimental Evidence from Uganda. *The Quarterly Journal of Economics 129*(2), 697–752.

Blattman, C. and E. Miguel (2010, March). Civil War. *Journal of Economic Literature 48*(1), 3–57.

Bolton, P. and D. S. Scharfstein (1990). A Theory of Predation Based on Agency Problems in Financial Contracting. *The American Economic Review 80*(1), 93–106.

Bowles, S. and A. Jayadev (2004). Guard Labor: An Essay in Honor of Pranab Bardhan. *mimeo*.

Buhaug, H. and G. Wischnath (2014, September). Rice or riots: On food production and conflict severity across India. *in Press, Political Geography*.

Burgess, R., O. Deschenes, D. Donaldson, and M. Greenstone (2011). Weather and Death in India. Technical report, MIT.

Burgess, R. and D. Donaldson (2010, May). Can Openness Mitigate the Effects of Weather Shocks? Evidence from India's Famine Era. *American Economic Review: Papers & Proceedings 100*(2), 449–453.

Burnicki, A. C., D. G. Brown, and P. Goovaerts (2007, May). Simulating error propagation in land-cover change analysis: The implications of temporal dependence. *Computers, Environment and Urban Systems 31*(3), 282–302.

Carpena, F., S. Cole, J. Shapiro, and B. Zia (2010). Liability Structure in Small-Scale Finance: Evidence from a Natural Experiment. *World Bank Policy Research Working Paper 5427*.

Chowdhury, P. R. (2005). Group-lending: Sequential financing, lender monitoring and joint liability. *Journal of Development Economics 77*(2), 415–439.

Ciccone, A. (2011, October). Economic Shocks and Civil Conflict: A Comment. *American Economic Journal: Applied Economics 3*(4), 215–227.

Coate, S. (2000). An Efficiency Approach to the Evaluation of Policy Changes. *Economic Journal 110*, 437–455.

Cole, S., X. Giné, J. Tobacman, P. Topalova, R. Townsend, and J. Vickery (2010, December). Barriers to household risk management: Evidence from India. *American Economic Journal: Applied Economics 5*(1), 104–135.

Cole, S., J. Tobacman, and P. Topalova (2008). Weather insurance: Managing risk through an innovative retail derivative. *mimeo*.

Collier, P. and A. Hoeffler (1998, October). On economic causes of civil war. *Oxford Economic Papers 50*(4), 563–573.

Collier, P. and A. Hoeffler (2004, June). Greed and grievance in civil war. *Oxford Economic Papers 56*(4), 563–595.

Collobert, R., J. Weston, and L. Bottou (2011). Natural language processing (almost) from scratch. *The Journal of Machine ... 12*, 2461–2505.

Conley, T. (1999, September). GMM estimation with cross sectional dependence. *Journal of Econometrics 92*(1), 1–45.

Cull, R., A. Demirgüç-kunt, and J. Morduch (2009). Microfinance meets the market. *Journal of Economic Perspectives 23*(1), 167–192.

Dasgupta, A., K. Gawande, and D. Kapur (2014). Anti-Poverty Programs Can Reduce Violence: India's Rural Employment Guarantee and Maoist Conflict. *mimeo*.

Dash, S. P. (2006). *Naxal Movement and State Power: With Special Reference of Orissa* (1st editio ed.). New Delhi: Sarup & Sons.

De Mel, S., D. McKenzie, and C. Woodruff (2008). Returns to Capital in Microenterprises: Evidence from a Field Experiment.

de Quidt, J., T. Fetzer, and M. Ghatak (2012). Group Lending Without Joint Liability. *mimeo*.

Dee D. P., Uppala S. M., Simmons A. J., Berrisford P., Poli P., Kobayashi S., Andrae U., Balmaseda M. A., Balsamo G., Bauer P., Bechtold P., Beljaars A. C., L. M. van de Berg, Bidlot J., Bormann N., Delsol C., Dragani R., Fuentes M., Geer A. J., Haimberger L., Healy S. B., Hersbach H., Hólm E. V., Isaksen L., Kållberg P., Köhler M., Matricardi M., McNally A. P., Monge-Sanz B.M., Morcrette J.J., Park B.K., Peubey C., de Rosnay P., Tavolato C., J. N. Thépaut, and F. Vitart (2011). The ERA-Interim reanalysis: configuration and performance of the data assimilation system. *Quarterly Journal of the Royal Meteorological Society 137*, 553–597.

Deshpande, R. S. and K. Shah (2010). *Agrarian Crisis and Farmer Suicides* (12 ed.). New Delhi: Sage Publications India Ltd.

Disdier, A. (2008). The Puzzling Persistence of the Distance Effect on Bilateral Trade. *The Review of Economics and Statistics 90*(1), 37–48.

Dixit, A. (2004). *Lawlessness and Economics: Alternative Modes of Governance*. New Jersey: Princeton University Press.

Donaldson, D. (2010). Railroads of the Raj: Estimating the Impact of Transportation Infrastructure. *NBER Working Paper 16487*.

Dube, O. and J. F. Vargas (2013, February). Commodity Price Shocks and Civil Conflict: Evidence from Colombia. *The Review of Economic Studies 80*(4), 1384–1421.

Duflo, E. and R. Pande (2007, May). Dams. *The Quarterly Journal of Economics 122*(2), 601–646.

Fearon, J. D. and D. D. Laitin (2003, March). Ethnicity, Insurgency, and Civil War. *American Political Science Review 97*(01), 75.

Feigenberg, B., E. Field, and R. Pande (2011). The Economic Returns to Social Interaction: Experimental Evidence from Microfinance. *mimeo*.

Ferrara, E. L. and M. Harari (2013). Conflict, Climate and Cells: A Disaggregated Analysis. *CEPR Discussion Paper 9277*.

Fetzer, T. R. (2013). Developing a Micro-Level Conflict Data Set for South Asia. *mimeo*.

Feyrer, J. (2009). Distance, Trade, and Income âĂŞ The 1967 to 1975 Closing of the Suez Canal as a Natural Experiment. *NBER Working Paper 15557*.

Fischer, G. and M. Ghatak (2010). Repayment Frequency in Microfinance Contracts with Present-Biased Borrowers. *mimeo*.

Gadenne, L. (2014). Optimal non-linear commodity taxation in developing countries: Theory and an application to India. *mimeo*.

Gawande, K., D. Kapur, and S. Satyanath (2012). Renewable Resource Shocks and Conflict in India's Maoist Belt. *mimeo*.

Geopolicity (2011). The Economics of Piracy: Pirate Ransoms and Livelihoods off the Coast of Somalia.

Ghatak, M. and T. W. Guinnane (1999). The economics of lending with joint liability: theory and practice. *Journal of Development Economics 60*(1), 195–228.

Giné, X. and D. S. Karlan (2011). Group versus Individual Liability : Short and Long Term Evidence from Philippine Microcredit Lending Groups. *mimeo*.

Glaeser, E. L. and A. Shleifer (2001). Not-for-profit entrepreneurs. *Journal of Public Economics 81*(1), 99–115.

Gomes, J. (2012). The Political Economy of the Maoist Conflict in India: An Empirical Analysis. *mimeo*.

Grossman, H. I. (2002, March). "Make us a king": anarchy, predation, and the state. *European Journal of Political Economy 18*(1), 31–46.

Guidolin, M. and E. La Ferrara (2007, December). Diamonds Are Forever, Wars Are Not: Is Conflict Bad for Private Firms? *American Economic Review 97*(5), 1978–1993.

Hamilton, J. (1989). A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica 57*(2), 357–384.

Hamilton, J. (1990). Analysis of Time Series Subject to Changes in Regime. *Journal of Econometrics 45*, 39–70.

Hansen, S. J. (2009). Piracy in the Greater Gulf of Aden: Myths, Misconception and Remedies. Technical report, Norwegian Institute for Urban and Regional Research Report.

Hegre, H. and N. Sambanis (2006, August). Sensitivity Analysis of Empirical Results on Civil War Onset. *Journal of Conflict Resolution 50*(4), 508–535.

Hoelscher, K., J. Miklian, and K. C. Vadlamannati (2012, June). Hearts and mines: A district-level analysis of the Maoist conflict in India. *International Area Studies Review 15*(2), 141–160.

Hornbeck, R. and P. Keskin (2014, January). The Historically Evolving Impact of the Ogallala Aquifer: Agricultural Adaptation to Groundwater and Drought âĂă. *American Economic Journal: Applied Economics 6*(1), 190–219.

Hsiang, S. M., M. Burke, and E. Miguel (2013, September). Quantifying the influence of climate on human conflict. *Science 341*(6151), 1–7.

Huffman, G. J., D. T. Bolvin, E. J. Nelkin, D. B. Wolff, R. F. Adler, G. Gu, Y. Hong, K. P. Bowman, and E. F. Stocker (2007, February). The TRMM Multisatellite Precipitation Analysis (TMPA): Quasi-Global, Multiyear, Combined-Sensor Precipitation Estimates at Fine Scales. *Journal of Hydrometeorology 8*(1), 38–55.

Hummels, D. (2007, August). Transportation Costs and International Trade in the Second Era of Globalization. *Journal of Economic Perspectives 21*(3), 131–154.

Imbert, C. and J. Papp (2015). Labor Market Effects of Social Programs: Evidence from India's Employment Guarantee. *forthcoming American Economic Journal: Applied Economics*.

Iyengar, R., J. Monten, and M. Hanson (2011). Building Peace: The Impact of Aid on the Labor Market for Insurgents. *NBER Working Paper 17297*.

Iyer, L., A. Mani, P. Mishra, and P. Topalova (2012, October). The Power of Political Voice: Women's Political Representation and Crime in India. *American Economic Journal: Applied Economics 4*(4), 165–193.

Iyer, L. and P. Topalova (2014). Poverty and Crime : Evidence from Rainfall and Trade Shocks in India. *Harvard Business School Working Paper 14067*.

Jain, S. and G. Mansuri (2003, January). A little at a time: the use of regularly scheduled repayments in microfinance programs*. *Journal of Development Economics*.

Jayachandran, S. (2006). Selling labor low: Wage responses to productivity shocks in developing countries. *Journal of Political Economy 114*(3), 538–575.

Jia, R. (2014, March). Weather Shocks, Sweet Potatoes and Peasant Revolts in Historical China. *The Economic Journal 124*(575), 92–118.

Kabeer, N. (2005). Is Microfinance a 'Magic Bullet' for Women's Empowerment? Analysis of Findings from South Asia. *Economic and Political Weekly 40*(44/45), 4709–4718.

Kalouptsidi, M. (2014, February). Time to Build and Fluctuations in Bulk Shipping. *American Economic Review 104*(2), 564–608.

Katyal, R., M. Sardana, and D. Satyanarayana (2001). *Estimates of District Domestic Product*. New Delhi: Socio-Economic Research Centre.

Khanna, G. and L. Zimmermann (2013). Guns and Butter ? Fighting Violence with the Promise of Development. *mimeo*.

Knack, S. and P. Keefer (1995, November). Institutions and Economic Performance: Cross-Country Tests using Alternative Institutional Measures. *Economics & Politics 7*(3), 207–227.

Kudamatsu, M., T. Persson, and D. Strömberg (2014). Weather and Infant Mortality in Africa. *mimeo*.

Kung, J. K.-s. and C. Ma (2014, November). Can cultural norms reduce conflicts? Confucianism and peasant rebellions in Qing China. *Journal of Development Economics 111*, 132–149.

Laffont, J. J. (2003, January). Collusion and group lending with adverse selection. *Journal of Development Economics*.

Laffont, J.-J. and P. Rey (2003). Moral Hazard, Collusion and Group Lending. *mimeo*.

Leon, G. (2009). Civil conflict and human capital accumulation: The long term effects of political violence in Peru. *mimeo*.

Leung, Y., Y. Ge, J. Ma, and J. Wang (2005). Measurement Errors and their Propagation in the Registration of Remote Sensing Images. In *Developments in Spatial Data Handling*, pp. 285–297. Springer Berlin Heidelberg.

Li, X.-H., Q. Zhang, and C.-Y. Xu (2012, March). Suitability of the TRMM satellite rainfalls in driving a distributed hydrological model for water balance computations in Xinjiang catchment, Poyang lake basin. *Journal of Hydrology 426-427*, 28–38.

Lilleor, H. and X. Giné (2005). Weather insurance in semi-arid India. *mimeo*.

Meza, D. D. and J. Gould (1992). The Social Efficiency of Private Decisions to Enforce Property Rights. *Journal of Political Economy 100*(3), 561–580.

Miguel, E., S. Satyanath, and E. Sergenti (2004, August). Economic Shocks and Civil Conflict: An Instrumental Variables Approach. *Journal of Political Economy 112*(4), 725–753.

Ministry of Rural Development (2008). The National Rural Employment Guarantee Act 2005 (NREGA)-Operational Guidelines. Technical report, Ministry of Rural Development, Government of India, New Delhi.

Ministry of Rural Development (2009). Joint Convergence Guideline. Technical report, Ministry of Rural Development.

Ministry of Rural Development (2012). *MGNREGA Sameeksha, An Anthology of Research Studies on the Mahatma Gandhi National Rural Employment Guarantee Act, 2005, 2006âĂŞ2012*. New Delhi: Orient BlackSwan.

Mirza, D. and T. Verdier (2008, June). International trade, security and transnational terrorism: Theory and a survey of empirics. *Journal of Comparative Economics 36*(2), 179–194.

Morgan, R. and D. Reiter (2013). How War Makes the State: Insurgency, External Threat, and Road Construction in India. *mimeo*.

Morrison, B. C. (2012). Grievance, Mobilisation and State Response: An examination of the Naxalite Insurgency in India. *Journal of Conflict Transformation & Security 2*(1).

Nichols, A. and R. Zeckhauser (1982). Targeting Transfers through Restrictions on Recipients. *American Economic Review: Papers & Proceedings 72*(2), 372–377.

Niehaus, P. and S. Sukhtankar (2013a, November). Corruption Dynamics: The Golden Goose Effect. *American Economic Journal: Economic Policy 5*(4), 230–269.

Niehaus, P. and S. Sukhtankar (2013b, August). The marginal rate of corruption in public programs: Evidence from India. *Journal of Public Economics 104*, 52–64.

North, D. C. (1968). Sources of Productivity Change in Ocean Shipping. *Journal of Political Economy 76*(5), 953–970.

Olken, B. and P. Barron (2009). The Simple Economics of Extortion: Evidence from Trucking in Aceh. *Journal of Political Economy 117*(3), 417–452.

Olson, M. (1993). Dictatorship, Democracy and Development. *American Political Science Review 87*(3), 567–576.

One Earth Future Foundation (2010). The Economic Costs of Maritime Piracy. *Working paper*.

One Earth Future Foundation (2011). The Economic Cost of Somali Piracy 2011. *Working Paper*.

Pandita, R. (2011). *Hello Bastar: The Untold Story of India's Maoist Movement*. New Delhi: Tranquebar Press.

Planning Commission (2005). Report of the Inter Ministry Task Group on: Investment Credit and Technical Support to Promote Self-employment in Agriculture, Horticulture, Afforestation, Dairying and Agro-processing. Technical Report January, Planning Commission, New Delhi.

Planning Commission (2008). Development Challenges in Extremist Affected Areas. Technical report, Planning Commission, New Delhi.

Prasad, P. (1987). Agrarian Violence in Bihar. *Economic and Political Weekly 22*(22), 847–852.

Radhakrishna, R. and S. Ray (Eds.) (2006). *Handbook of Poverty in India: Perspectives, Policies, and Programmes*. Oxford: Oxford University Press.

Rahman, H. and D. Sengupta (2007). Preliminary Comparison of Daily Rainfall from Satellites and Indian Gauge Data. *mimeo*.

Rai, A. and T. Sjöström (2010). Redesigning Microfinance. *Working paper*.

Rai, A. S. and T. Sjöström (2004). Is Grameen Lending Efficient? Repayment Incentives and Insurance in Village Economies. *Review of Economic Studies 71*(1), 217–234.

Ramana, P. V. (2007). *The Naxal Challenge: Causes, Linkages, and Policy Options* (1st editio ed.). Chennai: Pearson.

Rana, S. A. (2013). Terror financing or labour supply: the dominant channel of terrorist attacks in pakistan. *mimeo*, 1–51.

Ranjan, R. and V. Prasad (2012). Exploring Social Resilience in State Fragility: A Climate Change Perspective. In M. Hamza and C. Corendea (Eds.), *Climate Change and Fragile States: Rethinking Adaptation*, Number 16, Chapter 2. Bonn: United Nations University.

Rose-Ackerman, S. (2010). Corruption: Greed, Culture, and the State. *mimeo*.

Rosenzweig, M. and H. Binswanger (1993). Wealth, Weather Risk and the Composition and Profitability of Agricultural Investments. *Economic Journal 103*(416), 56–78.

Sanyal, P. (2009, January). From Credit to Collective Action: The Role of Microfinance in Promoting Women's Social Capital and Normative Influence. *American Sociological Review 74*(4), 529–550.

Sarsons, H. (2011). Rainfall and Conflict. *mimeo*.

Schneider, U., A. Becker, P. Finger, B. Meyer-Christoffer, Anja; Rudolf, and M. Ziese (2011). GPCC Full Data Reanalysis Version 6.0 at 0.5 Degree: Monthly Land-Surface Precipitation from Rain-Gauges built on GTS-based and Historic Data.

Sekhri, S. and A. Storeygard (2013). Dowry Deaths: Response to Weather Variability in India. *forthcoming, Journal of Development Economics*.

Shapiro, J. N., M. Callen, E. Berman, and J. Felter (2011). Do Working Men Rebel? Insurgency and Unemployment in Afghanistan, Iraq, and the Philippines. *Journal of Conflict Resolution 55*(4), 496–528.

Shortland, A. (2011). "Robin Hook" : The Developmental Effects of Somali Piracy. *DIW Discussion Papers*.

Shrivastava, A. (2014). Civil conflict with rising wages and increasing state capacity: The Maoist insurgency in India. *mimeo*.

Silva, J. M. C. S. (2011). poisson : Some convergence issues. *mimeo*, 1–8.

Silva, J. M. C. S. and S. Tenreyro (2006). The log of gravity. *The Review of Economics and Statistics 88*(November), 641–658.

Singh, P. (2013, January). Impact of Terrorism on Investment Decisions of Farmers: Evidence from the Punjab Insurgency. *Journal of Conflict Resolution*.

Smith, S. R., M. A. Bourassa, and M. Long (2011, July). Pirate attacks affect Indian Ocean climate research. *Eos, Transactions American Geophysical Union 92*(27), 225.

Spagnolo, G. (1999, January). Social relations and cooperation in organizations. *Journal of Economic Behavior & Organization 38*(1), 1–25.

Stopford, M. (2009). *Maritime Economics* (3 ed.). New York: Routledge.

UNCTAD (2009). Review of Maritime Transport. *Report by the UNCTAD secretariat*.

UNCTAD (2010). Review of Maritime Transport. *Report by the UNCTAD secretariat*.

Vanden Eynde, O. (2011). Targets of Violence: Evidence from India's Naxalite Conflict. *mimeo*.

Verma, S. (2011). Far Reaching Consequences of the Naxalite Problem in India. Technical Report July, Rakshak Foundation.

Voors, M. J., E. E. M. Nillesen, P. Verwimp, E. H. Bulte, R. Lensink, and D. P. V. Soest (2012, April). Violent Conflict and Behavior: A Field Experiment in Burundi. *American Economic Review 102*(2), 941–964.

World Bank (2013). The Pirates of Somalia: Ending the Treat, Rebuilding a Nation. Technical report, World Bank, Washington, DC.

Zimmermann, L. (2012). Labor market impacts of a large-scale public works program: evidence from the Indian Employment Guarantee Scheme. *mimeo*.