ABSTRACT

Title of dissertation: ESSAYS ON LABOR MARKETS
AND AGRICULTURE
IN DEVELOPING COUNTRIES

Fo Kodjo Dzinyefa Aflagah
Doctor of Philosophy, 2020

Dissertation directed by: Professor Jessica Goldberg
Department of Economics

This dissertation studies topics related to labor markets and agricultural cooperatives in developing countries.

In the first chapter, I study the Employment Tax Incentive (ETI), a South African wage subsidy to firms for newly hired young workers. This chapter provides a unique contribution to the literature on the effects of wage subsidies in developing countries. The ETI is one of only a few nationally implemented wage subsidies. Additionally, South Africa has the highest youth unemployment rate in the world. I use a difference-in-difference strategy based on age-eligibility and the start date of the policy to estimate its effects on youth employment outcomes in the first year of implementation. I find that even though the take-up rate of subsidies by firms was higher than expected, youth employment did not improve. I discuss the potential structural factors responsible for the seeming failure of this policy.

In the second chapter, I explore the importance of ethnic capital, a proxy for social capital, in determining labor market outcomes for South Africa's internal migrants. While the

role of social networks for the labor market outcomes of international migrants has been extensively studied, internal migrants have not received the same attention. Arguably, internal migration is an important phenomenon for developing countries. I look at how ethnic capital affects the employment probability and the occupation choice of South African internal migrants. I find that the employment probability of an internal migrant is driven by the education distribution of his ethnic group, even after accounting for local labor market conditions and unobserved ethnic group effects. I also find that occupation choice is related to the education distribution and the gender composition of the ethnic group at the destination.

In the third chapter, coauthored with Angelino Viceisza and Tanguy Bernard, we study how communication can be improved to solve coordination issues within agricultural co-operatives, using both a lab-in-the-field experiment (LFE) and a randomized controlled trial (RCT). Theory and conventional lab experiments suggest that non-binding pre-play communication can enhance coordination of interdependent agents in situations of strategic uncertainty, and that coordination should be easier to achieve in smaller groups. We test for these hypotheses in the context of Senegalese agricultural cooperatives which seek, but mostly fail, to jointly sell their outputs in order to secure higher unit prices. We show that revealing farmers' intended sales yielded enhanced coordination in larger groups where coordination is initially more difficult, and to higher income for small-scale farmers. We also show that participation in the LFE affected future behavior in the RCT.

ESSAYS ON LABOR MARKETS AND AGRICULTURE IN DEVELOPING COUNTRIES

by

Fo Kodjo Dzinyefa Aflagah

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2020

Advisory Committee:
Professor Jessica Goldberg, Chair
Professor Ethan Kaplan
Professor Guido Kuersteiner
Professor Sangeetha Madhavan
Professor Sergio Urzua

Dedication

To Mom.

Acknowledgments

I owe my gratitude to everyone whose support has made this thesis possible. I cannot name all but would like to thank some of you personally.

First, thanks to Professor Jessica Goldberg for her unconditional support throughout my time at UMD, from my admission up to this dissertation. Thanks to Professor Guido Kuersteiner for continuously supporting me in this dissertation. I would like to thank Professors Sergio Urzua and Ethan Kaplan for discussions that have helped shape the articles in my dissertation. Finally, thanks to Professor Sangeetha Madhavan for agreeing to serve on my thesis committee and sparing time to review this manuscript.

I would also like to extend my gratitude to the faculty and staff in the Economics Department of the University of Maryland for sharing knowledge, advice, and support throughout this year.

Thanks to all the students in the department for all these years together.

Thanks to Tanguy Bernard, Moussa Blimpo, Susie Godlonton, and Angelino Viceisza for countless hours of discussions.

Thanks to my family and friends (friends for life!) for the support, and for understanding all those unanswered calls and texts. Thanks to Marie Albertine for our discussions on the chapters in here.

Thanks Lou, for always being present and coping with me all these years!

Table of Contents

Dedication ii

Acknowledgements iii

Table of Contents iv

List of Tables vi

List of Figures viii

General Introduction 1

1 Failed Promises of a Wage Subsidy: Youth and South Africa's Employment Tax Incentive 5
- 1.1 Introduction 5
- 1.2 Wage subsidies in the literature 12
  - 1.2.1 Wage subsidies conceptually 12
  - 1.2.2 Wage subsidies in developing countries 14
- 1.3 Context 16
  - 1.3.1 South Africa's labor market and youth unemployment 16
  - 1.3.2 The Employment Tax Incentive 21
  - 1.3.3 Previous evaluations of wage subsidies in South Africa 23
- 1.4 Data and methodology 25
  - 1.4.1 Data: the Quarterly Labor Force Survey 25
  - 1.4.2 Identification strategy 29
  - 1.4.3 Suggestive evidence of parallel trends 34
- 1.5 Results 37
  - 1.5.1 Main results 37
  - 1.5.2 Robustness checks 45
  - 1.5.3 Heterogeneity of the effects 52
  - 1.5.4 Discussion: Unsuccessful scale-up? 55
- 1.6 Conclusion 59

2 Internal Migrants' Ethnic Capital and Labor Market Outcomes in South Africa 61
- 2.1 Literature 65
- 2.2 Methodology 70

2.3 Data . . . 77
2.4 Internal migration in South Africa . . . 79
2.5 Results . . . 87
2.5.1 Employment and ethnic capital . . . 87
2.5.2 Occupational choice and ethnic capital . . . 96
2.6 Conclusion . . . 100

3 Cheap Talk and Coordination in the Lab and in the Field: Collective Commercialization in Senegal 102
3.1 Introduction . . . 102
3.2 Strategic uncertainty in Senegal . . . 106
3.3 Framework, literature, and hypotheses . . . 108
3.3.1 Framework . . . 108
3.3.2 Related literature and hypotheses . . . 110
3.4 Study design . . . 112
3.4.1 LFEs . . . 113
3.4.2 RCT . . . 117
3.4.3 Sample description . . . 119
3.4.4 Empirical strategy . . . 120
3.5 Results . . . 123
3.5.1 Cheap talk and coordination . . . 123
3.5.2 Collective commercialization and farmers' income . . . 127
3.5.3 What information matters? . . . 130
3.5.4 Participation in games and day-to-day behavior . . . 133
3.6 Conclusion . . . 135

A Appendix for chapter 1 138

B Appendix for chapter 2 152

C Appendix for chapter 3 162

Bibliography 178

List of Tables

1.1 Placebo effect using 2013 first three quarters 37
1.2 Impact of the ETI on youth employment: main results 39
1.3 Impact of the ETI on youth employment: aggregate level regressions . . . 41
1.4 Effects of the ETI on labor force participation 42
1.5 Main effects of the ETI, comparing 29 and 31 years old over time 46

2.1 Race and internal migration incidence 80
2.2 Comparing migrants and non-migrants 81
2.3 Size of the language and province of origin groups 83
2.4 Occupation choice distribution 85
2.5 Average employment rate by language groups (recent migrants and established migrants) 86
2.6 Employment and ethnic group characteristics 88
2.7 Employment and ethnic group characteristics: IV estimations 95
2.8 Occupational choice and ethnic capital (marginal effects) 98

3.1 Intensive margin: treatment effect on quantities contributed to the group (correcting for MHT) 126
3.2 Effect on revenues from groundnut sales 129
3.3 Effect on quantity sold by treatment condition (correcting for MHT) . . . 131
3.4 Participation in LFEs and behavior in RCT 136

A.1 Cohorts and eligibility 141
A.2 Detail summary statistics of outcomes of interest and control variables . . 142
A.3 Effects on unemployment 143
A.4 Robustness checks 144
A.5 Testing the sensitivity of the results to the functional form of experience proxy 145
A.6 Heterogeneous effects by gender 146
A.7 Heterogeneous effects by race 147
A.8 Heterogeneous effects by level of education 148
A.9 Heterogeneous effects by industry 149
A.10 Effects and local labor market conditions 150

A.11 Impact of the ETI on youth employment: main results, sensitivity to alternative clustering . . . 151

B.1 Province of birth and internal migration incidence . . . 153
B.2 Language and internal migration incidence . . . 155
B.3 Province of birth and current province of residence . . . 155
B.4 Employment and ethnic capital, results at the ethnic group - municipality level . . . 156
B.5 Employment and ethnic capital, by specific and multiple language proficiency . . . 157
B.6 Employment and ethnic capital, excluding migrants moving to neighboring local municipalities or within same province . . . 158
B.7 Employment and ethnic capital: results for migrants aged 25-64 . . . 159
B.8 Employment and ethnic capital: results for the province of Gauteng . . . 160
B.9 Level of education and occupational choice . . . 161
B.10 First-stage of IV-estimates . . . 161

C.1 Sample distribution across RCT and LFEs . . . 162
C.2 Distribution of Baseline and Communication Coordination Games . . . 162
C.3 RCT treatments . . . 163
C.4 Balance tests with respect to the binary treatment - farmer level . . . 164
C.5 Balance tests with respect to the treatment conditions - farmer level . . . 165
C.6 Extensive margin: treatment effect on the probability of collective action . . . 166
C.7 PSP estimates: LFE cheap-talk effect in Table 3.1 - column 1 $(\beta = 0.66)$ . . . 167
C.8 PSP estimates: RCT cheap-talk effect in Table 3.1 - column 5 $(\beta = 0.99)$ . . . 167
C.9 Effect on collective commercialization - Tobit model . . . 168
C.10 Effect on collective commercialization - inverse hyperbolic sine transformation . . . 169
C.11 Intensive margin results for LFEs: testing sensitivity to farmer random effects . . . 170
C.12 Isolating the effect of group size from that of the threshold (LFEs) . . . 171
C.13 Robustness check: effect of the treatment on patience, generosity and risk . . . 172
C.14 Comparing large and small cooperatives . . . 173
C.15 Collective commercialization and aggregate intentions . . . 174
C.16 Collective commercialization and leader's intentions . . . 175
C.17 Comparing farmers attending games to farmers not attending, within LFE cooperatives . . . 176
C.18 Comparing farmers in LFE cooperatives to farmers in non-LFE cooperatives . . . 177

List of Figures

1.1 Amount of the ETI subsidy 23
1.2 Distribution of experience proxy 33
1.3 Testing parallel trends prior to the policy 35

3.1 LFEs – visual aid 114

A.1 Evolution of outcomes over time 138
A.2 Distribution of the proxy before the policy 139
A.3 Evolution of outcomes over time for women 140
A.4 Evolution of outcomes over time for men 141

B.1 Income distribution among migrants and non-migrants 152
B.2 Employment rates of migrants and non-migrants of same origin and language group 153
B.3 Occupational choices of migrants and non-migrants of same origin and language group 154

General Introduction

Development economics stresses the role of agriculture in the developing world, given the importance of the sector in the economies of developing countries. Sub-Saharan Africa, the poorest region in the world, relies on agriculture for 15% of its GDP, and about 52% of its jobs, although this importance has been decreasing over time. [1] The sector provided more than 60% of jobs and contributed to 20% of the GDP in the early 1990s. This shift in labor composition has been absorbed mainly by the tertiary sector, which has increased from about 45% to above 50% during the same period. This picture of the current state and the transitions in the economy of Sub-Sahara Africa is quite heterogeneous across countries, and motivates the importance of adopting different sets of policies to address economic questions in general, and labor market issues in particular.

The economy of Senegal, a West African country, has a structure and an evolution similar to the common trend across the continent. Agriculture contributes to 30% of current jobs, down from nearly 50% in the early 1990s, while the share of services in total employment has been increasing, from a little below 40% in the 1990s to 55% at the end of the 2010s. Additionally, informality is a prominent feature of the economy, with an estimated 90% of non-agricultural jobs being informal. South Africa is an example of a country with a different evolution, one that is in some ways typical of other Southern African countries

[1] All the statistics in this section come from the World Bank's World Development Indicators.

but in other ways includes labor market features that are linked to the country's unique history. Agriculture has always represented a lower share of the economy due to the effects and legacy of land ownership rules present during the apartheid period. Only 5% of individuals are currently employed in the agricultural sector, a proportion that was 11% in the early 1990s. The economy is dominated by the service sector, which currently contributes to 72% of jobs. Two main features of its labor market distinguish South Africa from other countries. The first is South Africa's large unemployment rate, estimated at 28% in 2018, one of the largest in the world; as a comparison, the average unemployment is 6.2% in Sub-Saharan Africa in the same year. [2] The second feature is the low share of the informal sector in South Africa, which accounts for about a third of non-agricultural jobs in the country.

In this dissertation, I study strategies that individuals and governments can take to increase job opportunities or to get better returns from existing activities, given the particularities of these two countries mentioned above. In South Africa, where jobs come mostly from the formal non-agricultural sector, I investigate whether a policy intervention, wage subsidies, and an informal institution, ethnic capital, can help individuals improve their labor market outcomes, in a context with important labor market imperfections. In Senegal, where agriculture still has a significant share in employment, I investigate whether an intervention through farmer cooperatives can help improve the welfare of members.

Young workers are particularly affected by the imperfections in South Africa's labor market. In 2018, 52% of job seekers aged between 15 and 24 were unemployed, making

[2] Unemployment is higher in the Southern African region, but still much lower than South Africa's. Across South Africa's neighboring countries, Namibia has unemployment rate currently at 20%, whereas this rate is 18% in Botswana.

South Africa the country with the highest youth unemployment rate in the world. [3]To address this particular issue, the government implemented the Employment Tax Incentive, a wage subsidy that reduces taxes paid by employers if they hire young workers. The policy offers a unique opportunity to contribute to our knowledge base on the effectiveness of wage subsidies implemented at scale, as most evidence from developing countries come from small-scale experiments, and because the subsidy represented a large share of the wages of covered workers. The first chapter of this dissertation assesses the effect of the policy and shows that the policy was not effective in improving labor market outcomes for young workers, despite the higher-than-expected take-up of the policy. The chapter discusses the potential reasons for the failure of this policy, in particular the fact that the policy did not consider the structural reasons for unemployment.

The second chapter looks at labor market outcomes for internal migrants in South Africa. Like international migrants, internal migrants move to new places in search of better opportunities, relying on their contacts to decide where to move to and to find opportunities once they arrive there. This aspect, however, has not been studied in the literature as much as international migration, despite the importance of the question. Internal migration is an important phenomenon in developing countries, and in particular, a major contributor to the increasing urbanization. Additionally, social networks and their use are prevalent in the context of developing countries. The second chapter proxies for social networks with ethnic groups. I follow Borjas (1992) by defining ethnic capital as the resources that an internal migrant might benefit from based on his ethnic group, and look at how the average characteristics of the ethnic group explain occupation choices and chances of

[3]The 15-24 age range is chosen to reflect the official ILO youth unemployment rate.

employment.

The third chapter of this dissertation is concerned with the agricultural sector in Senegal, and it investigates one way in which small farmers may get better access to markets: collective commercialization through their cooperatives. Researchers and policymakers have long considered cooperatives as a way by which farmers can get access to better markets by aggregating their output. Those cooperatives exist but try and mostly fail to aggregate their output to sell. The third chapter frames this question as a lack of coordination and investigates an intervention that can help address this failure. Using the case of peanut farmer cooperatives in Senegal and a series of lab-in-the-field and randomized controlled experiments, we test the effectiveness of "cheap talk" to enhance coordination and collective commercialization within these cooperatives. The chapter also highlights the interplay between "cheap talk" and group size in addressing coordination failure, a result that has not been explored previously in the literature.

Chapter 1: Failed Promises of a Wage Subsidy: Youth and South Africa's Employment Tax Incentive

1.1 Introduction

Youth unemployment is a global challenge. Around the world, young workers face unemployment rates two to four times those of the overall population. They struggle to transition into the labor market, even after investing in education to increase their chances of integration.[1] Addressing this challenge is a particularly important policy question for developing countries, given their relatively young populations and how the issue's significance can be important for economic growth.[2] High youth unemployment rates mean that a large share of the active population is not contributing to the growth of the economy or accumulating the experience that would in turn increase these workers' productivity. Human capital accumulated through schooling could also be depleted through long spells of unemployment. Finally, unemployment can produce negative externalities such as violence, crime, and social unrest, due to the frustration of those who are not able to get a

[1] Filmer and Fox (2014) note that transition into the labor market is a multifaceted issue, with some young individuals experiencing periods of economic activity while in school, many others transitioning via their parents' or relatives' business and farms, and some making it into the formal sector after experiencing long spells of unemployment.

[2] In 2014, the median individual was 18 years old in Africa, 7 years younger than in South Asia, the second-youngest region. The population in Africa is projected to stay young, with a projected median age of 25 in 2050. (Filmer and Fox, 2014).

job, conditions that further deter investment by businesses.[3] To address the issue, policymakers have considered interventions such as active labor market policies.[4] Economists have regained interest in the effectiveness of these policies due to their widespread use to mitigate the 2008 global recession; however, not enough is known about how these interventions affect employment, especially wage subsidies in developing countries (Card et al., 2015; Kluve et al., 2017; McKenzie, 2017).[5]

South Africa has the highest youth unemployment rate in the world. In the first quarter of 2018, 52% of South Africans aged 15 to 24 years were unemployed.[6] The youth unemployment rate is in part a consequence of the overall high unemployment rate in the country; in the first quarter of 2018, South Africa's overall unemployment rate was 27%.[7] South Africa's labor market in that regard remains an outlier, both among countries with the same level of development (especially Latin American countries) and among Sub-Saharan African countries.[8] Over the two decades following the end of apartheid in 1994, unemployment has risen continuously, except for a short period of strong growth in the 2000s that ended with the 2008 global recession. The situation is particularly severe for Black South Africans (30% of them are unemployed versus 11% of White South

[3]See Cramer (2010) for a discussion of the literature on youth unemployment and violence, with a focus on developing countries. Blattman and Ralston (2015) also discuss employment and social stability, emphasizing the role of employment programs.

[4]See the literature review for a discussion of active labor market policies.

[5]In the aftermath of the 2008 crisis, a report by the International Labor Organization and the World Bank revealed that almost all countries took some counter-cyclical measures, including 23 countries that had some wage subsidies, primarily in terms of reducing social contributions by employers. See http://ilo.org/wcmsp5/groups/public/---ed_emp/---emp_elm/documents/publication/wcms_186324.pdf

[6]The situation is particularly problematic given the size of South Africa's population. Based on World Bank statistics, South Africa has a population twice as large as any other country with a youth unemployment rate above 40 percent.

[7]Data are from the Quarterly Labour Force Survey Statistical release for the 2018$Q1$: http://www.statssa.gov.za/publications/P0211/P02111stQuarter2018.pdf

[8]By comparison, Nigeria, the largest economy in Sub-Saharan Africa, just preceding South Africa, had an unemployment rate of 19% in the third quarter of 2017.

Africans), which suggests that inequality in the labor market is a major contributor to income inequality in the country.[9]

South Africa has implemented the Employment Tax Incentive (ETI) to address its youth unemployment issue. It is a wage subsidy, initially proposed by Levinsohn (2007), which offers tax reductions to firms that hire young workers. The presence of a high youth unemployment, and a high retention in the formal sector were the rationale for this policy, as this could arguably serve as stepping stone into the labor market for young workers. [10] The government adopted the proposal and signed it into law in December 2013. The wage subsidy was phased in as a pilot program until 2016 and later extended due to higher than expected take-up by firms (National Treasury, 2016), although the extension was supposed to be conditional on effectiveness. To date, a few papers have attempted to analyze the ETI's impact on young workers' labor market outcomes (Ebrahim et al., 2017; Moeletsi, 2017; Ranchhod and Finn, 2016), but they have yet to produce conclusive evidence of the ETI's causal impact. In this paper, I investigate the extent to which the Employment Tax Incentive affected labor market outcomes of young workers.

By analyzing the ETI, I contribute to the literature on the impact of wage subsidies on youth employment in developing countries and provide one of only a handful evaluations of policies implemented at a national level. Similar policies have been evaluated in the literature, including the "Stage d'Initiation à la Vie Professionelle" (Initiation into the World of Work - SIVP) in Tunisia (Broecke, 2013), the post-2008 Employment Subsidy

[9] Income distribution in the country remains largely unequal, with one of the highest Gini coefficients in the world, and inequality, if anything, has worsened over time. See Sulla and Zikhali (2018) for a detailed report on inequality in South Africa.

[10] See section 1.3.1 for a further discussion of South Africa's labor market institutional features, and how a wage subsidy is a potential policy to address the high youth unemployment.

Program in Turkey (Balkan et al., 2016), the Youth Employment Subsidy in Chile (Bravo and Rau, 2013), and the First Employment Law in Colombia (Ariza and Cedano, 2017). The ETI is different from these policies mainly because it offers larger subsidies as a share of the wages paid; its value is as high as half the wages paid, whereas other subsidies studied amount to between 10 and 20 percent of wages. Furthermore, the ETI covers a broader share of workers, compared to the policy in Tunisia that targeted recent graduates or the policy in Chile that targeted poor young workers. It also covers a wider range of the income distribution, including earnings well above the median and the current national minimum wage. Finally, unlike Turkey's program that was first designed to counter the effects of a recession, South Africa implemented the ETI independently of the business cycle. There have been some evaluations of wage subsidies on youth employment in developing countries using small-scale randomized control trials (Galasso et al. (2004) in Argentina, Groh et al. (2016) in Jordan), including a pilot evaluation in the case of South Africa (Levinsohn et al., 2014).[11] There is a debate on how scaling up policies changes their effectiveness; as such additional analysis of at-scale policy would contribute to our knowledge on the subject when it comes to wage subsidies.

By studying the effect of the ETI, I also contribute to the literature on youth unemployment in South Africa. The alarming youth unemployment rate not only makes any potential policy of interest in itself, but analyzing the ETI can also help policymakers rethink that policy or inform future decisions. The importance of the structural features underlying South Africa's high unemployment rate casts doubt on the effectiveness of any policy

[11] An additional evaluation by de Mel et al. (2019) examines subsidies for small firms to hire additional workers, but without restriction on age for the workers to hire. The firms hired workers that were 31.5 years old on average.

that does not adequately address those features. Economists often mention four main features: minimum wage laws, skill mismatches between labor supply and labor demand, firing regulations, and unions (Banerjee et al., 2008).[12] Minimum wage laws in South Africa potentially deter employers from hiring, especially hiring young workers, as their productivity is likely below the statutory minimum.[13] The goal of a wage subsidy is to make these minimum wages less binding, by reducing the actual cost of labor. However, high firing costs are another feature in South Africa that could make firms reluctant to hire new young workers, even when a subsidy is offered, especially a temporary one. Firms perceive firing costs as very high, due to the legislation that imposes high hurdles on firing a worker and due to the unpredictable implementation of this legislation by labor courts (Bhorat et al., 2014). In fact, the initial policy proposal suggested a "no-questions-asked" dismissal period associated with the implementation of the ETI, but this suggestion never made it into the actual policy.

Finally, I contribute to the literature on wage subsidies in South Africa by providing reliable causal estimates of the impact of the ETI on young workers' labor market outcomes. In order for a firm to be able to claim the subsidy for of a given worker, the worker had to be recently hired (after October 2013) and aged 18 to 29 at the time of the claim. These criteria suggest using a difference-in-difference strategy to estimate the effects of the policy on the employment outcomes of young workers. I study individuals born between 1979 and 1988, who would have been between 26 and 35 in the first year of implementation of the policy. In this sample, the cohorts 1984-1988 are eligible for the subsidy, and

[12] See the section on South Africa's labor market for a more detailed discussion.

[13] This is a recurring point in evaluations of wage subsidies: firms report that they do not hire workers because of the restrictions on wages. See Galasso et al. (2004) and Groh et al. (2016).

I use the cohorts 1979-1983 as a control group. This is a departure from Ranchhod and Finn (2016), who uses the sample of all working age individuals for their estimation, but an approach that Moeletsi (2017) previously used. I improve on Moeletsi (2017) by controlling for experience, a key dimension along which eligible and ineligible cohorts differ. Intuitively, labor outcomes of young workers catch up with those of their older peers over time, as they accrue more experience; also, the relative importance of experience early in the life cycle has been documented in South Africa (Banerjee et al., 2008; Levinsohn, 2007). However, experience is not measured in the dataset available on workers' labor market outcomes. I address this by constructing a proxy for experience at an aggregate level. Finally, I use workers' information to analyze the ETI, instead of firms' data, as the design of the policy provides a criterion that allows estimation of causal effects on workers instead of firms. Ebrahim et al. (2017) studied the effect of the ETI using firms' administrative data. There are reasons to be cautious about the endogenous nature of firms' decision to claim the subsidy, even when a conditional difference-in-difference is used. For instance, firms that take up the subsidy likely anticipate that they will grow. The results by Ebrahim et al. (2017) are to be interpreted with caution because of this endogenous decision.

My results show that overall, the ETI did not lead to a significant increase in employment for young individuals.[14] Given the precision of my estimates, I can rule out improvement in overall employment in the order of 1 percentage point.[15] I find that men seem to have been more affected, as the estimates are consistent with larger effects than those

[14]Although these results are similar to Ranchhod and Finn (2016), it is worth stressing that I do find this result on a sample that is optimal for identifying the causal effects of the policy.

[15]These results rule out the effects found by Moeletsi (2017), suggesting the importance of controlling for experience when estimating the effect of the policy.

for women. Indeed, if anything, my results imply that the policy led to *decreased* employment for women, but for reasons that I discuss later, this should be interpreted with caution. I find no differential effects with respect to race or education, two other important dimensions that affect outcomes on South Africa's labor market (Banerjee et al., 2008; Kingdon and Knight, 2004). The evidence in this paper finally suggests that there is no heterogeneous effect with respect to industrial sectors or local labor market conditions. Taken together, these results imply that reducing labor costs may not be sufficient to induce firms to hire young workers. I argue that other constraints must be addressed in addition to reducing labor costs in order to improve employment rates for young workers. One such constraint is firing regulations in the country. [16] When firms make hiring decisions in a dynamic setting, with a wage subsidy that lasts at most two years, they take into account the probability that they might hire a worker for whom the costs of firing might exceed the value of the wage subsidy, together with the forgone profit. If those costs exceed the benefit from the wage subsidy, then firms would be unwilling to hire young workers. This reasoning is in accordance with the literature as McKenzie (2017) noted that wage subsidies tend to be more successful in settings where firms are not required to register workers, suggesting that labor regulations are an important barrier to the success of wage subsidies.[17]

The remainder of the paper is structured as follows. First, I discuss the literature about

[16]In the present analysis, I have not been able to exhibit which of these constraints are responsible for the absence of positive effects of the policy. The choice to present how firing regulations may have affected the effectiveness of the policy is purely arbitrary. In section 1.3.1, I discuss how the different constraints could potentially undermine the effectiveness of a wage subsidy.

[17]A related paper is Hardy and McCasland (2017) who studies the effects of a placement service for apprentices in Ghana. Among firms who expressed interest in hiring apprentices with no entry fees unlike what is regularly done, firms assigned no apprentice did not subsequently hire any worker. The paper cites frictions in screening apprentices' ability or skills as a barrier for these firms young apprentices.

wage subsidies, with a focus on empirical evidence in developing countries. Second, I present the main features of South Africa's labor market. Third, I explain the policy, and I present the data and my empirical strategy. Fourth, I present and discuss the results from my estimation. Finally, I conclude by discussing the lessons learned from evaluating this policy.

1.2 Wage subsidies in the literature

Wage subsidies are one type of active labor market policies, a category of policies that intervene in the labor market to correct imperfections, in order to improve employment prospects of some or all workers.[18] They are in contrast to passive labor market policies that provide income replacement to individuals that have lost or do not have a job.[19] [20]

1.2.1 Wage subsidies conceptually

Theoretical discussions of wage subsidies date back to Kaldor (1936). As mentioned in Katz (1996), the "basic idea behind (employer-side) wage subsidies is to reduce the costs to firms of employing the targeted group of workers thereby stimulating demand for these workers and raising their employment rates and earnings." This shifts out the demand for the labor of the targeted group, and the effects on employment depend on labor demand and labor supply elasticities. However, that assumes that markets determine equilibrium

[18]See Betcherman et al. (2004) and Card et al. (2015) for general meta-analyses, Kluve et al. (2017) for a survey of policies targeted at youth, and McKenzie (2017) for a survey of studies in developing countries.

[19]These include, for instance, welfare benefits and unemployment insurance.

[20]Besides wage subsidy programs, other active labor market policies are: (1) employment services that focus on producing better matches between jobs and job seekers, (2) training programs in specific skills that increase potential workers' employment prospects, (3) public work programs, by which the government directly provides short term jobs, and (4) micro-enterprise or self-employment development, which is assistance (either financial or technical) to help individuals to develop their own businesses.

wages. If wages are determined otherwise, it is not profitable for firms to hire workers whose marginal productivity is under the prevailing wage (e.g., minimum wage), but wage subsidies can induce firms to hire such workers by offsetting or reducing the perceived productivity gap (Wolff and Stephan, 2013).

Theoretical discussions highlight implementation issues and have long recognized the trade-off between general and targeted subsidies. A well-designed wage subsidy will minimize the "windfall" that arises when firms merely receive money for workers that they would have hired anyway. This consideration favors targeted subsidies over general ones. For this reason, Layard and Nickell (1980) suggested subsidizing "marginal employment" (subsidies conditional on increasing employment at the firm level), but recognized the difficulty of accurately identifying marginal employment. Katz (1996) made the case for subsidies administrated as tax credits, in order to ease the burden of administration. However, targeted subsidies may come at the expense of distortions in labor allocation or stigmatization of targeted workers. Targeted workers may be hired at the expense of other similar workers (substitution effects) or other workers that are fired (displacement effects).

Besides the incentive created for firms to hire targeted workers, wage subsidies can also be justified by their long run effects, as accumulated experience raises workers' productivity. Advocates of wage subsidies view this as a justification despite concerns of substitution among eligible and ineligible workers (Bell et al., 1999). This also hints at the fact that wage subsidies can be more successful if they are combined with other activities that enhance productivity, such as on-the-job training (Katz, 1996).

1.2.2 Wage subsidies in developing countries

McKenzie (2017) reviewed experimental evidence on the effectiveness of active labor market policies in developing countries. It follows from the discussion of experimental evaluations of wage subsidies (de Mel et al., 2019; Galasso et al., 2004; Groh et al., 2016; Levinsohn et al., 2014) that they have a limited overall impact on employment. However, they can be useful as a temporary policy tool for the creation of short-term jobs. One useful insight that potentially applies to the case of South Africa is that covering only registered firms, or requiring eligible firms to register their workers, seems to lower the effect. Groh et al. (2016) found a 38 percentage point increase in employment when firms are not required to register workers hired through the wage subsidy, whereas Galasso et al. (2004) found a 1.7 percentage point increase in employment when registration is required and firms face a penalty if they fire the worker after the end of the subsidy period. This suggests that wage subsidies are unlikely to be effective if firms have conditions on hiring, and points especially to the fact that Levinsohn's (2007) initial proposal suggested a "no-questions-asked" dismissal period that did not make it to the actual policy implementation. The effect in Groh et al. (2016) is however not long-lived. [21]

Evidence from nationally implemented policies comes in part from measures taken by countries to counter the effects of the 2008 global financial crisis. In Turkey, these programs were aimed at firms hiring new workers who were previously unemployed (Balkan et al., 2016) whereas in Mexico, the policy aimed to prevent firms from firing workers

[21] One aspect that is missing in the literature is the effect of wage subsidies coupled with other interventions, such as training for the workers. Groh et al. (2016) offered soft skills training in their context, but this did not lead to differential effects.

(Bruhn, 2016). In both studies, the policies increase employment, showing that wage subsidies can be effective countercyclical measures. In particular, Balkan et al. (2016) show stronger effect in regions with higher unemployment rates, and evidence from Bruhn (2016) shows that firms could use the subsidies to hire enough workers to return to their pre-crisis levels of employment.

Nationally implemented wage subsidies have different targeting. The policies in Chile (Bravo and Rau, 2013) and Columbia (Ariza and Cedano, 2017) targeted young low wage earners. The wage subsidy in Tunisia (Broecke, 2013) targeted recent college graduates. Unlike these studies, my paper looks at a policy that has a broader eligibility criterion, as all workers aged 18-29 were eligible. The post-2008 Employment Subsidy Program in Turkey analyzed in Balkan et al. (2016) had even broader coverage, as firms could claim subsidies for all women above 18 and for men aged 18-29. Wage subsidies sometimes have targeting based only on firms. The policy in Turkey analyzed in Betcherman et al. (2010) targeted the poorest provinces in the country while the one in Mexico (Bruhn, 2016) was available only for firms producing durable goods (e.g., firms in manufacturing and construction industries).

Among the previous studies, two have investigated whether their positive employment effects were driven by substitution or displacement effects, i.e. firms hiring workers, but not creating new jobs overall in the economy. Betcherman et al. (2010) conduct an indirect test examining economic activity by using energy consumption, whereas Bruhn (2016) compared hiring patterns by eligible and ineligible industries. Bruhn (2016) found no effects of substitution; although one should bear in mind that these results are estimated in the aftermath of a recession, implying a large of pool of unemployed individuals to hire from. I discuss later the potential implications of displacement or substitution effects given my results.

1.3 Context

1.3.1 South Africa's labor market and youth unemployment

Economists point to four features as the main reasons for the high unemployment rate in South Africa (Banerjee et al., 2008): minimum wage laws, skill mismatch, labor regulations regarding firing, and unions.

South Africa adopted a national minimum wage in 2018, and it has been in effect since January 2019. The level was set at ZAR 3500 a month, which was close to the median monthly earnings at the time.[22] Before 2019, minimum wages were defined at the sectoral level and negotiated between firms and workers' unions at the district council level. Studies of revisions to sectoral minimum wages found mixed results of these changes.[23] Dinkelman and Ranchhod (2012) found that the creation of a minimum wage for domestic workers who were not previously covered by such laws did not lead to a decrease in employment, whereas Bhorat et al. (2014) found that the institution of a minimum wage for agricultural workers led to a decrease in employment for these workers. The high unemployment in South Africa rate is consistent with a binding minimum wage that is higher than the marginal productivity of unemployed workers (Banerjee et al., 2008). Young workers are likely less productive than more experienced and likely older ones, because

[22] ZAR = South African Rand. Over the period that the policy covered, the exchange rate fluctuated between 10 to 15 rands for a dollar.

[23] The lack of consensus on the employment effects of minimum wage is not specific to South Africa. See for instance Card and Krueger (1994) and Neumark et al. (2014) for evidence in the US.

on-the-job training and experience increase productivity. As such, minimum wages may have a bigger impact on the employment of young workers. [24]

A skill mismatch between labor demand and labor supply could explain both high unemployment rates and the rise in unemployment in the decade following the end of apartheid. Labor supply has increased, primarily through the addition of unskilled workers, at the same time that the economy became more skill biased (Banerjee et al., 2008). There are two sides of the nexus between skills and employment outcomes. The first relates to the level of education. There are high returns to completing secondary school, but many youths do not finish their secondary education. The second relates to the skills that education indeed confers to workers. Firms complain that education qualifications convey little information on workers' skills. Indeed, Carranza et al. (2019) show that skill certification increases youth employment, by removing information frictions between workers and employers regarding skills, whereas Abel et al. (2019) show that a reference letter from previous employers increase callbacks and employment for women, although the latter findings likely confound information frictions about both skills and experience.

The mechanism behind the link between labor regulations and unemployment in South Africa is not clear, especially for firing regulations. The regulations appear to be no different from those in Latin American or European countries that have stringent labor regulations but much lower unemployment rates. However, firms cite high separation costs as a deterrent for hiring decisions. Benjamin et al. (2010) suggest that labor courts interpret the law inconsistently, raising the uncertainty of the outcome when a dismissal is

[24] Neumark and Wascher (2004) provide cross-country evidence that minimum wages decrease youth employment, with a larger effect when there is no specific minimum wage for workers and when labor standards and union coverage are higher. Gorry (2013) further provides theoretical evidence that minimum wage interacts with the ability of a young worker to gain experience, leading to larger effects.

challenged as unfair. Bertrand and Crepon (2019) further show that providing information on labor regulations to firms make these firms hire substantially more workers.[25] Because there is no age restriction on these regulations, they are likely to be more binding for firms when recruiting young workers, as these workers' productivity is likely lower on average and has greater variability, while their potential tenure is longer.[26]

The presence and strength of unions in the country reinforce the role of minimum wage and labor regulations as constraints for job creation. Labor regulations explicitly grant each worker the right to be represented by a union member, both during internal discussions regarding dismissals and during court hearings. With unions' support, workers have better representation in courts and are more likely to have favorable outcomes during trials. This likely raises firms' costs of dismissal. Prior to the institution of a national minimum wage, unions were also responsible for collective bargaining that sets wages at sectoral and district council levels.

Given its high unemployment rate, it is surprising that informal employment in South Africa is low, especially compared to other countries in Sub-Saharan Africa. The share of workers in informal employment is 34 percent of the total employment in South Africa, compared to 72 percent in Africa and 53 percent in Latin American and Caribbean countries (ILO, 2018). Plausible explanations include barriers to entry in terms of credit con-

[25] Bertrand and Crepon (2019) report effects 6 months after their intervention. One possible explanation of their results is that providing information about regulations to firms did not make these firms create jobs, but made them more confident in their hiring decisions, and consequently make them hire at a quicker pace than without the information. As such, the effects found could be driven by quicker hires instead of job creation.

[26] Firms might be able to identify some of the older workers who have high ability or high productivity because of their experience, as these workers are more likely to get and mostly keep jobs, which would show in their experience, for example on their resume. This is not the case for young workers, as they usually do not have the experience. With this greater uncertainty about young workers' productivity, there are higher risks that firms get a worker with a low ability among young workers and that this worker might be difficult to fire later, due to the regulations.

straints, lack of skills and a lack of entrepreneurial culture due to the legacy of apartheid (Kingdon and Knight, 2004). These constraints also account, to some extent, for why firms in the formal sector do not hire more. Magruder (2012) further makes the case that regulations in the formal sector spread to the informal sector; regarding wage settings (recall the discussions above), they could raise wage expectations overall which negatively impacts employment creation by small firms, thus possibly firms in the informal sector. Young workers are usually disproportionately represented in the informal sector (Filmer and Fox, 2014), so an unusually small informal sector may contribute to high youth unemployment.

The high unemployment rate in South Africa has often raised questions about how unemployed individuals meet their basic needs without labor income. One hypothesis is that unemployed workers may rely on other sources of income and consequently that part of unemployment might be voluntary. Transfers from social programs, especially the old age pension, are one such source that is often mentioned. It is not clear however how the presence of pensioners in a household affects the labor supply of working age members. Studies that claim that the old age pension reduces labor supply (e.g., Bertrand et al., 2003) ignore the effect on migration, as the pension seems to provide a relaxation of financial constraints on migration for rural households (Ardington et al., 2009; Posel et al., 2006). However, Abel (2019) suggests that the results may not generalize to the whole country. These mixed results cast doubt on the hypothesis that income from social programs alters labor supply behavior and leads to voluntary unemployment.

These features of South Africa's labor market call into question how likely it is that a wage subsidy will increase young workers' employment. Minimum wages create a wedge or gap between marginal productivity and wages, which is likely larger for young workers since they are less experienced and less productive. A wage subsidy lowers the costs of labor, in this case of young workers' labor. If we first ignore the dynamic aspect of their decision-making, the wage subsidy may induce firms to hire more young workers than otherwise, provided that the subsidy compensates for the gap between marginal productivity and the minimum wage or other wage paid. Skills mismatches contribute to the productivity gap; the lower the skills of young workers, the higher the wage subsidy needs to be to induce firms to hire young workers. If there are skills that young workers can acquire by on-the-job training, the wage subsidy then acts as a compensation for the employer's costs of training the worker. Firms without the subsidy would not have voluntarily incurred these costs, especially since there is no way to guarantee that the workers will stay at the same firm once they are trained. [27]

Labor regulations highlight the importance of considering the dynamics of firms' decisions and their uncertainty regarding young workers' productivity. [28] In South Africa, firing regulations imply that firms are uncertain that they could fire a worker once they have hired him or her, or that they are uncertain of the cost at which they can fire the worker. In the case of a permanent subsidy, these regulations would have little influence on whether firms hire young workers. When the subsidy is temporary, however, the firm

[27] In theory, a solution to the costs of training faced by the employer (together with the risk that the worker leaves once he or she is trained) is that the firms pay a lower wage at the beginning of the worker's tenure, and then increase the wage with tenure. However, wage floors in terms of a minimum wage could make this impossible.

[28] The fact that firms are more uncertain of young workers productivity that they are of older ones plays into how minimum wages affect the effectiveness of a wage subsidy. However, if there were no regulation preventing form firing workers, firms can experiment with a wage subsidy, assured that they can fire the worker once he or she is hired and the firm learns about his or her productivity. The uncertainty about young worker's productivity is thus more relevant in relation with labor regulations.

needs to account for the possibility that it gets stuck with a worker that would be difficult to fire; the productivity of the worker may well be above the actual costs of labor for the firm during the subsidy period, but not afterwards. The presence of unions would tend to reinforce the point on labor regulations, as unions represent workers in courts during unfair firing trials. In the end, the decision to hire workers targeted by the subsidy will depend on whether the subsidy offsets the previously perceived productivity gap, accounting for potential on-the-job training, but labor regulations make firms less likely to hire despite the subsidy, especially given the uncertainty of young workers' productivity.

1.3.2 The Employment Tax Incentive

Levinsohn (2007) first proposed a wage subsidy to tackle youth unemployment in South Africa.[29] The wage subsidy is justified by the fact that young workers face high unemployment rates, but those who experience smooth transitions to work remain employed in the future. The original proposal argued for an individual account for young job seekers that they could use with any prospective employer. This framing seems to present youth unemployment as a low supply of labor issue instead of a low demand for labor. However, it is firms that need to be incentivized since they do not create jobs. In the end, the subsidy was implemented as a reduction in tax liabilities for employers. The initial proposal also included a probationary period during which a "no-questions-asked" dismissal policy would be in place, but this provision was not adopted.

[29] The other policy suggested is a reform on immigration policy that would encourage an influx of educated migrants especially from other African countries. The reason for such a policy is that skilled and unskilled labor are complement rather than substitutes. An increase in the level of education of the population propelled by an immigration reform is thus likely to increase employment prospects for low skilled labor.

The proposal was enacted into law as the Employment Tax Incentive Act which was signed in December 2013. The bill made it clear from the onset that the wage subsidy was a temporary policy, active from 2014 to 2016, that would be subject to a review and considered for a non-guaranteed extension. Specifically, policymakers expected to spend ZAR 5 billion over 3 years to create 178,000 new jobs for youths at a cost of approximately ZAR 28,000 per job, while subsidizing 423,000 jobs.[30] The policy has since been extended twice, in 2016 and in 2018, and is still in implementation until 2029. The decision to extend was based primarily on firms' higher-than-expected take-up of the subsidies. During the first two years of implementation, more than ZAR 6.3 billion was claimed, with an estimated 650,000 jobs claimed in the fiscal year 2014/2015 alone (National Treasury, 2016).

The wage subsidy has eligibility criteria on both the firm and the worker side. The firm needs to be registered for Pay-As-You-Earn (PAYE) for employee income tax purposes. Firms in the public sector are not eligible. Firms can also only claim deductions if they do not owe the South African Revenue Service any money. The bill requires that no older worker is displaced as the consequence of the young worker being hired, upon penalty of ZAR 30,000 and the possibility of losing the firm's eligibility to claim the incentive in the future. Finally, the employee's wage needs to be in accordance with minimum wage regulations. For their part, employees need to be between 18 and 29 years old, for the firm to claim the subsidy on their behalf. The other important restriction regarding the employee is that he or she may not be related to the employer.

[30] As a comparison, the government social protection programs amounted to ZAR 130 billion for the fiscal year 2013/2014. These programs primarily include the old age pension and the child support grant.

The employer can claim the wage subsidy as a reduction in taxes owed the National Treasury.[31] The subsidized amount per eligible worker hired depends on the wage he or she is paid, as shown in Figure 1.1. The amounts shown correspond to the subsidies for full-time workers. For part-time workers, the subsidy is proportional to the earnings that would have been paid if the worker was working full time, at the same wage rate. The subsidies can be claimed by an employer for the same worker for up to two years. The value of the subsidy falls by half in the second year it is claimed by the same firm, for the same worker.

Figure 1.1: Amount of the ETI subsidy

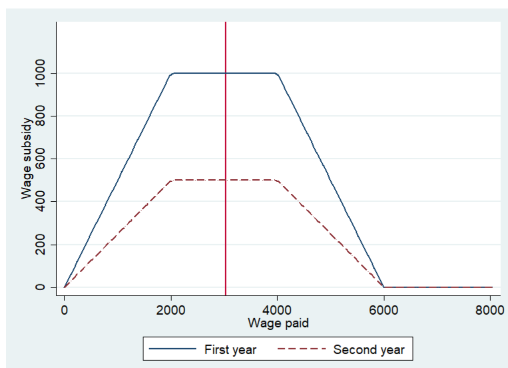

Notes: The vertical line denotes the median monthly earnings the year before the ETI was implemented.

1.3.3 Previous evaluations of wage subsidies in South Africa

Prior to the implementation of the Employment Tax Incentive, Levinsohn and Pugatch (2014) used a search model estimated using the Cape Area Panel Study data to perform

[31] Employers are subject to monthly withholding, including income tax that they pay to the South Africa Revenue Service. While paying the amount that they owe, firms deduct the subsidy and pay only the net amount due for withholding for all their workers. In case their total withholding is below the total subsidy owed, the excess is paid back to the employer as part of the reconciliation that employers must submit twice a year.

a prospective analysis of the impact of the policy. The study focuses on how the search behavior of workers responds to the subsidy, but their model does not incorporate firms' decisions, even though the statutory incidence of the subsidy is on firms. Levinsohn and Pugatch (2014) show that a wage subsidy can increase reservation wages, although when combined with the increase of the value of searching for a job, the net result is an increase in youth employment. They estimated that a ZAR 1,000 subsidy would decrease youth unemployment by 12 percentage points.

Levinsohn et al. (2014) later conducted an experiment piloting a wage subsidy voucher for young workers. A sample of 4000 youths aged 20-24 at baseline was divided into a wage voucher group and a control group. Follow-up surveys show that individuals who were offered the voucher were 7 percentage points more likely to be employed one year later (the period during which the voucher was redeemable) and 10 percentage points more likely to be employed two years later. Attrition is important, however, and may be driving the medium term effects. Besides this point, the authors noted that their results may be driven by the control groups turning down more job offers, especially when they are in households that have other employed individuals.

After the ETI was implemented, Ranchhod and Finn (2016) carried out a short term evaluation using the Quarterly Labor Force Survey. Using a difference-in-difference estimation strategy, they found no effects in the first year of policy implementation. Moeletsi (2017) later identified a flaw in their definition of the treatment group. Ranchhod and Finn (2016) defined the treatment group by current age (the below 30 group), rather than by age at the time of implementation. Moeletsi corrected this by defining the treatment group in term of fixed cohorts (born after 1985). Moeletsi's approach is an improvement, but it fails to account for workers' experience, which is likely to bias his estimates up, since young workers will catch up with the older ones over time as they gain more experience.

Ebrahim et al. (2017) use firms' administrative data to compare firms that claimed the tax incentive to firms that did not, before and during the implementation of the policy. They use a combination of matching and difference-in-difference methods. Their results indicate that, on average, firms that claimed the incentive hired more youths than to firms that did not. However, they found that hiring of non beneficiaries (workers aged 30 or above) increased by about the same amount. This most likely means that they capture the fact that firms that took up the subsidy would have hired workers even without the ETI. Using administrative records, they compute the deadweight loss due to subsidized jobs that would have been created anyway, given their estimates. Their results indicate that only 8% of claims can be attributed to newly created jobs, a very small share compared to similar policies in other countries.

1.4 Data and methodology

1.4.1 Data: the Quarterly Labor Force Survey

The Quarterly Labor Force Survey (QLFS) provides the best available information on South African labor market outcomes at the worker level, given its frequency and its national scope.[32] Since 2008, the QLFS has collected information from a nationally representative sample, at a quarterly frequency. It is a rotating panel in which one fourth

[32]Other available datasets are the National Income Dynamics Survey (NIDS), a panel survey conducted every two years since 2008, and the General Household Survey, which is conducted once a year as a general purpose survey but also collects labor market information.

of the sample is replaced each quarter. [33] The sampling unit is the residential address, whereas the unit of observation is the household. As such, in case of residential turnover, the new household is surveyed in place of the older one. Each household is interviewed in a maximum of four consecutive quarters.

In order to obtain repeated and independent cross-sections and avoid issues of correlated measures within households over time, I restrict the sample used in my analysis to the first time each dwelling unit in the survey is observed in the data, which I call the incoming rotation group. The rotation groups are designed such that each group has "the same distribution pattern as that which is observed in the whole sample" (Statistics South Africa, 2008) and this ensures that the incoming rotation group constitutes a nationally representative sample. Using the incoming rotation group avoids attrition issues, especially if individuals or households leave the survey in ways that are related to eligibility to the ETI.[34] Papers such as Verick (2012) have used algorithms to recover the panel structure of the dataset, specifically by tracking changes in the structure of the households interviewed in the same dwelling place. As I am not interested in changes within individuals but in comparing cohorts over time (see more on the methodology in the section below), the repeated cross-section component is suitable for the analysis. Also, given the short

[33] This dataset can be compared to the US Current Population Survey, which also has a rotating panel structure. According to the United States Census Bureau, "households from all 50 states and the District of Columbia are in the survey for 4 consecutive months, out for 8, and then return for another 4 months before leaving the sample permanently". (see https://www.census.gov/programs-surveys/cps/technical-documentation/methodology.html). For the QLFS, dwelling units are in the survey for four consecutive quarters and then leave the sample permanently.

[34] In order to have a sample of independent cross-sections, studies in the US that use the CPS usually resort to the outgoing rotation group, i.e. households that are leaving out the survey for four months or indefinitely, but more so because these households are asked additional questions. In the case of the QLFS, since there is no additional advantage of using the equivalent outgoing rotation group, using the incoming rotation makes more sense. It further avoids attrition issues, to the extent that these issues are not properly accounted for by the weight corrections performed by Statistics South Africa.

length of time that individuals stay in the panel (four consecutive quarters at most), the panel dimension of the dataset is not useful for my estimation strategy.

The QLFS is designed to be representative at the province level, and within the province level, at the metropolitan/non-metropolitan area level. [35] The sample is based on a stratified two-stage sampling design; in the first stage, primary sampling units (PSUs) are chosen with probability proportional to their size, and in the second stage, dwelling units are systematically drawn with equal probability from each sampled PSU. At the first stage, the sampling is stratified by the province-metro/non-metro areas. This survey design produces inverse probability weights that can be used to produce statistics that are representative of population level parameters. Statistics South Africa further adjusts those weights for non-responses, so that the information from the survey is consistent with the information available at the time on the structure of South Africa's population. [36] I treat those final weights as population weights for my estimation purposes.

There have been changes over time in the methodology of the survey, especially in the master sample from which primary sampling units are drawn. From 2008 to 2014, the QLFS used a master sample based on the 2001 census data. A change occurred in 2015, with the master sample rebased to the 2011 census. This poses problems for using time variation that includes data before and after the update, as any break in the structure of the data can be a confounding factor that prevents interpretation of the estimates as causal effects of the policy. There is still debate about the implications of these changes in the

[35] Metropolitan municipalities are important municipalities that enjoy an autonomy in their administration, as opposed to other areas. There as 8 such metropolitan areas, with three of them in the province of Gauteng. Provinces are consequently often divided into metropolitan municipalities and non-metropolitan areas.

[36] See Statistics South Africa (2008) for a discussion of the computation of the weights.

survey on measurement of employment (see Kerr and Wittenberg, 2019, for a discussion). As such, I only use data from 2014 and earlier, which use the 2001 census' master sample, and cover the first year of the policy's implementation.

I consider as employment outcomes overall employment (employment in any firm), private employment and private formal employment. I consider workers employed in the private sector when they work for a private firm and in private formal employment when they report that their employer deducts income tax from their salary. The goal of the ETI is to increase employment for young workers. The subsidy does so by targeting private firms that are registered for tax purposes. Analysis of private formal employment gives us the direct effects of the policy on the targeted outcome. However, the underlying goal of the policy is to increase overall employment for youths, which may extend beyond private formal employment. For example, youths could respond through changes between jobs, such that overall employment does not change for young workers but private formal employment increases. Another example is that some workers may be encouraged to search and find informal jobs, but not formal jobs, which would increase employment beyond any effect on private formal employment. As such, I consider any private employment and any employment as additional employment outcomes.

I additionally consider participation in the labor market as an outcome, for two reasons. First, the subsidy may have induced young workers to look for jobs, if these workers believed that what their decreased costs to firms may increase their chances of being hired. Second, because of this first point, looking at both employment outcomes and participation will give us an idea of the effects of the subsidy on unemployment.

1.4.2 Identification strategy

Because the policy has a clear start date and clear eligibility rules for workers whose wages can be subsidized, I use a difference-in-difference strategy to estimate the effects of the Employment Tax Incentive on young workers' employment outcomes.[37] A firm can claim the subsidy if it hired a worker after October 2013 and the worker was between 18 and 29 years old at the time he or she was working at the firm. Although this criterion is labeled in terms of current age, it is more intuitive to reason in terms of cohorts. Table A.1 summarizes the cohorts who were eligible and the time when they were eligible. Using current age as the distinguishing criterion between the treatment and control groups for the difference-in-difference analysis could be misleading. Because eligibility stops once the individual turns 30, using current age rules out the effects of the policy on workers after their eligibility. The justification for a wage subsidy was to propel young workers into the labor market; if they stay employed even after their eligibility ends, this should be part of the effects of the policy. Treated workers could stay in employment after their eligibility because they keep the job that they were recruited for under the policy, or because they get other jobs, using the experience that they accumulated from subsidized employment. Arguably, these are not major concerns as I am only interested in the first year of implementation, but I adopt the most accurate way to define eligible and ineligible

[37] There is always the concern that agents can anticipate future policies and act accordingly. This can lead to bias in the estimation of the effects of these policies. In the case of the ETI, there seems to have been a long debate at the parliament on this policy, as early as in 2011; also, there was no clear announcement that predated the signing of the ETI act. Given this, such anticipations are not a concern in this specific case. However, I provide a discussion below on how the placebo test in periods before the policy addresses to some extent the anticipation concerns.

individuals. [38]

The treated group is the group of individuals who are 29 or younger as of 2013Q4. Because the QLFS only reports the year of birth, this translates into individuals born in or after 1984 (see Table A.1).[39] The control group is composed of individuals who were 30 or above at the time of implementation, meaning that they were born in 1983 or before. For the estimation, I need to choose a sample that increases the chances of finding the effects of the policy, with individuals in the sample not too far from the cutoff. For this reason, unlike Ranchhod and Finn (2016) who use the whole sample of individuals of working age, and in line with Moeletsi (2017), I use a sample of individuals 5 years above (control group) and 5 years below (treatment group) the cutoff provided by eligibility of workers for the subsidy. The main sample in my estimations is thus the group of individuals born between 1979 and 1988.

One caveat in the use of the difference-in-difference approach in the present case is the fact that the control group in my analysis is potentially affected by the policy. As discussed in the literature review above, when wages are subsidized for part of the population, firms could substitute ineligible workers for eligible ones, meaning that ineligible workers are potentially negatively affected by the policy, or negatively treated. In that case, the difference-in-difference will provide at best, the effect of the wage subsidy on

[38]Studies of nationally implemented wage subsidies which targeted individuals based on age (e.g., Balkan et al., 2016) likely estimate a lower bound of the effects of these policies. Once an individual passes the age of eligibility, they become part of the control group in those studies' empirical strategy and thus contribute to the construction of the counterfactual. To the extent that there are returns to the experience accumulated because of subsidized employment, this overestimates what would be the employment for the control group in the absence of the subsidy, and biases the estimate of the policy down.

[39]Although the questionnaire of the QLFS asks for the exact date of birth, the publicly available dataset only contains the age in years. Discussions with staff at DataFirst suggest that this may have to do with the fact that an individual's national identification number is composed of his or her day of the birth, information that needs to be removed for anonymity purpose. Consequently, I can only consider in my analysis the age at the time of the survey or the year of birth.

eligible workers relative to ineligible ones, including both jobs created by firms for young workers as well as jobs that they obtain at the expense of ineligible ones. The interpretation of such effects would be that firms reacted to the policy, with no direct way to tell what fraction of jobs are created, and what fraction is due to substitution. The design of the policy and the regulations in South Africa imply, however, that such substitutions should be limited. [40]

Table A.2 presents summary statistics on employment outcomes and other characteristics for the treatment and control groups, before and after the policy. Workers in the young cohort saw their employment increase for all measures of employment, though these before and after statistics do not necessarily reflect causal effects of the program, for reasons that I discuss below.

I estimate the following equation to obtain the difference-in-difference estimates using a linear probability model: [41]

$$E[Y_{icpt}|X_{icpt}] = \beta_0+\delta_0 d_c+\delta_1 d_c{*}post_t+\beta_1 exp_{cpt}+\beta_2 exp^2_{cpt}+\gamma X_{icpt}+prov_p+Year_t+Qtr_t \tag{1.1}$$

Y_{icpt} is the employment outcome of interest, for individual i in cohort c, who lives in province p and was interviewed at quarter t. d_c is a dummy for young cohorts that is 1 for cohorts c =1984-1988 and zero for cohorts c =1979-1983. In my main specification, I use data for $t = 2011Q1, \cdots, 2014Q4$. $post_t$ is a dummy variable for periods during and after $2013Q4$, the first five quarters of implementation of the ETI. δ_1 is the difference-

[40] I discuss more this issue in relation to the results in section 1.5.1.

[41] The absence of an error term in the equation below reflects the fact that it is a linear probability model that I estimate.

in-difference parameter, the parameter of interest. The equation controls for experience with a proxy constructed at the year of birth-province level, exp_{cpt}. Province of residence ($prov_p$) is included to capture province specific time-invariant characteristics that may be correlated with employment opportunities. X_{icpt} are additional control variables, including race, gender, and education.

The QLFS does not report information on labor market experience. However, it is important to control for experience in the equation above. The treatment and the control group differ on the experience that they have acquired in the labor market, with older individuals having greater experience, and it is known that accumulated experience affects labor market outcomes, typically with a quadratic functional form. I construct a proxy for experience at an aggregate level. For example, in 2012, 39.2% of individuals born in 1984 and living in Gauteng were employed; this implies that, on average, the group of individuals born in 1984 and living in Gauteng has accrued .39 year of labor market experience in 2012.[42] To calculate this experience proxy, I use data from all labor force surveys conducted since the older cohort entered their working age then compute the average employment ratio for each year of birth cohort at the province level. I next sum these average employment ratios to obtain my proxy for experience. Experience at time t for the cohort c living in province p is then the cumulative exposure to the labor market that this group has had until time $t-1$. Figure 1.2 below presents the distribution of this proxy for the young and the old cohorts, before the policy, with more detailed information in Figure A.2.

[42] I consider the frequency of the surveys (biannually until 2007, then quarterly afterwards) in computing the experience proxy. The example here uses the annual employment ratio for illustration purposes.

Figure 1.2: Distribution of experience proxy

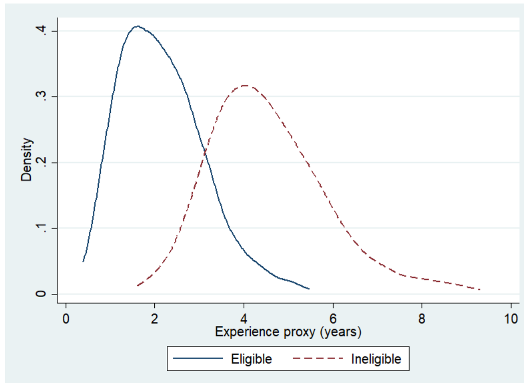

All the estimates that I present are weighted by the adjusted inverse probability weights that Statistics South Africa provides for the QLFS. Although the use of weights is common for the estimation of parameters describing variables' distribution, there is a long debate in economics and social science in general on whether those weights should be used when running regressions (see e.g., Solon et al., 2015, for a discussion on the subject). The consensus seems to be that weights should not matter, especially when the variables describing the sampling scheme, and thus the weights are included in the model that is estimated. If they do not matter, then not using them should be preferred, as the OLS estimator is BLUE, while the weighted least squares estimator is not. However, in the case of difference-in-difference, the regression estimated to obtain the estimator is often guided both by economic intuition. In the end, it is a difference in means that is estimated and as such, the means should be estimated to be representative of the population. Although the variables that I include in the model are guided by economic knowledge of labor markets in general and South Africa's labor market in particular, the goal is to adjust for any characteristics that may be affecting labor market outcomes differently for the young and the old cohorts. All this favors the use of weights in the present case.

1.4.3 Suggestive evidence of parallel trends

The identifying assumption for the difference-in-difference strategy is that the treatment and the control groups would have evolved in the same way if no policy intervention were made, such that the changes in the outcomes of the control group are a valid counterfactual for the treatment group, in the absence of the policy. Although this assumption cannot be tested directly, it is possible to check if, before the policy, employment outcomes for the young and the older cohorts evolved in similar ways. There are two ways to implement this test. The first is to estimate differences between the two groups, conditional on observables, in periods before the policy, and test whether they are significantly different over time. I use the following equation for this test:

$$E[Y_{icpt}|X] = \beta_0+\delta_0 d_c+\sum_{q=2011Q1}^{2013Q2}\delta_q d_c * 1_{\{t=q\}}+\beta_1 exp_{cpt}+\beta_2 exp_{cpt}^2+\gamma X_{icpt}+prov_p+Year_t+Qtr_t \tag{1.2}$$

Y_{icpt} is the employment outcome of interest, for individual i in cohort c, who lives in province p and was interviewed at quarter t; d_c is a dummy for eligible cohorts that is 1 for cohorts $c =$1984-1988 and zero for cohorts $c =$1979-1983. I use data for the periods $t =$ 2011-$2013Q3$. In the equation above, δ_0 is the difference between young and old cohorts in $2013Q3$. $\delta_q + \delta_0$ is the difference for each of the quarters $q = 2011Q1, \cdots, 2013Q2$. The test of parallel trends before the policy is thus the test of joint significance of the coefficients δ_q. Figure 1.3 presents these coefficients. Although the figure points to differences in the coefficients δ_q for all the outcomes over time, with most of these coefficients being positive, the confidence intervals associated with the estimates do not allow to rule out that they are jointly equal to zero. As a result, I do not reject that there were parallel trends in outcomes between the eligible and the ineligible cohorts before the policy. The coefficients and confidence interval suggest overall however that the test is a bit imprecise.

Figure 1.3: Testing parallel trends prior to the policy

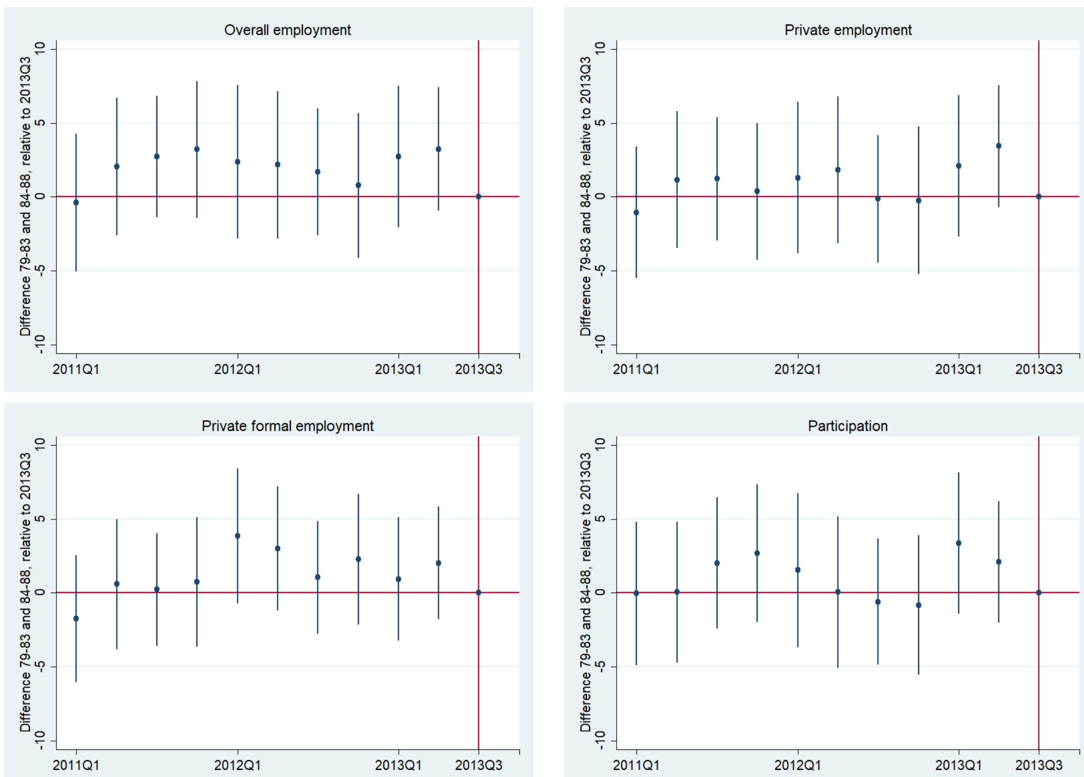

Notes: The graphs above present the result of estimating Equation 1.2. The coefficients plotted here are the coefficients of interaction between year-quarter dummies and the dummy for young cohorts. They represent the difference between the old and the young cohort, relative to the difference in $2013Q3$, the last quarter before the implementation of the ETI. The bars are the 95% confidence intervals.

A similar test is whether there has been a differential change in employment outcomes between the eligible and the ineligible cohorts, in periods just prior to the implementation of the policy. I perform this test for the first three quarters of 2013, i.e. the three quarters immediately before the implementation of the ETI. This is intuitively a placebo test of a change before the policy was implemented. If the test detects a change, then it is unlikely that eligible and ineligible cohorts had the same trend. To implement this test, I run the following regression, using data for $t = 2011Q1, \cdots, 2013Q3$:

$$E[Y_{icpt}|X] = \beta_0^P + \delta_0^P d_c + \delta_1^P d_c * placebo_t + \beta_1^P exp_{cpt} + \beta_2^P exp_{cpt}^2 + \gamma^P X_{icpt} + prov_p + Year_t + Qtr_t \tag{1.3}$$

where $placebo_t$ is a dummy that is 1 for the quarters $2013Q1$ to $2013Q3$ and zero for 2011 and 2012. If the parallel trend assumptions hold, then the coefficient δ_1^P should be zero. The estimations in Table 1.1 show that I cannot reject that this coefficient is equal to zero. In particular, the estimates for private employment indicate that if anything, young workers may have been catching up to the older workers. This points to the importance of controlling for experience and suggests that my estimates are likely an upper bound on the true effects of the policy.

Results of equation 1.3 also provide evidence against the idea that firms or workers took anticipatory actions in advance of the policy implementation. Such reactions would include firms firing young workers with the expectation of rehiring them and thus benefiting from the wage subsidy, or firms postponing their hiring decisions regarding young individuals. Both actions would show up as a relative decline in employment outcomes for young workers, leading up to the policy start date. While the policy had a provision against displacement of older workers (see the earlier presentation of the ETI), there was no explicit mention of re-hiring, leading to concerns that there might have been some temporary separations between firms and workers. The presence of such effects would imply a negative coefficient δ_1^P. The fact that this coefficient is not statistically different from

Table 1.1: Placebo effect using 2013 first three quarters

	Overall employment		Private employment		Private formal employment	
	(1)	(2)	(3)	(4)	(5)	(6)
δ_1^P, placebo in 2013	1.174	0.193	2.283	1.323	0.519	-0.268
	(1.357)	(1.289)	(1.433)	(1.336)	(1.330)	(1.224)
Experience proxy	12.245	8.360	9.696	6.081	10.125	4.824
	(1.197)***	(0.861)***	(1.073)***	(0.821)***	(0.985)***	(0.747)***
Experience proxy square	-0.685	-0.551	-0.487	-0.422	-0.487	-0.339
	(0.113)***	(0.090)***	(0.104)***	(0.082)***	(0.098)***	(0.078)***
Controls	N	Y	N	Y	N	Y
R^2	0.04	0.16	0.02	0.13	0.03	0.14
N	39,407	39,407	39,407	39,407	39,407	39,407

Notes: Data are from the QLFS 2011-2013Q3. Sample: individuals from the incoming rotation group, born in 1979-1988. δ_1^P is the coefficient for the placebo test before the policy (see Equation 1.3). Standard errors in brackets are clustered at the stratum level. Regressions are weighted using the sample weights from the QLFS and estimated using a linear probability model. Controls are race (Blacks, Colored, Asians/Indians and Whites), gender, education (no education, some primary, primary completed, some secondary, secondary completed and post-secondary), marital status, the presence of an individual older than 65 and the presence of another individual who is working, all included in the regression equation as dummy variables.

zero, and if anything slightly positive for employment and private employment, helps rule out these concerns.

1.5 Results

1.5.1 Main results

Table 1.2 presents the results of the estimation of Equation 1.1. The specifications in the first two columns do not control for potential experience, while the two last columns do. The first two specifications imply that the overall employment for young workers increased by between 1.3 and 1.8 percentage points, relative to an employment ratio of 44 percent before the policy began. Although these results are not statistically significant, I cannot rule out effects as large as 3 percentage points, and this includes an increase of 2 percentage points, the size of the effects found in Moeletsi (2017). Once I control for potential experience, my estimates suggest that the policy led to a *decrease* in the probability of employment, but not a statistically significant one, of between .2 and .7 percentage points. More importantly, I can rule out benefits as small as a 1 percentage point gain. Notably, controlling for experience reduces the estimated impact of the ETI.

The estimates for private employment and private formal employment are similar in their magnitude and their precision to the estimates for overall employment. Even though only firms that were private and registered for income tax purposes were eligible to claim the subsidy for workers that they hire, the policy could have resulted in reallocation of labor instead of job creation, such that effects on private employment or overall employment are different from the effects on private formal employment. Similar results for all these outcomes imply that the policy did not lead to such reallocation.

Table 1.2: Impact of the ETI on youth employment: main results

	(1)	(2)	(3)	(4)
A: Overall employment				
δ_1	1.842	1.353	-0.206	-0.735
	(0.965)*	(0.891)	(0.984)	(0.893)
Experience proxy			12.120	8.549
			(1.118)***	(0.741)***
Experience proxy square			-0.716	-0.622
			(0.090)***	(0.071)***
Average for young cohort prior to the policy:		43.9%		
R^2	0.01	0.15	0.03	0.15
N	56,453	56,453	56,453	56,453
Controls	N	Y	N	Y
B: Private employment				
δ_1	1.157	0.752	-0.137	-0.765
	(1.046)	(0.968)	(1.067)	(0.979)
Experience proxy			9.524	5.854
			(1.002)***	(0.688)***
Experience proxy square			-0.507	-0.441
			(0.085)***	(0.064)***
Average for young cohort prior to the policy:		35.8%		
Controls	N	Y	N	Y
R^2	0.01	0.13	0.02	0.13
N	56,453	56,453	56,453	56,453
C: Private formal employment				
δ_1	0.672	0.446	-0.485	-0.782
	(1.005)	(0.905)	(1.005)	(0.927)
Experience proxy			9.860	4.390
			(0.954)***	(0.638)***
Experience proxy square			-0.492	-0.347
			(0.084)***	(0.060)***
Average for young cohort prior to the policy:		27.2%		
R^2	0.00	0.14	0.02	0.14
N	56,453	56,453	56,453	56,453
Controls	N	Y	N	Y

Notes: Data are from the QLFS 2011-2014. Sample: individuals from the incoming rotation group, born in 1979-1988. δ_1 is the difference-in-difference estimator (see Equation 1.1). Standard errors in brackets are clustered at the stratum level. Regressions are weighted using the sample from the QLFS and estimated using a linear probability model. Controls are race (Blacks, Colored, Asians/Indians and Whites), gender, education (no education, some primary, primary completed, some secondary, secondary completed and post-secondary), marital status, the presence of an individual older than 65 and the presence of another individual who is working, all included in the regression equation as dummy variables.

In the previous section on parallel trends before the implementation of the ETI, I noted that if anything, young workers' outcomes were catching up to those of workers a few years older. If that trend continued even in the absence of the policy, young workers may have continued catching up with older workers. Then, the effects estimated with a difference-in-difference strategy would tend to over-estimate the effects of the policy. This further reinforces the interpretation that the policy did not have meaningfully large positive effects.

One concern with the preceding analysis is that I use as a control for experience a proxy constructed at an aggregate level, the birth year-province (*cp*) level, while analyzing an outcome at the individual level. Measurement errors in aggregation could lead to bias in the implied effect of the policy. For this reason, I re-estimate Equation 1.1 at the level at which the proxy for experience is constructed. Table 1.3 presents the results of the aggregated level regression. There are 10 birth cohorts in my sample (5 in the eligible group and 5 in the ineligible cohort) and 9 provinces in South Africa, which results in 90 units of observation. I use 16 quarters for my main estimations (2011-2014); this amounts in the end to 1440 observations, as shown in Table 1.3. The estimates in Table 1.3 are similar to the ones at the individual level (Table 1.2, columns 3 and 4), suggesting that measurement errors due to aggregation are unlikely to explain my results.

Table 1.3: Impact of the ETI on youth employment: aggregate level regressions

	Overall employment		Private employment		Private formal employment	
	(1)	(2)	(3)	(4)	(5)	(6)
δ_1	-0.206	-0.858	-0.137	-0.934	-0.485	-0.857
	(1.061)	(0.822)	(1.051)	(0.894)	(1.212)	(0.913)
Experience proxy	12.120	8.316	9.524	5.627	9.860	4.505
	(1.033)***	(0.804)***	(1.053)***	(0.672)***	(1.392)***	(0.629)***
Experience proxy square	-0.716	-0.609	-0.507	-0.435	-0.492	-0.355
	(0.095)***	(0.068)***	(0.100)***	(0.061)***	(0.134)***	(0.058)***
Controls	N	Y	N	Y	N	Y
R^2	0.46	0.61	0.32	0.54	0.31	0.63
N	1,440	1,440	1,440	1,440	1,440	1,440

Notes: Data are from the QLFS 2011-2014. Sample: individuals from the incoming rotation group, born in 1979-1988. Data are aggregated at the birth cohort-province level. δ_1 is the difference-in-difference estimator (see Equation 1.1). Standard errors in brackets are clustered at the stratum level. Regressions are weighted using the sample weights from the QLFS and estimated using a linear probability model. Controls include race (Blacks, Colored, Indians, Whites), gender, education (no education, some primary, primary completed, some secondary, secondary completed, post-secondary), marital status, the presence of an individual older than 65 and the presence of another individual who is working, all included in the regression equation as dummy variables.

Table 1.4 presents the effects of the policy on young workers' labor force participation. I look at this outcome in isolation from the others as it reflects the reaction of individuals to the prospect of a wage subsidy. Given that they are now a relatively cheaper labor, young workers may have internalized the increased value of job searching and responded by increasing their labor force supply.[43] The estimates, however, show that there was no increase in labor force participation for young workers as a result of the ETI. In fact, the estimates for labor force participation are similar to the ones for employment outcomes, suggesting that young workers did not increase their labor force supply when the ETI became available. One possibility is that these young workers may have correctly antici-

[43]By looking at labor force participation, I investigate the effects of the ETI on the extensive margin of the labor supply.

pated that the policy would not increase their chances of employment. By combining the results on employment and labor force participation, it follows that the ETI did not affect unemployment for young workers. [44]

Table 1.4: Effects of the ETI on labor force participation

	(1)	(2)	(3)	(4)
δ_1	0.701	0.322	-0.502	-0.952
	(0.947)	(0.904)	(0.995)	(0.926)
Experience proxy			5.601	6.330
			(1.004)***	(0.772)***
Experience proxy square			-0.377	-0.412
			(0.084)***	(0.073)***
Controls	N	Y	N	Y
R^2	0.01	0.09	0.01	0.09
N	56,453	56,453	56,453	56,453

Notes: Data are from the QLFS 2011-2014. Sample: individuals born in 1979-1988 from the incoming rotation group. δ_1 is the difference-in-difference estimator (see equation 1.1). Standard errors in brackets are clustered at the stratum level. Regressions are weighted using the sample weights from the QLFS and estimated using a linear probability model. Controls are race (Blacks, Colored, Asian/Indians and Whites), gender, education (no education, some primary, primary completed, some secondary, secondary completed, post-secondary), marital status, the presence of an individual older than 65 and the presence of another individual who is working, all included in the regression equation as dummy variables.

Overall, my results suggest that the Employment Tax Incentive did not produce to the (large, positive) effects expected from the policy. My results are similar to those in Ranchhod and Finn (2016), despite using a more appropriately defined sample that does not include much older and more experienced workers than those targeted by the ETI. These results, however, stress the importance of controlling for experience when accounting

[44] Unemployment is defined as employment among workers who supply their labor force, and neither employment nor participation is affected by the policy. Looking at the the effects of the policy on employment and labor force participation separately to infer the effects on unemployment avoids analyzing a sub-sample (active workers) that is potentially changing over time because of the policy, if the policy had an effect. I additionally analyze the effects on unemployment, and the results in table A.3 show that indeed youth unemployment was not affected by the policy.

for differences between the eligible and the ineligible cohorts. Not doing so leads to an over-estimation of the effects of the ETI.

There have been a number of concerns raised in the literature about wage subsidies that I investigate here to explore the extent to which they may confound or qualify my findings (Card et al., 2015; McKenzie, 2017). The first concern is displacement effects. Because of the targeting, the wage subsidy can lead to displacement effects, whereby workers that are not targeted by the policy are fired and replaced by targeted ones. This may be less likely in South Africa given the design of the ETI and features of the labor market. The ETI Act includes a penalty for firms that fire an older worker to hire one who qualifies for the subsidy, which should deter such displacements. There are also high perceived dismissal costs in the country. Workers have the right to challenge unfair dismissal, and in this specific case, older workers could challenge unfair dismissal to hire a young worker. Previous research (e.g., Benjamin et al., 2010) suggests that the disruptions that these challenges cause to firms is a serious concern in general, and this would have discouraged firms to fire older workers to recruit younger ones. A final point is the absence of large positive effects which further casts doubt that displacement effects might have occurred. If there were such displacement effects, employment outcomes of the ineligible workers would have fallen compared to a situation where no policy was implemented. The difference-in-difference estimator would then capture both jobs obtained by eligible workers, and jobs that were displaced, overstating the effect on eligible workers. The fact that I can rule out even small positive effects provides evidence against such concerns. It thus seems in the end unlikely that there would have been such displacement effects.

A second concern is substitution effects. When wage subsidies are targeted, firms may hire eligible workers at the expense of other workers, which is labeled the substitution effect. This is different from displacement effects because firms do not fire workers. If there were substitution of old workers by young workers, the difference-in-difference estimates would capture these effects, since it is estimated out of differences in outcomes for young and older workers. The job creation effects would then be lower than the effect estimated here. But the fact that the estimates are low implies that substitution is not widespread. All young workers are eligible, and as such, substitution between young workers would bot be induced by the ETI.

The third concern is deadweight loss: firms could take up the wage subsidy and hire workers that they would have hired anyway, even in the absence of the wage subsidy. This concern was in the mind of policymakers when devising the policy, as reflected in the newly hired requirement to be able to claim the subsidy, along with the penalty for displacement. However, none of these requirements prevent a firm that would have hired young workers anyway from claiming the subsidy. The combination of high claims by firms and no effects on employment suggest that the claims of the subsidy are entirely driven by this behavior. This is consistent with estimates in Ebrahim et al. (2017) using firms' administrative data, which indicate that that firms that claimed the ETI increased hiring of young workers, but also hiring of older workers. Another interpretation of their results is that firms claimed the subsidy as they expanded and hired both young and old workers, independently of the ETI.

1.5.2 Robustness checks

Before discussing the robustness of my results based on the main strategy as described by section 1.4.2, and in particular as summarized in Equation 1.1, I offer alternative estimations of the effects of the policy based on the definition of eligibility. I argued in section 1.4.2 that the best way to define the eligible and the ineligible groups is by looking at cohorts. Using current age to define those groups should however play little difference, especially given that the effects are estimated only in the first year. [45] I present in Table 1.5 the results of an equation similar to Equation 1.1, but with eligibility defined in terms of age. To avoid issues of systematic differences between the eligible and the ineligible groups, I restrict the sample to individuals aged 29 or 31, with individuals aged 29 in the eligible group. The estimates suggest results that are globally in line with my estimates using cohorts to define eligibility. Eligible workers did not experience significant improvement in their employment outcomes, relative to ineligible workers. The estimated effects seem to suggest possible larger positive effects in overall employment, but if anything negative effects for private formal employment, the main target of the policy. Overall, the estimates become less precise, as the sample size shrinks due to the age restriction. I take this overall as evidence that the results are robust to defining eligibility in terms of current age or cohorts.

I test the robustness of the results to some concerns that may confound these results. The first concern is that firms can claim the subsidy for periods of different lengths depending

[45] This rules out possible returns to experience that would be best captured by defining eligiblity in terms of cohorts, allowing eligible individuals to stay in the estimations as ”treated”, after they are no longer eligible for the subsidy.

Table 1.5: Main effects of the ETI, comparing 29 and 31 years old over time

	Overall employment		Private employment		Private formal employment	
	(1)	(2)	(3)	(4)	(5)	(6)
Age 29, after 2013Q4	1.261	-0.113	0.778	-0.224	-0.551	-1.596
	(1.830)	(1.647)	(1.828)	(1.628)	(1.711)	(1.512)
Age 29 (base: age 31)	-3.366	-2.175	-1.964	-0.772	-1.436	-0.556
	(1.209)***	(1.134)*	(1.114)*	(1.027)	(1.021)	(0.935)
R^2	0.00	0.15	0.00	0.14	0.00	0.15
N	11,082	11,082	11,082	11,082	11,082	11,082
Controls	No	Yes	No	Yes	No	Yes

Notes: Data are from the QLFS 2011-2014. Sample: individuals from the incoming rotation group, aged 29 or 31 at the time of survey. The results presented are the coefficients for individuals aged 29, and its interaction with a dummy variable for the periods after the start of the policy. Standard errors in brackets are clustered at the stratum level. Regressions are weighted using the sample from the QLFS and estimated using a linear probability model. Controls are race (Blacks, Colored, Asians/Indians and Whites), gender, education (no education, some primary, primary completed, some secondary, secondary completed and post-secondary), marital status, the presence of an individual older than 65 and the presence of another individual who is working, all included in the regression equation as dummy variables.

of the age of the worker. Firms can claim the subsidy for workers for at most two years, or until these workers become 30 years old. This implies that firms can claim the subsidy for two years only for workers that are hired at 28 or below, which translates to workers born in 1986 or afterwards. I use individuals born between 1986 and 1988 as my eligible group, and compare them to individuals born between 1981-1983.[46] From the perspective of the value of the subsidy for firms (i.e. ignoring productivity concerns), it is more valuable for firms to hire workers for whom they can claim the subsidy for a longer period, unless these workers are less productive. This would imply larger effects for these workers. Overall, the the effects on employment for younger workers that are eligible for two years (Table A.4, column 2) are still not significant. Even when considering the point estimates and the largest effects that they imply, there is no clear pattern regarding differential effects on these younger workers. The effects estimated on private formal employment and private employment are consistent with larger effects on these young workers, whereas the opposite is true for overall employment. Overall, these results are not consistent with larger effects on workers for whom firms could claim the subsidy longer.

Second, I address the concern that there might be differential returns to education for the eligible and the ineligible cohorts. The system of education has changed over time in South Africa, along changes in the political landscape with the end of the apartheid regime, although the changes in education sometimes predated the political end of apartheid.[47]

[46] The latter group is chosen both to have a 3 year cohorts band for ineligible workers, just as for the workers that are in the eligible group. The cohorts in the middle, those born 1984 or 1985 are workers that are eligible for less than two years.

[47] One characteristic of the system of education under the apartheid system was racial segregation of schools. School desegregation started as earlier as 1990 in the province of Gauteng, years before the end of apartheid regime.

The sample in my analysis spans several cohorts that potentially were differentially affected by these changes. These individuals were as young as 6 and as old 15 in 1994, the year marking the end of the apartheid regime. To the extent that there have been changes in the quality of education, and thus changes in the returns to education, this may confound the estimates of the effects of the policy, as these are differential changes for young and old cohorts just as the change induced by the ETI. I address this concern by allowing the employment returns to education to differ by time (Table A.4, column 3) or by cohorts (eligible, and ineligible; Table A.4, column 4) in Equation 1.1. The results from these specifications are similar to the main results, which implies that different returns to education is not a concern for my estimates of the effects of the ETI.

I then address a concern related to the effects of experience on employment outcomes. My main specification assumes the same experience profile for eligible and ineligible workers. It may be the case that worker eligible for the subsidy, who are on average younger, are on a different experience profile. I address this concern by allowing the experience proxy to enter Equation 1.1 differently for eligible and ineligible individuals. The results over the three employment outcomes (see table A.4, column 5) are similar to the main results, suggesting that my results are not sensitive to differential experience profiles.

I finally test the sensitivity of my results to the functional form of my experience proxy. The quadratic relationship between employment outcomes and the experience proxy in Equation 1.1 is primarily motivated by its common use in such equations. However, given the fact that I do not observe experience itself, it could be argued that a different functional form may be more suitable. Table A.5 presents the results of estimations when I consider several alternative specifications. These specifications include a linear function, a cubic function, and the use of deciles of the experience proxy to account for a potentially non-parametric relationship. The results using the cubic relationship and the non-parametric relationship indicate results that are overall similar to my main specifications (quadratic functional form). However, the results are compatible with higher effects when using a linear relationship. This indicates that not including the quadratic term of the experience proxy acts as omitting a relevant variable that biases my results upwards. The downside of this analysis is that my results are sensitive to the functional expression of my proxy for experience. [48]

In the end, across all the results, experience seems to matter a lot for my results (see presentation of the main results in Table 1.2). As I argued in the methodology section, I adopt this empirical strategy to account for time-varying differences in experiences that may confound the changes introduced by the wage subsidy policy, especially given that eligible and ineligible workers are at different phases of the employment profiles. The importance that experience plays could have been expected given the importance of experience for labor outcomes and the quadratic functional form. Nonetheless, there remains the issue that control variables should not matter in a difference-in-differences setting, as the strategy removes any pre-existing differences between differences. Control variables

[48] I address an additional concern related to the precision of the results. My main estimations throughout this article cluster standard-errors at strata level, to take into account the design of the QLFS. Other considerations also govern the choice of the level of clustering are the fact that outcomes maybe correlated within units, such as local labor markets, or the units of treatment assignment. These additional concerns would favor clustering at the cohort level, given that whole cohorts are treated; the cohort X time level, given that the treatment occurred at specific dates; the province X cohort level or the province level to take into consideration the correlation of outcomes within labor markets. The province X cohort level would also address to fact that the experience proxy is measured at that level. To assess the influence of such concerns on the results, I consider alternative clustering at the province X cohort level, and at the cohorts level. In the later case, given the number of cohorts (10), I use a wild bootstrap to perform the tests. Results presented in Table A.11 show that there is no meaningful changes in the results due to these alternative clustering.

should matter only if they are necessary to establish suggestive evidence that the parallel trend assumptions hold only conditional to these control variables. It is the case that control variables do diminish the implied placed effects reported in 1.3, especially for private employment. The results overall are, however, not significant, with or without the inclusion of experience. It is also true that the effects estimated, besides in one specification for overall employment. But again, in conjunction with the results on participation, it would be hard to argue how the policy led to an increase in overall employment, without a meaningful change in private formal employment. I take this collectively at the least as the lack of robust evidence that the policy improved employment outcomes for eligible young workers, even as the inclusion of the experience proxy seems to affect the results a lot.

Regarding other control variables, they can be divided into two categories: those that are time-invariant, such as race, gender, and those that are time-varying, such as education, the structure of the family, and marital status. Control variables that do not change over time would not matter for the estimates of the results but would lead to more precise estimates. Among control variables that are time-varying, education is one that could potentially be problematic. Indeed, young individuals (who are eligible for the policy) can choose to stay longer in school since education is another important variable explaining employment outcomes in South Africa. Investing more in education could be seen as a response to the greater prospect of employment caused by a wage subsidy. So, to the extent that the policy affects employment through education, its inclusion in the equation will cause a reduction of the estimated effect of the ETI. However, it is unlikely that the indirect effect through education would be observed in the immediate future. The reasons are that I focus only on the first year of implementation of the ETI, and that eligible individual in my main sample are at least 26 years old, with less room to increase their education. Also, the effects on participation would suggest that such concerns are not present in this setting. The structure of the household and marital status are also variables that could respond. It is less clear how immediate, such responses could be to the policy. An additional concern is related to the pre-trend tests that I conducted in section 1.4.3. The results that I obtained, including the placebo test in 2013 seem imprecise. I obtain large coefficients, but for which I cannot reject the hypothesis that they are different from zero. Also, these coefficients seems to be if anything, larger than some of the estimates of the effects of the policy. Although this would suggest that the results may be in the end not precise, this is not consistent with the much smaller standard errors that I obtain for the estimations of the effects of the policy.[49]. The comparison of the precision of the placebo test and the actual estimates could actually be misleading, as the placebo tests are by construction less powered compared to the actual estimates, given that the actual estimates use a larger sample, on the time dimension. Also, the comparison of the placebo effects to the actual estimates would preclude the fact that the estimates are obtained conditional on the acceptance of the hypothesis of parallel trends. In the end, the conclusion are to be considered within the limits of the precision of my pre-policy trend, which is evidenced by the large coefficients for the placebo tests, although the signs of the placebo tests suggest that if anything, my results are an overestimation of the effects of the policy.

[49]That said, some of the estimates are compatible with important increases in employment, although this is not the consistent conclusion across the estimations.

1.5.3 Heterogeneity of the effects

It is possible that a null effect on the average young worker masks a combination of positive and negative effects on different subgroups of youth. Therefore, I explore the heterogeneity of the effects of the policy along dimensions that are regularly considered in studies of South Africa labor market (Banerjee et al., 2008): gender (Table A.6), race (Table A.7), and education (Table A.8). Women, Blacks (referred to simply as Africans) and individuals with low education have worse employment outcomes. Of the three, arguably only the difference in employment outcomes by level of education is potentially related to differences in the workers' productivity. [50]

Across those three dimensions, the most puzzling result is along the gender dimension. Table A.6 presents the results of Equation 1.1, estimated separately for men and women. The overall effect is a combination of negative effects on women's employment and positive effects on men's employment. The results suggest that private formal employment decreased by 3 percentage points for women; the most optimistic effect implied by these estimates is a reduction by .9 percentage points of the probability of being employed in a private formal firm, because of the policy. These results seem to be a combination of decreased employment for young female workers and increased employment for older female workers (see Figure A.3). However, for most outcomes, the increase for older

[50]The theoretical predictions on the heterogeneous effects are not clear; i.e. for education for instance, whether individuals with low education or those with high education might gain more from a wage subsidy. To the extent that the difference in employment by education reflects only differences in productivity, the wage subsidy could induce firms to hire now those who otherwise were not profitable for these firms, i.e. individuals with low education. This would imply a higher effect on individuals with low level of education. Similarly, firms could hire more workers with high level of education as the wage subsidy potentially help them expand their business, especially since some of these workers with high education were previously unemployed.

women seems to have started before the policy. Moreover, the trends in the outcomes seem rather erratic to support evidence of parallel trends, even though tests similar to the ones presented in the identification strategy section do not rule out that the two cohorts of women had similar trends before the policy. Given these concerns, the result that the wage subsidy decreased employment for young women should be read with caution.

The trends in outcomes look more similar for young and older male workers (see Figure A.4); the estimates for the effects of the policy, though positive, are not statistically significant. Given the precision of the estimates, positive effects as small as 3 percentage points can be ruled out as the increase of overall employment for young men; but similarly, these estimates cannot rule out decreases in overall employment of 2 percentage points.

Overall, the analyses by the worker's level of education and race do not reveal any heterogeneity. Table A.8 presents estimates of Equation 1.1 separately for individuals without and with a matric, the exit exam at the end of secondary school. These results are really similar and imply no differential effects with respect to the level of education. Table A.7 presents the results by race, separately for White and Black South Africans, the two major race groups in South Africa. The point estimates for Whites are larger. In particular, the estimates imply an increase by 1.6 percentage points of the probability of private formal employment for eligible Whites, because of the policy. These results are however imprecise and do not rule out null effects, nor effects equal to those estimated for Black South Africans. In conclusion, there is no heterogeneity of the effects with respect to race.

I then perform the analysis of the heterogeneity by main industry. As the data that I use for the analysis is a household survey, there is little information that would allow heterogeneity analysis with respect to firms characteristics. Respondents however give the industry in which they are employed, and the information is coded to have the 1-digit main industry, per the Standard Industrial Classification of industries. I investigate whether employment has changed in some industries in ways that would suggest an impact of the policies for firms in those industries. I use the main industries as possible outcomes, and estimate for each of the industries, equations similar to 1.1. Table A.9 shows the results of the estimation. Employment across all the industry has not differently changed for eligible workers compared to ineligible ones. This implies no differential effects of the policy by industry. One limitation of the heterogeneous analysis by industry as performed here is that the different industries as outcomes, are considered as independent outcomes. These outcomes are however related, as a worker is employed in one only of these industries. The evidence on the heterogeneity by industry as presented here is thus suggestive.

One additional analysis that I perform is to test whether results are heterogeneous with respect to local markets conditions. Firms that are in more dynamic labor markets may have better opportunities to hire young workers because of the policy. I use the unemployment rate of workers above 35 to capture the strength of the labor market, which I interact with the main variable of interest. The model estimated is thus a triple difference. The main variable of interest in this case is the triple interaction term. Results presented in Table A.10 shows that the triple interaction is not statistically significant, which implies that there was no heterogeneous effects with respect to local labor market conditions.

1.5.4 Discussion: Unsuccessful scale-up?

An important implication of the present analysis is the fact that the ETI does not replicate the effects of the small pilot of a wage subsidy studied by Levinsohn et al. (2014). The results in Levinsohn et al. (2014) suggest an increase in employment by 7 percentage points, an effect that is ruled out by most of the estimates that I found in my analysis. It is an open question as to what factors may have caused the difference in effects, as there are substantial differences in the scale and the details of the pilot wage subsidy and the ETI. The pilot was a voucher for workers to search for a job and lasted only 6 months. However, most of these comparisons would favor the ETI as the more impactful (e.g., high take-up in the ETI, vs. few claims of vouchers by the firms). The pilot was implemented as a voucher to increase search by workers, rather than hiring by firms, suggesting that the pilot targeted labor supply, while the ETI was meant to address labor demand. One specific finding in the pilot study is that that the effects seem to be driven by a reduction in the number of offers turned down. [51] This is a hint that job creation may have not been the driver of the results of the pilot, explaining why a wage subsidy might work at a small scale but not when implemented at scale. Such comparison of course is qualified by the inherent differences in the pilot wage subsidy vouchers and the ETI implemented at scale. Considering that these are both wage subsidies, this would suggest that there

[51] Levinsohn et al. (2014) noted that the reduction in the number of offers turned down is especially true in households with other employed members, who in general are the least likely to be in wage employment and when they do, have higher wages on average. This echoes the idea that the structure of households affect labor market outcomes in South Africa. The paper also frames this result as the consequence of a higher perceived probability of success for jobs for which they receive information through their family. This contradicts however the idea that turning jobs in the control group was indeed a choice to get higher paying jobs. The results do not provide clear evidence on earnings effects, and more importantly do not offer insight on job offers to get as to what the earnings for the jobs turned down are.

are some scale-up challenges, besides the difference in the design of the two policies. I provide below a discussion of considerations that may have prevented the replication of the effects of a wage subsidy once implemented at scale.

Banerjee et al. (2017) offer a framework to think through why a pilot RCT does not replicate when the intervention is scaled up. There are potentially six types of reasons that could explain the failure for an intervention to scale up: market equilibrium effects, spillover effects, political reasons, context-dependence, implementation challenges, and pilot site-selection bias.

Market equilibrium concerns are conceptually important for wage subsidies, e.g., with the possibility of changes in the composition of employment without no changes at the aggregate level, i.e., without any job creation. However, these concerns do not seem to be the explanation in the case of the ETI. There have not been fundamental changes in labor market conditions following the implementation of the policy, as the level of unemployment in the country has, if anything, remained the same as in previous years. Equilibrium effects that would be consistent with the results found in this paper would have been a wage decrease for ineligible workers, such that the ETI does not offer a relative advantage for young workers. Reports from the Labour Market dynamics in South Africa for 2014 indicated that wages have remained unchanged overall compared to their level in 2013, implementation challenges, which would suggest that market equilibrium effects are unlikely. [52]

Spillover effects in the context of a wage subsidy could be of two forms: substitution

[52]See the report here: `https://www.datafirst.uct.ac.za/dataportal/index.php/catalog/536/download/7390`.

effects (negative spillovers) or complementarity between young and older workers labor (positive spillovers). I have previously discussed that substitution may be unlikely in the South African context and given the regulations around the ETI. This implies that negative spillover effects are unlikely. Positive spillovers are one explanation that would be compatible with no relative effects, i.e., no improvement of young workers compared to older ones, but improvement for both eligible and ineligible workers. This could be justified by a complementarity between the labor of young workers and that of older workers. Overall in 2014, labor market outcomes have not improved, and in fact, growth overall has stalled in that year, implying that there was not a sudden rise in employment for the eligible workers. The question remains, though, as to what would have been the situation, were 2014 a period of strong growth, such that the ETI could have acted as a buffer. In fact, business cycles could well be an explanation as to why the wage subsidy did not matter, given the low growth in 2014. I cannot, however, provide conclusive evidence as to the relationship between business cycles and the effectiveness of wage subsidy in South Africa. My exploration using provincial differences in unemployment rates did not lead conclusive results (see Table A.10 and related discussion above). In all, it seems that spillover effects might be unlikely, but that the lack of results in 2014 could well be related to the business cycle in that year.

For political reactions, there could be, on the one hand, reactions on the workers side to the policy or a reaction by firms. A reaction from firms would, however, have been inconsistent with the higher-than-expected take-up of the subsidy that the government reported. It is not clear how a political response on the workers side could explain the result. It is the case that unions initially opposed the idea of a wage subsidy, and thus could have mobilized workers against this policy. It is not clear, however, how workers would follow such an indication to oppose the policy. Jobless workers are not covered by unions and thus may have no relation with the unions. Additionally, these workers would have likely taken an offer that a firm made to them.

Context dependence refers to the idea that a policy could be successful in a particular place but not somewhere else. Small-scale RCTs of wage subsidies have consistently found positive effects (see discussion in the literature review). This suggests that the positive effects in the wage subsidy pilot study in South Africa are an unusual feature that would fail to replicate elsewhere. This also relates to the next point, pilot site-selection bias. There is no indication that the pilot was implemented in areas that would favor the effectiveness of a wage subsidy, in ways that would not replicate when it scaled up to the whole country.

The last point in the framework by Banerjee et al. (2017) is implementation challenges, a point that would not be consistent with the higher take-up that the policy mentioned. If scaling up the wage subsidy was a challenge, then we would have expected that fewer firms take up the subsidy compared to the pilot. It is the contrary that the implementation of the ETI indicates.

From this illustrative discussion, it seems that one potential explanation as to why the ETI is not effective and does not replicate the findings from the small scale wage subsidy is that 2014 was a year of slow growth. It is not possible, however, to ascertain that implementing the wage subsidy in a different year would lead to different results. Overall though, this result suggests the importance of considering the scale of policies when discussing their effectiveness.

1.6 Conclusion

In this paper, I evaluate the impacts of the Employment Tax Incentive, a wage subsidy implemented by South Africa's government. The policy was designed to give incentives to firms for hiring young workers by lowering their labor costs. The wage subsidy was valuable to firms as it could amount to half the wage paid to a young worker. It was originally intended to last only two years, but it has been extended twice, and is currently scheduled to last until 2029. The basis of this renewal has been the higher than expected take-up of the subsidy by firms.

Using the ETI's age-eligibility restrictions and start time, I evaluate its impact on young workers by comparing individuals below and above the birth cohort cutoff for eligibility, before and after the start of implementation. My results show that young workers' employment did not increase as a consequence of the subsidy; I rule out effects as small as a 1 percentage point increase in overall employment.

Conceptually, the wage subsidy would have narrowed the productivity gap of young workers (due in part by the lack or mismatch of young worker's skills), and compensated employers for any initial training that these workers would have required. However, firms would have hired young workers only if the amount of the subsidy was large enough to offset this productivity gap as well as the training costs for employers. Firing regulations combined with the temporary nature of the subsidy (the subsidy could be claimed for up to two years for the same worker) would have further reduced the chances that firms indeed hired these young workers. Because the wage subsidy did not lead to improvements in employment of young workers, these other constraints (skills and labor regulations) should be considered for further policy actions, possibly in conjunction with the wage subsidy. Indeed, there have been some promising experiments on certifying skills (Carranza et al., 2019) and improving firms' information on labor regulations (Bertrand and Crepon, 2019). Even though these experimental studies find improvements in labor market outcomes, my analysis suggests that we should exercise caution when generalizing the findings of these experiments. Indeed, the present analysis suggest how the results of the wage subsidy implemented at national scale are different from those of the pilot study by Levinsohn et al. (2014).

One of the limitations of this study is the difficulty of drawing concrete policy recommendations. At the very least, my research implies that a wage subsidy, as currently implemented through the ETI, may not help improve labor outcomes for young workers. However, this paper does not provide clear guidance for how to improve the policy, or definitive evidence of why it was not effective. The discussion in section 1.3.1 provides a review of the factors in the structure of South Africa's labor market could potentially shape the effectiveness of a wage subsidy. My analysis has shown no heterogeneous effects by race or education, parts of the features discussed, or by industry, which would have given an indication of ways to improve the policy. However, it could be worth recognizing and exploring changes to the policy along these dimensions, e.g., subsidies to specific industries, to investigate pathways to a successful wage subsidy. Other policies that directly address the features presented in section 1.3.1, such as promoting access to quality education, could be considered as well.

Chapter 2: Internal Migrants' Ethnic Capital and Labor Market Outcomes in South Africa

The role of social networks in economic decisions has been recognized and has received a lot of attention from economists (see e.g., Granovetter, 2005; Jackson, 2011, for an overview). Individuals consider others' choices and benefit from others' resources when they decide what is in their own best interest. The role of social networks has been extensively studied in two areas: migration and labor markets. As the decision to migrate requires that the prospective migrant examines the potential gains in the destination against the costs of moving, social networks are potentially helpful in lowering costs by providing information to newcomers and in facilitating their integration at the destination. As for labor markets, there is ample evidence that individuals often find their jobs via friends, family, or other contacts (Ioannides and Datcher Loury, 2004). [1] Both theoretical and empirical evidence shows that social networks affect migration decision and labor market outcomes, with a particular focus on the integration of international migrants in the labor market in their destination. Internal migrants do not seem to have received the same attention, although they potentially face the same problems at their destination that international migrants face in their host countries.

[1] See section 2.1 on the literature review for more on the empirical evidence.

The question of the role of social capital for internal migrants' labor market performance is particularly relevant for developing countries. Internal migration is an important phenomenon in developing countries, as this migration is generally from rural to urban areas, and accounts for a sizable part of the growth of urban populations (Lall et al., 2006). Also, social networks have particular importance in developing countries due to imperfect and sometimes missing markets (Munshi, 2003).[2] In rural areas in developing countries, individuals typically rely on their friends and family members for loans, which are either cheaper compared to the high-interest rate loans offered by existing institutions, or the only choice in places where such institutions do not exist.[3] Individuals also use their social networks to insure themselves against risk, as suggested by the relative smoothness of consumption even though income is really volatile.[4]

South Africa is a good setting to look at the importance of networks for labor market outcomes. Burns et al. (2010) highlight the importance of networks in the labor market for the whole population in South Africa. The relevance of networks for employment opportunities in South Africa can be appreciated by looking at both job seekers and firms in the country. As suggested by the high share of discouraged workers, those who are unemployed are unlikely to look actively for jobs, which could also be related to the already high unemployment rate.[5] Data from the National Income Dynamic Survey (NIDS) re-

[2] As mentioned in Chauang and Shecther (2015), "social networks are arguably both more necessary and more frequently used by individuals in developing countries." The paper provides an overview of the importance of social networks in developing countries, with an emphasis on financial aspects and information sharing. Munshi (2014) provides a similar review on the role of social networks in developing countries.

[3] The World Bank reports that within Sub-Saharan Africa, about 2/3 of individuals who borrowed any money in 2017 borrowed from family or friends (World Bank, 2018).

[4] This holds both at the village aggregate level (Townsend, 1994) as well as within villages when specific networks (loans and gifts) are considered (De Weerdt and Dercon, 2006; Fafchamps and Lund, 2003; Udry, 1994).

[5] The unemployment rate was 26.5% in the fourth quarter of 2016 (South Africa Statistics: `http://www.statssa.gov.za/?p=9561`). For comparison, this rate was 13.9% in Nigeria in the third quarter

veals that 48% of employed individuals found their current job from a relative or a friend. Evidence from Standing et al. (1996) reveals that a large share of firms (41%) rely on their existing employees to recruit new workers. These statistics highlight the likely contribution of networks for labor market outcomes in South Africa. Indeed, Burns et al. (2010) find that, overall, the chances of employment of an individual are enhanced by 3 to 12% by his network.

In this article, I study the role of the ethnic capital on the labor market performance of internal migrants in South Africa. Social networks refer to actual contacts that individuals have and to which they turn or might turn to request information or services. As individuals are more likely to interact with individuals that have the same origin, ethnicity, or race as them, the ethnic group may be a proxy for potential contacts that and individual has. Borjas (1992) studies the extent to which individuals benefit from the ethnic group of their parents for their education and labor market choices, defining these resources available from the ethnic group as the ethnic capital. Similarly, I define the ethnic capital of internal migrants as resources that they might get from individuals of the same ethnic group, and I study the extent to which they benefit from this capital on the labor market. [6]

This article contributes to the literature in two dimensions. First, I focus on internal migrants, unlike most studies that focus on international migrants. Internal migrants move

of 2016 (Nigerian Bureau of Statistics: `http://www.nigerianstat.gov.ng/report/481`). The high share of discouraged workers is a sign that individuals who are willing to work resort to passive methods of job search instead of actively being on the market.

[6] Madhavan and Landau (2011) notes that, because population in African cities are comprised mainly of people that migrated at different period times there and usually not long ago, there is a fluidity in the concept of new arrivals of migrants and established populations. This calls into questions whether the concept of social and ethnic capital, as described in the literature on international migration (e.g., Munshi, 2003) can be applied in for internal migration for African cities. This implies as well however that it is an empirical question as to whether ethnic identity is a factor that affects labor market outcomes, potentially through formation of ethnic capital. In the end, the results discussed in this chapter should be considered with a critical perspective on the notion of ethnic capital.

to locations where finding a job is uncertain, and where they may lack the information necessary to get the best opportunities on the labor market. As such, they may use their social contacts to compensate for this lack of information. Given the prevalence of the use of social networks in South Africa for labor markets, as suggested by Burns et al. (2010), it is an interesting question to see how internal migrants' labor market outcomes are related to their social networks' characteristics.

In addition to looking at internal migrants, I study occupation choice, not just the probability of employment. Little attention has been paid in the empirical economic literature to the importance of social networks for occupational choice. Two articles that consider occupational choice are Patel and Vella (2013) and Guerra and Mohnen (2016). Patel and Vella (2013) look at the probability that an incoming migrant to the USA chooses the occupation that is the most popular among migrants from the same country. Guerra and Mohnen (2016) use asymmetries in peers' influences to identify the effect that social interactions play on the occupational choice in 19^{th} century London. In this article, I assess how occupation choices by internal migrants are related to the characteristics of individuals of the same origin and the same ethnic group who have established themselves previously in the migrant's current place of residence.

The structure of the article is as follows. I first give a brief overview of the literature on social networks and labor markets. I then present the methodology used in this article. After presenting the data and a description of the sample of internal migrants that I use, I discuss the results of the estimations.

2.1 Literature

Ioannides and Datcher Loury (2004) provide an excellent overview of the role of social networks in the labor market, both from sociological and economic perspectives. Montgomery (1991) and Calvo-Armengol and Jackson (2004) provide theoretical frameworks to understand the role of networks in job search and their implications for the members of the network. Montgomery (1991) starts from the perspective of employers who do not observe the ability of potential workers. They might assume, however, that the ability of current workers reflects that of the members of their networks, or that those workers have better information about the skills of their peers. In either case, it is beneficial for employers to rely on the network of their current workers to hire future employees. This implies that an individual who has a larger or relatively more employed network is more likely to find a job. Calvo-Armengol and Jackson (2004) start from workers' perspective by considering a network that individuals rely on for information about job opportunities, taking the network as given. The model generates a positive correlation between the employment status of one's network and his own employment probability. The dynamics of the model highlights however a negative short term effect of a large network, as unemployed individuals compete for job openings that are passed onto them by individuals employed in the network.

Social networks are used for job search in part because of the lack of other sources of information (Chauang and Shecther, 2015). As such, individuals who don't have access to information about job openings are more likely to use their social networks to get a job. Migrants form one such category of individuals, especially just after their arrival in their destination. They may lack information about job availability at their destination and rely on their social network for such information. The study of the role of social networks for migrants has also been in part motivated by the observation that international migrants choose to settle in a few destinations within the host country.[7] This concentration follows an ethnic pattern as migrants from the same country are over-represented in specific places. The literature on information networks presented above would imply that the presence of social contacts could lead to positive benefits. However, the lack of contacts with the native population could prevent the acquisition of the country-specific human capital and hurt their labor market success. Even though such perspective could lead migrants not to choose to live in enclaves, a theory of ethnic goods suggests that migrants would still make this choice if ethnic goods are available in those enclaves at a lower cost, because of economies of scale (Chiswick and Miller, 2005). The effect of living in an ethnic enclave on labor market outcomes is thus an empirical question.

Although the high percentage of individuals finding jobs through their network is suggestive that networks might have an effect on labor market outcomes, the endogenous formation of networks plagues the estimation of the effect of social networks or social contacts. People choose their networks and this choice is related both to their unobservable personal characteristics as well as those of the group they choose to belong to. Edin et al. (2003) provide a theoretical model of the choice of (the ethnic concentration of) the place where migrants live. The model illustrates how selection into a network can lead to

[7] Borjas (1999) reports that in the US, almost three-quarters of immigrants resided in only six states in 1998. Edin et al. (2003) states that "the share of the foreign-born population living in the three largest metropolitan areas outstripped the share of the native population by 18 percentage points in 1997" in the case of Sweeden. In the case of Denmark (Damm, 2009), 71% of nonrefugee immigrants live in the capital area, where 26% of the overall population lives.

bias in estimations. If high-skilled individuals choose large networks, the effect estimated using OLS would lead to an overestimation of the effect of the network, as high-skilled individuals would be more likely to find a job anyway.

Besides the endogenous choice issue, most data available to researchers, especially for labor market outcomes, do not include actual contacts between individuals. As such, researchers have to decide both on how to infer network connections and how to deal with endogenous sorting into networks. Studies generally use one of two proxies to infer who individuals are likely to have as contacts. The first one is geographic proximity. The second is the ethnicity or the language spoken. In the case of international migrants, the country of origin is used to denote ethnicity. The identification assumptions to account for the endogenous sorting problem are related to these choices for social network proxies. In particular, studies using geographic proximity usually choose units of residence such that other confounding factors can be ruled out, whereas studies relying on the ethnic enclaves try to account for sorting by using quasi-experiments.

Bayer et al. (2008) is a leading example of the use of geographic proximity for the study of the effect of social networks on labor market outcomes. The study assumes that interactions between individuals are very local, and shows that two individuals who live in the same census block are more likely to work at the same place than neighbors living slightly further apart. Andersson et al. (2014) applied the same methodology but focus on international migrants, and find that migrants are more likely to work with a neighbor if they live in an immigrant neighborhood.

Studies on migrant networks' effects often rely on refugee dispersal programs for quasi-experiments. Edin et al. (2003) use the Swedish refugee program which places refugees across the country, in a quasi-random way, conditional on information available to the refugee office at the time of placement. This leads the refugees to be exposed to a network that is exogenous to their own characteristics. Damm (2009) uses a similar approach with the Danish refugee dispersal program: the stock of refugees placed since the beginning of the program induces exogenous variations in the size of the enclave in which refugees are first placed. The approach in Beaman (2011) is similar to Damm (2009), but she uses the annual flow of refugees placed by one agency in the US to obtain exogenous variation in the network size. Unlike those studies, Munshi (2003) uses variation in historic rainfall in the community of origin to instrument for the size of the network of Mexican migrants in the US.

Edin et al. (2003) and Damm (2009) find negative sorting in enclaves, meaning that low skilled individuals are the ones that choose most often to live in large immigrant enclaves. However, after correcting for this selection, living in a larger enclave is associated with higher earnings. In the Swedish case, the gains from living in an enclave are present only for low skilled workers, while high skilled workers also benefit from enclaves in Denmark. Munshi (2003) finds that an exogenously larger share of the network in the destination increases employment in general, and in particular employment in better paid nonagricultural jobs. The effects estimated are heterogenous, with respect to gender, education, time since arrival and age, networks benefiting more to those for whom it would have been difficult to find a job otherwise. Beaman (2011) differs from the other studies in that it takes into account the dynamics of the networks following the theoretical work by Calvo-Armengol and Jackson (2004). The model suggests a negative effect of a larger network in the short run, as newly arrived migrants who are unemployed compete for job openings. Beaman (2011) finds that the number of migrants arrived in recent years decreases the chances of employment for newcomers, whereas the number of established migrants increases these chances.

The literature has mostly focused on the probability of getting a job and earnings as the main outcomes of interest. Less importance has been given to occupation choice, which seems to have been more studied in the sociology literature, as mentioned in Ioannides and Datcher Loury (2004).[8] Hofmeyr (2010) studies the question in South Africa by defining occupational enclaves in the manufacturing sector as occupations where a specific ethnic group is over-represented and finds network effects in the chances of working in an ethnic enclave. However, this approach doesn't allow to detect which occupations exhibit higher network effects, and it requires that one chooses a threshold for the specialization index, above which one occupation is considered an enclave for a particular network group. The study also only covers the manufacturing sector. Two papers related to my current approach are Patel and Vella (2013) and Guerra and Mohnen (2016). Each looks at the probability of getting jobs in specific occupations and demonstrate that networks play a role in such probabilities.

Studies on social networks and labor market outcomes related to developing countries have focused on the reasons why firms might find it beneficial to use referrals to fill their vacancies. Three reasons are usually mentioned (Fafchamps and Moradi, 2015): screening of potential candidates, saving on search costs and monitoring (or solving the

[8]The general question here is whether there is heterogeneity in the jobs that individuals get and how social networks is related to those different type of jobs. The industry can be used to study the different type of jobs. Considering the industry has the appealing feature that a worker hears about opening in his firm or other similar firms and pass the information to members of his network. However, looking just at the industry doesn't allow the comparison of the position that the individuals have in their respective jobs. I can do this looking at the occupation choice.

moral hazard problem). In experimental games, Beaman and Magruder (2012) find that participants with high ability could refer other high ability candidates but choose to do so only when their bonus for the referral is tied to the performance of the person referred, suggesting that there is a screening model in play. Fafchamps and Gubert (2007) do not find evidence for the screening hypothesis, as in the case of the Gold Coast Army, the individuals referred were not less likely to desert compared to those recruited by other means. Heath (2018) finds support for the moral hazard model in the case of the garment industry in Bangladesh. When a referral is made, firms tie increase in wages of the referrers to the performance of the worker referred, inducing the latter to provide more effort, even though the workers referred are not on average of better quality than workers not recruited through referrals.

2.2 Methodology

The set-up is a country that has J local municipalities, denoted $j = 1, \cdots, J$. There are in total $J = 230$ such municipalities in South Africa. Considering the different languages and provinces of South Africa, I define $k = 1, \cdots, K$ groups among which each South African belongs to uniquely one. A group k is the set of all individuals who were born in a specific province and who share a primary language. There are 11 language groups and 9 provinces in the country, leading to $K = 99$ such groups, which I refer to as ethnic group in this article. An ethnic group is defined by birth and language and spreads across many municipalities, through migration. The assumption is that a migrant benefits from the information or resources of individuals from his ethnic group who live in his current municipality of residence. [9]

Since I am interested in labor market outcomes, the individuals that I consider are working-age migrants (15-64 years old) who are in the labor market (workers and individuals looking for jobs) and discouraged workers. The definition excludes individuals still in school and those that are neither looking for a job nor willing to work.

Employment and ethnic capital

Consider an internal migrant i living in a local municipality j and who belongs to an ethnic group k. His employment probability depends on his characteristics and the characteristics of his current local municipality of residence. Personal characteristics include the level of education, age, and sex. Characteristics of the place of residence capture the local labor market's conditions. I denote by y_{ijk} the binary variable that takes the value 1 if the migrant i is working and zero if not. This decision can be modeled in a latent variable formulation as follows:

$$y_{ijk} = 1\, if\ y^*_{ijk} = X_i\beta + W_j\gamma + u_{ijk} \geq 0\ ;\ 0\ otherwise \tag{2.1}$$

To the extent that a migrant benefits from his ethnic group in finding a job, the model should account for the characteristics of the group. There are potentially two ways to incorporate this feature in the labor market participation equation above. On the one hand, the average employment rate (which is the average of the binary dependent variable

[9] One possibility would have been to consider the last province of residence in building a proxy for social capital. However, the previous residence, potentially different from their place of birth, is the result of the individual's choice, which is thus endogenous. It is certainly more prone to individual's choice than the place where they were born. For these reasons, I chose the place of birth instead.

in equation 2.1) of the social group (in this case of individuals of the same ethnic group living in the same municipality as the migrant) could be incorporated as an explanatory variable in the model. This approach seems, however, not internally coherent, due to the different nature of the latent variable and the binary variable. On the other hand, the model could include the leave-out-mean of the latent variable on the right-hand side of the equation. This leads to the following formulation:

$$y^*_{ijk} = X_i\beta + W_j\gamma_1 + \gamma_2 \frac{1}{N_{jk}-1} \sum_{i' \in j, i' \in k} y^*_{i'jk} + u_{ijk} \tag{2.2}$$

N_{jk} represents the size of the ethnic group k living in the local municipality j. Equation 2.2 has the features of a spatial model, where the dependence across individuals happens within an ethnic group and a specific municipality. I propose here to estimate reduced form parameters of the model in equation 2.2. Specifically, summing the latent variable in equation 2.1 over all individuals, natural independent variables in equation 2.2 are the leave-out-means of the characteristics of the ethnic group in the destination municipality. The model should also account for error terms at the group level, as they show up while averaging the latent variable. However, including an error term or fixed effects at ethnic group X municipality level would prevent the identification of parameters at the group level due to multicollinearity. I instead introduce fixed effects that are additive in the ethnic group component and the municipality component. The model to be estimated reads:

$$y_{ijk} = \ 1 if\ y^*_{ijk} = X_i\beta + W_j\gamma_1 + \bar{X}_{jk}\delta + \mu_j + \lambda_k + u_{ijk} \geq 0\ ;\ 0\ otherwise \tag{2.3}$$

μ_j is the unobserved local municipality effect, and λ_k is the unobserved ethnic group effect. The parameter of interest, in the end, is δ. It reflects the extent to which a migrant's behavior in the labor market is related to the average characteristics of others in his ethnic group that live in his local municipality. Evidence of the importance of the group is provided by the statistical significance of this parameter. $\bar{X}_{jk}$ represents ethnic capital. The introduction of the fixed effects at the ethnic group level and the residence level implies that the parameter δ is identified by the variation within the ethnic group and within the local municipality. The fact that those two fixed effects enter in an additive way allows for the remaining variation that is used for identification. An important feature, especially related to the introduction of local municipality fixed effects is that, the identification of the parameter of interest controls for the endogenous choice of destination by the migrants, which may have biased the results.

There remain, however, some challenges to identification. Migrants' choices are potentially related not only to the local municipality's characteristics (local labor market conditions) in absolute, but also related to the extent to which members of their ethnic group perform in this particular labor market. The local municipality fixed effect accounts for the global (or common to all groups) conditions of the local market. This problem could have been solved using fixed effects at the ethnic group - municipality level (e.g., ω_{jk}), but these fixed effects would be perfectly collinear with ethnic capital, $\bar{X}_{jk}$. My study assumes this away through the additive error term (ethnic group and municipality components). Assuming that unobservable characteristics at the ethnic group - municipality level are additive in the two components (i.e. $\omega_{jk} = \mu_j + \lambda_k$) would allow identification of the ethnic capital parameters. Identification would also be achieved if the additive term captures most of the variation of the specific component that explains the choice of a municipality by members of a particular ethnic group. To partially address the concern that I do not control for unobserved effects at the ethnic group - municipality level, I take into account the potential correlation of unobservables at the said level. I do this by clustering all standard errors at the ethnic group - municipality level.

I conceptually assume that individuals are in one of the three following groups: recent internal migrants, established internal migrants and non-migrants. [10] The analysis estimates the model in equations above for recent migrants. When I use the distinction between internal migrants and non-migrants, I refer to the experience of recent migrants relative to all others. My measure of ethnic capital $\bar{X}_{jk}$ is the average characteristics of established migrants aged between 25 and 64 years old in a given municipality. It is plausible that migrants would benefit from information from long-established individuals as opposed to more recent arrivals. This is also motivated by the fact that, in the previous literature, the effects of information networks are found in the long term, but not in the short term. In the end, my evidence for the importance of the ethnic group in the employment probability is the extent to which recent migrants' employment outcomes are explained by the characteristics of long-established individuals of the same ethnic group.

Occupation and ethnic capital

One contribution of this article is the emphasis on the occupation choice of migrants (see sections 2.3 and 2.4 for more details on the occupation variable), not just the probability of employment. I extend the analysis presented above to the explanation of the occupation

[10] Refer to the data section for more explanation on the implementation of this distinction.

choice and its relationship with the ethnic group's characteristics. The goal here is to see how the likelihood that an internal migrant chooses a particular occupation varies with the characteristics of individuals of his ethnic group who live in his current municipality. Suppose that there are $q = 1, \cdots Q + 1$ occupations that an individual can choose. Define the variables $Y^q_{ijk} = 1$ if a migrant chooses the occupation q and 0 otherwise. Similar to the framework presented above, one could model the choice of a particular occupation q as follows:

$$Y^q_{ijk} = 1 \; if \; y^{*q}_{ijk} = X_i\beta^q + \overline{X}_{jk}\delta^q + \mu^q_j + \lambda^q_k + \varepsilon^q_{ijk} \geq 0 \tag{2.4}$$

My main interest is in the relationship between migrant occupation choice and the characteristics of his ethnic group. In the formulation above, δ^q represents the reduced form parameters of the ethnic group's effect on the occupation choice. A significant δ^q would be my evidence that migrants' occupation choices are influenced by ethnic capital in their municipality of residence.

In order to compute the marginal effects from equation 2.4, I use a SUR model where each equation represents a linear probability model obtained from equation 2.4. Specifically, I consider the following model:

$$P[Y_{ijk}] = E(Y^q_{ijk}) = X_i b^q + \overline{X}_{jk} d^q + n^q_j + m^q_k, q = 1, \cdots Q \tag{2.5}$$

This system of equation is estimated with a SUR model. The advantage of a SUR model over equivalent linear probability models is that it takes into consideration the correlation among the error terms that are obtained in each equation. Compared to marginal effects that are obtained using a multinomial logit model, the present approach doesn't require the assumption of independence of irrelevant alternatives to be satisfied. The estimation of the marginal effects using SUR models has been suggested by Zellner and Lee (1965) and Heckman and Macurdy (1985).

My analysis is agnostic as to whether the results are driven by labor demand or labor supply. Ethnic capital, as defined throughout this chapter, can move both the supply and the demand for labor, in an analogy to exogenous variables or instruments that would be able to move only one of the two. Consistent with the idea that ethnic capital may provide information to migrants on opportunities, ethnic capital can move labor market outcomes by increasing the supply of labor. Consistent with the idea that firms may use performance by workers as a proxy for the performance of prospective employees (as initially expressed in Montgomery (1991) analysis), firms could use the ethnic group of a migrant as a proxy for his or her ability. Still, similarly, this could induce firms to use workers to recruit more employees by explicitly asking them to make referrals. In the end, what would be problematic would be that the results are driven entirely by labor demand, but not through a request for referrals. If firms discriminate by using average characteristics of the ethnic group, this will make ethnic capital significant in the equation estimated using the methodology above. However, unless the way this discrimination operates is through firms using their workers for referrals, the mechanism would not be ethnic capital per se. The only way that my analysis tries to account for this is with the use of the additive fixed effects to control for performances of an ethnic group within a specific municipality. My results are thus to be interpreted within the limits of this methodology. The concerns around the distinction between labor supply and labor demand are relevant in line with what my methodology allows me to do in terms of identification.

2.3 Data

The data that I use in this study is the publicly available 10% sample of the 2011 census in South Africa.[11] The census data contains information that allows inferring both the employment and migration status of interviewees. The measures of the employment status include the one-digit International Standard Classification of Occupations (ISCO) code for occupation. Regarding migration, the census data reports whether the interviewees have been living at the same place since 2001. I define migrants as individuals who have moved after 2001 to a new local municipality and were still living there in 2011. Thus, migrants considered in my study are relatively recent arrivals; they have migrated less than 10 years ago to their current place of residence. I consider in my main sample analysis all Black South Africans (or by convention, ”Africans”) internal migrants aged between 15 and 64 years.

The reason for considering only Africans in the present study is that labor market behavior and migration patterns are race-specific in South Africa (see e.g., Garlick et al., 2016). Africans represent 80% of the population. Migration among Africans is linked to the political and economic history of South Africa. As early as Europeans discovered mines in the country, they installed a structure that constrained access to land for Africans, making them more likely to turn to labor in the mining sector. Additional measures that constrained movement made permanent migration to mining and urban areas impossible.[12] These restrictions were later removed, especially with the abolition of apartheid. When

[11] The dataset has been obtained from the data portal of the University of Cape Town. The link to the dataset: `https://www.datafirst.uct.ac.za/dataportal/index.php/catalog/485`.

[12] Pass books were required outside homelands, where Africans were living, besides when going out to work, and the length of the stay for work-related reasons was limited.

those restrictions on people's movement were in place, a system of circular migration was observed: people would migrate to the place where they were working, staying there alone while working, as they were not permitted to move with their families. They had to go back to their place of origin, the homelands, once the contract was finished. This circular migration was expected to disappear and lead to more permanent migration following the end of apartheid. However, Posel and Casale (2003) show that there is evidence that the circular migration of Africans is still ongoing in South Africa.

These features of Africans' migration patterns raises questions on how adequate the different datasets on South Africa are to study issues related to migration, and how they affect the results that are obtained. The census data report information on individuals at a specific time. Migrants are identified as individuals that are living at a place where they were not living 10 years ago. This definition includes all the permanent migration that occurred in the last 10 years and still ongoing, but misses individuals that are in their place of origin, returning from a spell of temporary or circular migration. Such individuals would be classified as non-migrants, under-estimating the true count of the migrants in the economy. However, as my interest is in the behavior of migrants and not the selection process, I need, for my results to be unbiased, that individuals that are returning to their place of origin at the time of the census data collection to be not systematically different from other migrants. This should not be an issue since the selection here is mostly related to the time of collection of the data, which does not present any particular feature that one should worry about in the present context of the interaction between ethnic group and labor market outcomes.

There is a separate concern, which is the presence at the origin of "permanent return migrants" who, at the time of the collection of the census data, do not intend to return to their previous migration destination. These return migrants do not provide a bias in the definition of migrants for the purpose of this study but constitute a channel through which individuals can get information on the place to migrate, and potentially kicks into how ethnic capital at the place of origin helps explain the choice of the destination. Ethnic capital from these return migrants at the place of origin could be correlated to the ethnic capital from established migrants. This would then be an issue for the estimation in this chapter, as it will contribute to the endogeneity of ethnic capital at the place of destination. There would be an issue whether these return migrants are positively or negatively selected. Both successful migrants could share their positive experiences, just as unsuccessful ones could offer cautionary tales about the migration experience. Ideally, characteristics of these returns would have helped in an equation into the selection of the destination of migrants. In the end, though, the results of my analysis should be interpreted within the limitations of how my methodology achieves identification.

2.4 Internal migration in South Africa

Describing internal migrants

The share of internal migrants in the labor force is 11%. Black South Africans represent the major group in the population, but the second most mobile when it comes to internal migration. Indeed, only 11% of economically active Black South Africans have migrated, whereas 20% of White South Africans have also migrated (see Table 2.1 for more details). Internal migrants tend to be better educated than the general population. Considering the

Table 2.1: Race and internal migration incidence

	Share in the population	Internal migration incidence
Black (African)	76.97	10.73
Coloured	10.03	5.03
Indian or Asian	2.71	9.82
White	10.10	19.83

Notes: Data: South Africa 2011 census.

whole labor force in the country, the share of individuals holding a post-secondary degree in the whole population of working age is 14%, whereas 23% of the internal migrants hold a post-secondary degree. This comparison also holds within racial groups. For Africans (see Table 2.2), 17% of Black South Africans who migrated have a post-secondary degree, as compared to 10% only of those who did not migrate. Internal migrants are younger, more likely to be male and less likely to present physical disabilities (Table 2.2).

Among Africans who are in the labor market, internal migrants are on average 31.6 years old, while those who did not migrate are 35 years old. They have been living in their current place for three years on average. The proportion of women among migrants is 45%, whereas women represent 52% of non internal migrants. Internal migrants are less likely to present disabilities. As an illustration, 9% of internal migrants have vision issues compared to 11% of those who did not move. Although data on the exact income or wages are not available, self-reported gross income (including all sources of incomes) of individuals are reported in brackets. A comparison of the distribution of income shows that internal migrants tend to earn more, as the distribution of their income is skewed to the right compared to that of non-migrants (see figure B.1).

One striking aspect while comparing migrants to those who have not moved is that 60%

Table 2.2: Comparing migrants and non-migrants

	Non-migrants	Migrants	p-value of difference
No schooling	0.052	0.027	0.000
	(0.222)	(0.162)	
Some primary	0.115	0.072	0.000
	(0.320)	(0.258)	
Completed primary	0.049	0.036	0.000
	(0.216)	(0.187)	
Some secondary	0.388	0.352	0.000
	(0.487)	(0.478)	
Matric	0.312	0.341	0.000
	(0.463)	(0.474)	
Post-secondary	0.084	0.172	0.000
	(0.277)	(0.378)	
Person is male	0.489	0.548	0.000
	(0.500)	(0.498)	
Age in completed years	35.223	31.586	0.000
	(11.586)	(9.130)	
Any difficulty in Seeing	0.113	0.090	0.000
	(0.317)	(0.286)	
Any difficulty in Hearing	0.024	0.021	0.000
	(0.154)	(0.142)	
Any difficulty in Communication	0.007	0.006	0.011
	(0.082)	(0.076)	
Any difficulty in Walking or climbing stairs	0.021	0.014	0.000
	(0.142)	(0.117)	
Any difficulty in Remembering / Concentrating	0.026	0.020	0.000
	(0.160)	(0.140)	
Any difficulty in Self Care	0.008	0.005	0.000
	(0.090)	(0.070)	
N	916,395	123,885	

Notes: Data: South Africa 2011 census. Sample: Africans aged 15-64, active in the labor force or discouraged workers.

of migrants are the head of the household in which they live. This is explained by the structure of the households, as those headed by internal migrants tend to be smaller: households headed by internal migrants have, on average, 2.3 persons (compared to 3.5

for households that are headed by non-migrants). 43% the households with an internal migrant head are a single-member household, and 46% other households have just between one and three other persons. In addition, internal migrants that are head of households are 10 years younger than those who are not migrants. This suggests that people moving within the country are less likely to do so as a household, but more likely to be individual persons who move alone, and later form their households. This echoes the results of Garlick et al. (2016) that show that only one third of individuals involved in internal migration between 2008 and 2016 were in households that moved as a whole.

Migrants' ethnic group

The two features that define proxies for networks in my study are the province in which individuals were born and the language that they speak. Table 2.3 presents the 99 different ethnic groups, as defined in section 2.2. These groups vary considerably in their size. Out of the 99 different groups, the five major ones make up close to 60% of all internal migrants, whereas the 30 smallest cells provide less than 1% of the internal migrants. Tables B.1 and B.2 provide the rate of migration given the language spoken and the province of birth. The migrants were primarily born in Eastern Cape (24%), Limpopo (22%) and Kwazulu-Natal (17%). They primarily speak IsiXhosa (24%), IsiZulu (22%), and Sepedi (14%). There is a strong correlation between those two characteristics. 86% of the IsiXhosa were born in Eastern Cape, whereas close to 70% of IsiZulu were born in Kwazulu-Natal and 80 % of Sepedi in Limpopo.

Regarding the destination that they choose, migrants tend to move mostly to the wealthiest

Table 2.3: Size of the language and province of origin groups

Language spoken	Province of birth									
	WC	EC	NC	FS	KN	NW	GA	MP	LI	Total	Percentage
IsiXhosa	807	25,363	272	613	723	356	1,034	158	111	29,437	23.8
IsiZulu	89	848	96	625	19,033	319	3,307	2,728	593	27,638	22.3
Sepedi	45	129	33	120	110	272	1,251	1,456	14,297	17,713	14.3
Setswana	50	121	1,534	592	185	7,524	2,081	445	697	13,229	10.7
Sesotho	77	1,145	90	6,180	275	596	2,161	510	830	11,864	9.6
Xitsonga	21	224	16	37	51	171	565	1,433	4,337	6,855	5.5
Tshivenda	16	18	7	11	17	34	264	52	4,288	4,707	3.8
English	134	926	79	332	778	391	776	341	707	4,464	3.6
SiSwati	9	44	17	17	83	35	170	2,559	140	3,074	2.5
IsiNdebele	16	92	36	40	232	146	483	1,367	498	2,910	2.3
Afrikaans	185	318	247	297	114	179	386	118	150	1,994	1.6
Total	1,449	29,228	2,427	8,864	21,601	10,023	12,478	11,167	26,648	123,885	
Percentage	1.17	23.59	1.96	7.16	17.44	8.09	10.07	9.01	21.51		

Notes: Data: South Africa 2011 census. Sample: African migrants aged 15-64, active in the labor force or discouraged workers. The provinces are : WC : Western Cape, EC : Eastern Cape, NC : Northern Cape, FS : Free State, KN : Kwazulu-Natal NW : North West, GA : Gauteng, MP : Mpumalanga, LI : Limpopo.

and most populated province of the country, Gauteng (see Table B.3). This province contains Johannesburg, South Africa's largest city. Gauteng is the destination of 42% of internal migrants. Kwazulu-Natal, including the main city of Durban, accounts for the destination of 13% of the internal migrants. The less common destinations are Northern Cape (2%) and Free State (4%). Migrants primarily choose to go to Gauteng or to a local municipality that is in the same province. For example, Eastern Cape is the province where most migrants were born. Of those, 28% choose to go to a geographically close province, Western Cape, while 27% stay in the same province. 22% choose to go to Gauteng. In total, 40% of those who migrate have chosen to stay in the same province as the one they were previously living in.

Labor market outcomes of internal migrants

The definition of occupational choice builds on the major group of occupations, the one-digit level of the International Standard Classification of Occupations (ISCO). One particularity of South Africa's labor market is the high unemployment rate and the high share of discouraged workers. These high numbers have been at the source of questions as to whether part unemployment is voluntary, and the factors that explain why a large share of the jobless choose not to actively search for a job. To take into account this feature of the labor market, instead of studying the choice of occupation conditional on being employed, or conditional on being active in the labor market, I adopt a definition of occupation choice that augments the original variable by two categories for unemployed and discouraged workers. Table 2.4 presents the distribution of this variable. Internal migrants are less likely to be unemployed or discouraged workers compared to the whole population. 60% of the whole active population is employed, but this share is 75% within internal migrants.

This description of migrants characteristics reveals that migrants are selected positively, as they are better educated, and, in part due to this selection, have better labor outcomes. However, other factors come into play, as the opportunities available at the destination, and ethnic capital, which is the factor that my study seeks to highlight. This process can potentially be used as evidence as that migration is motivated by economic reasons.

Table 2.4: Occupation choice distribution

	Whole population	Internal migrants	Black internal migrants
Manager	4.3	7.25	4.21
Professional	3.88	6.7	4.79
Technician	5.83	7.55	6.23
Clerk	7.05	10.06	7.98
Sales and services	9.47	12.22	12.19
Skilled agriculture	0.51	0.65	0.5
Craft and related traders	7.05	8.67	8.85
Plant and machine operators	4.04	4.66	4.86
Elementary workers	10.16	10.68	11.92
Domestic workers	5.97	6.12	7.01
Unemployed	30.75	20.74	25.57
Discouraged workers	11.00	4.71	5.88

Notes: Data: South Africa 2011 census. Sample: individuals aged 15-64, active in the labor force or discouraged workers.

Labor market outcomes and ethnic capital

Descriptive evidence of the importance of the ethnic group for the employment probability comes from the observation the employment rate of recent internal migrants is correlated with the employment rate of established migrants of the ethnic group that have been residing for a longer time in the same local municipality. Table 2.5 presents the average employment rate of recent internal migrants across language groups, along with the average employment rate of established internal migrants of the same ethnic group. Recent migrants speaking primarily Afrikaans have the highest employment rate (71%). At the same time, Afrikaans established migrants also experience the highest employment rate (81%). Similarly, recent migrants speaking IsiXhosa have the lowest employment rate (57%) and establisehd migrants speaking IsiXhosa have the lowest employment rate among established migrants. The details for all language and origin groups are provided as a graph (figure B.2). This graph suggests, like the comparison between IsiXhosa and Afrikaans speakers, that a high employment rate of established migrants is associated with a high employment rate of migrants of the same ethnic group.

Table 2.5: Average employment rate by language groups (recent migrants and established migrants)

	Migrants				
	Recent		Established		
Language groups	Mean	s.d	Mean	s.d.	N
Afrikaans	0.71	0.17	0.81	0.39	2,002
English	0.68	0.15	0.78	0.42	4,484
IsiNdebel	0.60	0.14	0.72	0.45	2,917
IsiXhosa	0.57	0.12	0.64	0.48	29,511
IsiZulu	0.57	0.12	0.68	0.47	27,712
Sepedi	0.61	0.13	0.68	0.47	17,777
Sesotho	0.62	0.13	0.71	0.45	11,897
Setswana	0.64	0.15	0.74	0.44	13,274
SiSwati	0.62	0.13	0.71	0.45	3,086
Tshivenda	0.69	0.13	0.71	0.46	4,731
Xitsonga	0.62	0.13	0.69	0.46	6,868

Notes: Data: South Africa 2011 census. The averages are computed from the migrants' point of view. The average for established migrants is thus the employment rate of established migrants of the same ethnic group, weighted by the size of the migrants' group.

I provide similar evidence for the occupational choice. Just as for the employment rate, I plot the share of each occupation among recent migrants against the same share for established of the ethnic group (figure B.3). Overall, the share of recent migrants choosing a particular occupation tends to rise as the share of established migrants choosing this occupation rises. These results provide suggestive evidence that the more employed or represented in an occupation the ethnic group of a migrant at the destination is, the more likely the migrant is to get employed and to make similar occupation choices. Such correlation could just be due to confounding factors, such as the local labor market conditions that affect both recent migrants and established migrants. In the next section, I provide evidence that the benefits that migrants get from their ethnic group persist even after controlling for some of the confounding factors.

2.5 Results

2.5.1 Employment and ethnic capital

Table 2.6 presents the results from the estimation of equation 2.3. The equation is intended to produce reduced form estimates of the relationship between a migrant employment probability and the characteristics of his ethnic group. As expected, individual characteristics matter for the chances of being employed: a higher education level and being a male are associated with higher chances of employment, as well as being a man. There is also a positive relationship between age and the chances of being employed. Conditional on individual characteristics, ethnic group characteristics significantly explain the chances of employment of an internal migrant. The F-test shows that those characteristics are jointly significant (Table 2.6, specification 2), even after controlling for unobservable factors at the ethnic group and local municipality of residence level (specification 3).

The effect of the distribution of education at the ethnic group level shows that internal migrants benefit from both ends of the distribution, i.e. both from the share of their ethnic group with high education as well as the share with no or little education. These results are in comparison with the share of the ethnic group that has a secondary education. [13] On

[13] As the sum of the shares of low education (no education or some primary), secondary and post-secondary degree holders equal to one, those three variables cannot be added to the model. As such, I

Table 2.6: Employment and ethnic group characteristics

	(1)	(2)	(3)
Migrant own characteristics (dummy variables)			
No education or some primary	-0.016	-0.025	-0.031
	(0.007)**	(0.006)***	(0.005)***
Post-secondary	0.206	0.206	0.201
	(0.005)***	(0.005)***	(0.005)***
Person is male	0.171	0.168	0.167
	(0.006)***	(0.006)***	(0.006)***
Age 25-44	0.210	0.211	0.212
	(0.005)***	(0.005)***	(0.005)***
Age 45-65	0.279	0.282	0.285
	(0.009)***	(0.008)***	(0.007)***
Ethnic group average characteristics			
Ethnic group share with primary or no education		0.182	0.072
		(0.042)***	(0.021)***
Ethnic group share with post-secondary education		0.197	0.136
		(0.035)***	(0.024)***
Ethnic group share of male		0.070	0.069
		(0.031)**	(0.018)***
Ethnic group share 45-64		-0.067	-0.020
		(0.033)**	(0.019)
R^2	0.11	0.11	0.13
N	123,885	121,797	121,793
Significance of ethnic group characteristics (F-stat)		10.03	11.84
P-value of joint significance test		0.00	0.00
Ethnic group fixed effects	No	No	Y
Local municipality fixed effects	No	No	Y

Notes: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Data: South Africa 2011 census. Sample: all Black South African internal migrants aged 15 to 64 in 2011. Averages for established migrants are computed for those aged 25 or above. The table presents the marginal effects obtained from a linear probability model explaining the probability of being employed by individual characteristics and the characteristics of established migrants of the same ethnic group. Specification 1 contains only individual characteristics, whereas 2 and 3 introduce ethnic group characteristics and fixed effects for the ethnic group and municipality. Standard errors are clustered at the ethnic group x municipality level. The F-tests presented are for the joint significance of the characteristics of established migrants.

have chosen to remove the share of the ethnic group that has secondary education. This choice was motivated by the fact potential benefits of ethnic capital from individuals with low and high levels education, which would translate into a nonlinear relationship between the average level of education of the ethnic group and the chances of employment.

the one hand, an internal migrant's chances of getting employed increase with the share of members of his ethnic group that have no level of education or just primary education. The effect persists even after controlling for local labor market conditions, even though the estimates are roughly divided by two. From the most conservative estimates, a 10 percentage point increase of the share of primary or no education holders increases the chances of being employed by .7 percentage points. On the other hand, the share of members that have a post-secondary degree increases the chances of getting employed, an effect that is twice as large as the effect of the share of individuals with primary education or less. Finally, a higher proportion of males among previously established migrants translates as well into higher chances of employment for new migrants. Although the age composition of the population seems to matter, suggesting that the older the ethnic group is, the less migrants benefit from it in terms of employment opportunities, it does not remain significant once local labor market conditions are accounted for.

One potential issue is the fact that I perform my analysis using individual-level data, while I am interested in the effects of variables measured at the ethnic group - municipality. An analysis at the ethnic group - municipality level would address some of the concerns related to unobservables at that level. Table B.4 presents the results of the estimates at that aggregate level. Averages of all the variables are computed to obtain employment rates as the dependent variable, and averages of independent variables. Regressions are weighted by the size of each cell (ethnic group - municipality level). The results support the idea that ethnic capital matters when considering the aggregate level for the regressions, even after accounting for unobserved ethnic group and municipality effects (see columns 2 and 4). However, the results exhibit some abnormalities that potentially qualify this analysis.

Regarding individual characteristics, a large share of migrants with no education or some primary education seems to lead to a higher employment rate, when the base is secondary education. This is counter-intuitive, and also at the odds with the regressions at the individual level. Given that the conclusion that ethnic group matters holds both with and without these individual characteristics, the issue with the individual variables is unlikely to affect this conclusion. Still, this may warrant further investigations.

Ethnic capital or language proficiency?

My analysis assumes that the language a migrant primarily speak helps identify his or her ethnic group. This assumption excludes the possibility that a migrant learns the language that he or she speaks mostly, a situation that may happen when migrants choose to learn, speak primarily or as a second language a language specific to the place where they move. Previous literature (Deumert et al., 2005; Deumert and Mabandla, 2009) has exhibited a change in languages for internal migrants in South Africa, mostly stressing that the dialect of internal migrants experiences some changes to incorporate elements of the local language. To see the extent to which my results are sensitive to these concerns, I perform several analyzes based on language proficiencies, summarized in Table B.5. I first check whether the results are different with respect to the proficiency of the most prevalent language in the municipality of residence. Column 5 in Table B.5 presents the results for migrants who do not speak the most prevalent language, whereas column 6 shows the results for migrants who speak that language. The results are mostly similar when it comes to the distribution of education; if anything, the effect of the share of individuals with a post-secondary degree is slightly larger for migrants proficient in the most prevalent language in their place of residence. Other variables capturing ethnic capital are significant only for migrants who speak the most spoken language. These results suggest that migrants who speak the most prevalent language in the place where they move to benefit more from their ethnic capital. The fact that those who do not speak the most prevalent language also benefit from their ethnic capital implies that the results of the importance of ethnic capital are not entirely driven by those who master the language spoken where they move. In the end, concerns that individuals may learn contribute to the results presented in this paper, but these concerns are not the sole drivers of my results. I explore two additional concerns related to language proficiency. I look at internal migrants who speak English either as their first or second language. Results presented in column 2 show little difference compared to the results for all migrants. Similarly, I look at migrants who speak two languages. The results (Table B.5 column 3) are mostly similar. If anything, they suggest a weaker effect of the share of established migrants with a post-secondary degree and a stronger effect of the share of primary education holders. I take these results overall as suggestive evidence that the results indeed reflect ethnic capital and not merely language proficiency.

Robustness: alternative sample analyses

In the analysis so far, I have defined internal migrants as individuals who have changed the local municipality where they reside in 2001 or afterward. This definition incorporates individuals that have just moved inside the same province, and individuals that have moved to a neighboring local municipality. Hypothetically, an individual could have changed the place where he lives while still keeping the same job. This is more likely to happen when an individual moves to a neighboring local municipality, or moves within the same province. It will be erroneous to attribute anything to the ethnic group for the success in the labor market of those migrants. To make sure that those individuals do not drive my results, I present in Table B.6 the analysis where I excluded them in the sample of internal migrants. The analysis here thus includes only individuals that have moved to a non-neighboring local municipality and in a different province. The results show that the distribution of education and the gender composition of a migrant ethnic group matters for his own chances of getting employed, even after controlling for local labor market conditions and unobservable at the migrant ethnic group. This suggests that the results observed above are unlikely to be driven by people moving while still keeping their jobs. An additional issue is that I don't distinguish the reasons why individuals migrate. However, my main focus is on individuals who are migrating for economic reasons, i.e. in search of a job. Including people who migrate for other reasons may lower the actual role of the ethnic group in finding a job. To get a sense of how not identifying the economic migrants might affect the results, I conduct two separate analyses. In the first analysis, I considered only migrants that are between 25 and 64 years old, whereas my main sample contains individuals aged 15 to 64. Arguably, younger individuals may have migrated for other reasons, such as just accompanying their families, and not necessarily in search of a job themselves. The results in Table B.7 show that my conclusions for all the migrants hold in particular when I consider only migrants that are 25 years older or more.

In the second analysis, I focus on the main province of South Africa, Gauteng. Besides being the most populous province (24% of the population) and the main destination of internal migrants (42% of them), Gauteng is the richest province of the country, as it concentrates 34% of the GDP of the country. [14] According to the Labor Force Survey, the province receives 71.8% of migrant workers. This suggests that migrants choosing to settle in Gauteng are more likely to be economic migrants. Therefore, analyzing only migrants in Gauteng should provide results that are more related to economic migrants. Table B.8 presents the results of equation 2.3 estimated only for migrants who are living in Gauteng. The results show that the ethnic groups matter for migrants in Gauteng, with some specificities. The effects of the lower end of the distribution of education are similar, whereas the effect of the upper end of the distribution, the share of individuals holding a post-secondary degree, is stronger (50% higher).[15] In all, the previous two analyzes provide evidence that the results of the role of ethnic capital persists when analysis is restricted to individuals more likely to migrate for economic reasons.

Robustness check: IV analyses

So far, I have used as independent variables in my regressions the characteristics of established internal migrants from the same ethnic group. However, as internal migrants decide where they move to, this variable is prone to endogeneity. As mentioned in the methodology section, including local municipality fixed effects contributes to addressing

[14] See the provincial review 2016. The Real Economy Bulletin:
`https://www.tips.org.za/images/The_REB_Provincial_Review_2016_Gauteng.pdf`

[15] These results could be linked to the fact that migrants to locate in Gauteng are in general better educated. The share of migrants that have a post-secondary degree in Gauteng is 20% whereas this share is 15% in the other provinces. As the specification controls for the migrant own level of education, the higher effect in Gauteng can be in part attributed to the interaction effect found between one's own level of education and the distribution of education.

that endogeneity issue. However, it does not entirely solve the problem, especially to the extent that the choice of residence is related to the ethnic group's presence in the destination, in ways that the additive fixed effects do not capture. I follow previous research (e.g., Borjas, 1992; Patel and Vella, 2013) and use as instrument variables the values of these characteristics in the past, in my case, 10 years prior (using South Africa's 2001 census data). As an illustration, I instrument for the share of established migrants who hold a post-secondary diploma in 2011 by using the share of established migrants that held a post-secondary diploma in 2001 in the local municipality of residence. The identification strategy here assumes that migration decision responds to recent migration experience in the ethnic group; and that historical migration patterns have an effect on the current migration patterns, only through the recently established migrants.

Table B.10 shows the results of the instrumental variable estimations. I estimate a linear probability model equivalent to equation 2.3.The first stage regressions (Table B.10) show that historical values of the characteristics of established migrants are highly correlated with their current values, suggesting that the instruments are not weak. The IV estimations also confirm that ethnic capital is associated with better chances of getting employed for internal migrants, suggesting that the results observed above are robust to the endogeneity of migrants choice of their destination. The results also show that if anything, OLS are biasing the effect of ethnic capital downwards.

Overall, the results suggest that the characteristics of his ethnic group can explain an internal migrant's probability of employment. These results thus show that the importance of social networks that has been highlighted by Burns et al. (2010) holds for internal

Table 2.7: Employment and ethnic group characteristics: IV estimations

	OLS (1)	IV (2)	OLS (3)	IV (4)
Migrant own characteristics (dummy variables)				
No education or some primary	-0.025	-0.031	-0.031	-0.038
	(0.004)***	(0.004)***	(0.004)***	(0.004)***
Post-secondary	0.206	0.206	0.201	0.201
	(0.003)***	(0.004)***	(0.003)***	(0.004)***
Person is male	0.168	0.173	0.167	0.170
	(0.003)***	(0.003)***	(0.003)***	(0.003)***
Age 25-44	0.211	0.206	0.212	0.208
	(0.003)***	(0.003)***	(0.003)***	(0.003)***
Age 45-65	0.282	0.272	0.285	0.279
	(0.005)***	(0.006)***	(0.005)***	(0.005)***
Ethnic group share (fraction)				
Ethnic group share with primary or no education	0.182	0.284	0.072	0.300
	(0.012)***	(0.025)***	(0.017)***	(0.141)**
Ethnic group share with post-secondary education	0.197	0.405	0.136	0.589
	(0.018)***	(0.068)***	(0.021)***	(0.213)***
Ethnic group share of male	0.070	0.060	0.069	0.082
	(0.011)***	(0.027)**	(0.015)***	(0.109)
Ethnic group share 45-64	-0.067	-0.135	-0.020	-0.206
	(0.014)***	(0.085)	(0.016)	(0.244)
R^2	0.11	0.11	0.13	0.13
N	121,797	99,410	121,797	99,410
Significance of ethnic group characteristics: statistic.	77.00	132.36	18.69	22.90
P-value for joint significance test	0.00	0.00	0.00	0.00
Ethnic group fixed effects	No	No	Yes	Yes
Local municipality fixed effects	No	No	Yes	Yes

Notes: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Data: South Africa 2011 census. Sample: all Black South African internal migrants aged 15 to 64 in 2011. Averages for established migrants are computed for those aged 25 or above. Instruments for the IV estimates are the values for established migrants, from the 2001 census. The table presents the marginal effects obtained from a linear probability model explaining the probability of being employed by individual characteristics and the characteristics of established migrants of the same ethnic group. Results from OLS and IV estimations are compared in the table, with specifications that include or not fixed effects for the ethnic group and the municipality. Tests for the joint significance of the characteristics of established migrants are presented. F-tests are used for OLS estimations and χ^2 tests for the IV estimations.

migrants in particular as well. [16]

My results also suggest that migrants' social capital matters not only for the migration decision, but also for the employment opportunities of the migrants. The importance of social networks in the migration decision in South Africa has been suggested by Kok et al. (2003) and Gelderblom and Adams (2006); it has also been highlighted by Reed (2009) and Stapleton (2015) who have quantified this importance. Reed (2009) suggested that social networks play a role only for the choice of destination, but not the probability of finding a job, but the present analysis shows a link between the characteristics of the ethnic group and the chances of getting a job.

2.5.2 Occupational choice and ethnic capital

I now present the results on the role of the ethnic capital in the occupation choice. [17] Table 2.8 presents the results of the estimation of equation 2.5, which quantifies the relationship between the occupation and the ethnic group characteristics. The coefficients are the marginal effects of the variables on the probability that a migrant chooses a particular

[16] The results presented in my analysis are not directly comparable to those in Burns et al. (2010), as their methodology differs from mine. They use a network measure which is the national employment rate of the ethnic group multiplied by the size (number of members) of the ethnic group in the municipality of residence. This network measure helps include fixed effects to account for correlated effects. In the end, the parameter that they obtain (coefficient of the network measure) does not have a direct interpretation, but helps compute the share of employment that is due to social networks. In my analysis, however, I present evidence of how ethnic capital affect these employments, with reduced form estimates of the relationship between individuals' employment probability and their ethnic capital. That said, results in Burns et al. (2010)seem to suggest, however, that migrants exhibit smaller benefits from their networks than non-migrants.

[17] I choose to include unemployed and discouraged workers as a category for the occupation variable. As mentioned in section 2.4, this is motivated by the importance of these categories in the labor market. An alternative would have been to consider two models. The first model would consider the choice between being discouraged workers, unemployed, and employed. The second model would consider the choice between the different occupations (in the usual sense) in the labor market. Modeling jointly these choices help consider the idea that some job seekers or jobless individuals are passed information about job openings, or are made job offers that they turn down (see, e.g., Levinsohn et al., 2014).

occupation. The estimation controls for unobservable effects at the ethnic group level as well as the local municipality level. This implies that the results are identified holding local labor market conditions constant as well as unobservable factors at the ethnic group level that might explain differential choices of occupation, for reasons that are not related to the presence of the ethnic group in the particular local municipality. From the results, the ethnic group matters for the choice of occupation, although not for all occupations. The χ^2 test for the joint significance of the characteristics of the ethnic group reveals that they play no significant role in the probability of being employed in skilled agriculture or in sales and services. [18]

Looking specifically at what characteristics matter at the ethnic group level, regarding the distribution of education, the share of individuals with a post-secondary degree is the variable most frequently significant. It increases the chances that an internal migrant works as a manager, a professional, a clerk, or a plant operator. Having a high share of established ethnic contacts that have low or no education helps for some occupations as well, namely plant operators, elementary occupations, and domestic workers. The gender and age distribution also matters for the choice of occupations. A higher share of men in the networks is associated with higher chances of being employed in craft and in plants, while there are lower chances of working as a domestic worker. In the case of the latter occupation, it is worth noting that 76% of domestic workers are women, which may explain the effect of the gender composition on the choice of this occupation. Similarly,

[18]The shares for the different categories of the occupation variables are quite different, as low as .5% for skilled agriculture, which implies smaller variation in these variables. This potentially has implications in terms of power of the statistical analysis. That said, besides skilled agriculture, the other category for which ethnic capital does not play a significant role is services an sales, an occupation that 12% of migrants choose. It is less likely that power issues apply to the results for services, but they could well explain the lack of results for skilled agriculture.

Table 2.8: Occupational choice and ethnic capital (marginal effects)

Occupation (dependent variable)	*Ethnic group share (independent variables)*				$\chi^2 stat.$
	Primary education	Post-secondary education	Aged 45-65	Male	(p-value)
Manager	0.003	0.025	0.009	-0.010	11.48**
	(0.008)	(0.010)**	(0.007)	(0.007)	(.022)
Professional	-0.010	0.025	0.010	0.003	12.11**
	(0.008)	(0.010)**	(0.008)	(0.007)	(.016)
Technician	-0.003	0.016	-0.026	0.006	12.89**
	(0.009)	(0.011)	(0.009)**	(0.008)	(.012)
Clerk	0.009	0.045	0.001	0.002	12.64**
	(0.010)	(0.013)**	(0.010)	(0.009)	(.013)
Sales and service	-0.003	-0.011	0.014	-0.007	2.24
	(0.012)	(0.015)	(0.012)	(0.011)	(.691)
Skilled agriculture	0.001	-0.004	0.001	0.002	2.45
	(0.003)	(0.004)	(0.003)	(0.002)	(.653)
Craft and related traders	-0.013	0.023	0.001	0.043	26.01***
	(0.011)	(0.014)	(0.011)	(0.010)**	(.000)
Plant and machine operators	0.017	0.027	-0.013	0.043	44.70***
	(0.008)*	(0.011)*	(0.008)	(0.007)**	(.000)
Elementary occupations	0.044	-0.017	-0.011	0.013	19.39***
	(0.012)**	(0.015)	(0.012)	(0.011)	(.000)
Domestic worker	0.027	0.008	-0.005	-0.026	17.28***
	(0.009)**	(0.012)	(0.009)	(0.008)**	(.001)
Unemployed	-0.053	-0.094	0.006	-0.045	38.71***
	(0.016)**	(0.020)**	(0.015)	(0.014)**	(.000)
N	121,780				

Notes: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Data: South Africa 2011 census. Sample: all Black South African internal migrants aged 15 to 64 in 2011. Averages for established migrants are computed for those aged 25 or above. The table presents the results of the estimation of a SUR model where each equation represents an occupation. Each row represents an occupation choice. Discouraged workers are the occupation group that has been chosen as the reference. Each equation controls for ethnic group fixed effect as well as local municipality fixed effects. The last column presents for each regression the χ^2 test for the joint significance of the characteristics of the ethnic group, with the associated p-values in parentheses.

78% of individuals craft workers and 73% of plant operators are male, which can explain the effect of the gender composition on those choices. [19]

One interesting analysis is to relate the importance of the ethnic capital to the education

[19] The story here is choice of occupation for which the gender gap is large are associated with the gender composition of established migrants. This suggests as well that there might be different components of the ethnic capital for men and women.

level required for the different occupations. Ioannides and Datcher Loury (2004) note that the role of social networks in finding a job may differ with the skills required in the occupations, depending on how difficult it is to infer the productivity of the incoming workers based on those skills. Regarding the probability of finding a job, it is often shown that less skilled workers benefit more from their networks than skilled workers. Table B.9 shows the distribution of education for each occupation, allowing a ranking by the level of skills required. “Professionals” have the highest share of workers with a post-secondary degree, followed by the technicians, whereas the domestic workers and the elementary occupations workers are the least likely to have such a degree. Comparing the results of the estimations suggest that social networks matter more for occupations that require less education. Overall, putting together the estimations in Table 2.8 and the distribution in Table B.9, the results suggest that the ethnic group characteristics especially matter for occupations that require less education.

These results are, however, small in their magnitude, and the elasticities they imply are very low. An illustration is the effect of the share of individuals having a post-secondary diploma on the probability of being employed as a clerk. The estimate implies that a one percentage point increase in the share of the ethnic group with a post-secondary diploma leads to an increase by .04 percent the probability of being employed as a clerk, which is the highest effect induced by the education distribution. While an increase of the share of individuals holding a post-secondary diploma by 1% (the national share is 10%) is economically significant, the implied increase in the probability of being employed as clerk is not, as this represents the occupation that 7% of migrants choose. This result seems to be in line with the study of occupation enclaves and social networks by Hofmeyr

(2010), which suggests that internal migrants are less likely to be in occupations where their social group is over-represented. One caveat to this study and Hofmeyr (2010) is that actual social contacts are not observed, and to the extent that our definition includes individuals that are not relevant for migrants' networks, the present estimates might be biased downwards when it comes to the role that social networks play in the choice of occupation.

2.6 Conclusion

Social networks play an important role in different areas of economic decisions, especially in labor markets and migration. This importance has been highlighted both theoretically and empirically and is thought to be more prevalent in developing countries (Chauang and Shecther, 2015). In this paper, I focus on one dimension of social networks, the ethnic group, and study how the group characteristics affect labor market outcomes for internal migrants in South Africa. My results show that a migrant's employment probability is related to the education distribution and the gender composition of established individuals from his ethnic group. The results hold even after I control for unobservables at the place of residence or the social group considered. I also provide evidence that the choice of occupation by a migrant is associated with the characteristics of the social group in the place of residence. Results on the occupational choice show that the ethnic group's characteristics in terms of education and gender induce migrants to choose specific occupations. The results in magnitude are however generally small, although IV estimations for employment probability suggest that they might be larger.

This article suffers, however, from some limitations that may qualify the results. The main limitation comes from the choice of data. The main advantage of using the census data is the size of the sample obtained, especially for the representation of very small groups, like the groups obtained by crossing the place of residence and the language/province of birth groups. This gain in representation comes with limitations on the information collected. One such information not collected is the actual links between individuals. The results here are to be interpreted as evidence of the importance of networks in labor markets, while not quantifying the social network effect as defined in the peer effects literature. To the extent that interaction happens primarily within one's own language group, the results presented here represent the effects on migrants of the choices and behavior of potential contacts that they have in the place where they now reside.

Besides this aspect, the results presented here are reduced form evidence of the relationship between the ethnic group characteristics labor market outcomes for migrants. The results presented here hold when I control for unobserved common characteristics to the ethnic group, but at a national level, although the additive local municipality fixed effect allows for the control of some local ethnic group common effect. This suggests that further analyses will help better understand the role that social networks play in the insertion of internal migrants in the labor market in South Africa in particular, and in developing countries in general.

Chapter 3: Cheap Talk and Coordination in the Lab and in the Field: Collective Commercialization in Senegal

This chapter is co-authored with Tanguy Bernard and Angelino Viceisza.

3.1 Introduction

Many activities depend on coordination among a sometimes large number of agents; for example group projects, the meaning of language, and navigating traffic. From a game theory standpoint, coordination is required whenever there exist multiple, potentially Pareto ranked, equilibria and theory does not clearly predict behavior. In fact, without prior knowledge about others' intended actions, strategic uncertainty will likely steer agents away from activities that require coordination. This in turn leads to coordination failure; that is, agents choosing risk-dominant strategies over Pareto-dominant ones, thus leading to suboptimal outcomes (e.g., Rosenstein-Rodan, 1943).

Coordination games have featured prominently in the literature.[1] Moreover, with coordination failure at the heart of certain poverty traps (Wydick, 2007), a key policy question has been how to enhance coordination (Hoff, 2001). For example, Devetag and Ortmann

[1] Section 3.3 will discuss our work in the context of prior literature. Some overviews are Ochs (1995), Goeree and Holt (2002), Camerer (2003), Devetag and Ortmann (2007), Van Huyck and Battalio (2008), and Cooper and Weber (2019).

(2007) point at several factors that may improve coordination: (1) smaller group sizes (e.g., Van Huyck et al., 1990); (2) lower attractiveness of the secure relative to the risky action (e.g., Brandts and Cooper, 2006); and (3) certain types of cheap talk/communication in certain types of games (e.g., Cooper et al., 1992; Farrell, 1987). Yet, as Avoyan and Ramos (2017) argue, while communication has been lauded for its potential to enhance coordination, there are also instances where it has had no effect (e.g., Cooper et al., 1992). One aspect that has received limited attention is the interaction between the number of players (i.e., group size) and communication (one exception is Feltovich and Grossman, 2015). In addition, most prior literature on coordination and its drivers (in particular communication) is theoretical and/or based on conventional laboratory experiments. Cooper and Weber (2019) also conclude that coordination games like ours (i.e., threshold games) and those that relate to problems in organizations (of which farmer cooperatives are an example) require further study. So, to better understand drivers of coordination, it is useful to complement existing findings with field evidence.

In this paper, we explore the impact of cheap talk on coordination and whether this effect varies with group size, in a field context where strategic uncertainty has historically led to coordination failure.[2] We study farmer cooperatives that seek to sell members' agricultural production collectively – a setting that represents much of the developing and developed world. Farmer groups have received considerable attention in policy and research. For example, the UN designated 2012 as the year of cooperatives, since 85% of farms worldwide are small family farms and most of them participate in cooperatives. Barrett

[2]Throughout, we use the following terms interchangeably: (1) cooperatives, farmer groups, and producer organizations; (2) farmers, smallholders, members, players, and subjects; (3) cheap talk, communication, intentions, plans/planned actions; and (4) coordination and collective commercialization/selling.

(2008) argues that mechanisms aimed at facilitating smallholder organization (e.g., cooperatives) are central to stimulating market participation. At the same time, cooperatives are not always able to achieve the goal of selling collectively (e.g., Bernard et al., 2008; Fafchamps and Hill, 2005; Ragasa and Golan, 2014). Our primary goal is not to promote collective commercialization, as we are agnostic regarding its effective capacity to increase farmers' income.[3] Instead, we seek to understand why, in groups designed for the purpose of collective commercialization, farmers are unwilling to do so. In particular, because those same farmers believe that collective commercialization could lead to higher incomes.

Based on data from the context in question (Bernard et al., 2014) as well as prior literature (e.g., Devetag and Ortmann, 2007), we hypothesize that (1) beliefs about others' sales to the group are key to one's decision to sell through the group; (2) larger groups have more difficulty coordinating; and (3) communicating others' intentions to sell through the group (i.e., cheap talk) can increase the likelihood of collective commercialization, particularly in larger groups.[4] In order to test these hypotheses, we work with a sample of 79 Senegalese groundnut-producing cooperatives and close to 2800 individual group members.

[3]Collective commercialization is often promoted in Sub-Saharan Africa on the grounds of imperfect competition between traders in remote and thin markets, which in turn leads to rent extraction from farmers. By aggregating farmers' production, collective commercialization offers the possibility to (1) reach distant but larger markets (e.g., due to lower transaction costs, Key et al., 2000) or (2) negotiate better with local traders. As reviewed by Dillon and Dambro (2017), the empirical evidence is mixed. For example, Bergquist (2017) finds a low pass-through of price premiums from traders to farmers in Kenya, while Casaburi and Reed (2017) find a high pass-through in Sierra Leone. Bernard et al. (2008) find positive but limited impacts of cooperatives in Ethiopia and Ashraf et al. (2009) find positive, but unsustained, impacts in Kenya.

[4]Bernard et al. (2014) find that among a sample of 27 groundnut-producing cooperatives in Senegal, 67% of group members believe that, if presented with the opportunity, other members would by-pass sales through the group and sell individually to a trader for a potentially lower, but more certain payoff. There may of course be other reasons why cooperatives are (un)able to sell collectively; e.g., Casaburi and Macchiavello (2019) and Hill et al. (2014) discuss commitment sanctions and deferred payments respectively.

First, in a subsample of these cooperatives, randomly selected group members participate in neutrally framed, high-stakes coordination games (i.e., artefactual field experiments in the terminology of Harrison and List, 2004, hereafter, lab-in-the-field experiments/LFEs) where we exogenously vary (1) the experimental group size, (2) the threshold required for coordination, and (3) information on others' intentions to coordinate (i.e., cheap talk). Second, six months later, we implement a comparable naturally-occurring intervention (i.e., a natural field experiment in the terminology of Harrison and List, 2004, hereafter, randomized controlled trial/RCT) with the full sample of cooperatives. Shortly after the harvest period, we collect all group members' intentions to sell through the cooperative in the upcoming commercialization season. In a random subset of these cooperatives, members' intentions are revealed in a public meeting right before commercialization starts.[5] Since these are existing cooperatives, unlike in the LFEs, it is infeasible to exogenously manipulate the group size and coordination threshold. Thus, both the LFEs and the RCT are key for testing our hypotheses. Finally, we supplement these experiments with survey data and cooperative records (i.e., administrative data).

Our findings are as follows. First, revealing aggregate intentions (i.e., cheap talk) improves coordination in both the LFEs and the RCT, particularly in larger groups. These results are relatively robust to (1) farmers' preferences, (2) a series of placebo tests confirming that intentions are balanced at baseline, (3) potential social desirability bias, which we rule out by comparing survey responses to group-administrative data, and (4) correcting for multiple hypothesis testing using the Holm-Bonferroni approach. Second, there

[5] This form of communication is different from the casual, mostly one-on-one "chit chat" that group members typically engage in because it enables them to know aggregate group intentions, which is a good indication of potential success in collective commercialization.

is suggestive evidence that the intervention positively impacts revenues (welfare) among farmers who produce smaller quantities. Third, behavior in the "lab" (LFEs) transfers to behavior in the "field" (RCT). We exploit the random selection of individuals and find that, in the same cooperative, LFE participants were more likely to engage in day-to-day collective commercialization than non-LFE participants. Consistent with the literature on transfer (e.g., Cooper and Van Huyck, 2016), but perhaps contrary to Voors et al. (2012), participants learn from the LFEs.[6] Our study thus also contributes to the literature linking results across the spectrum of experiments; see for example, Levitt and List (2007), Camerer (2015), Viceisza (2016) and the references within.[7] While our findings support the use of LFEs in combination with RCTs, researchers should take heed that participation in LFEs may affect subjects' future behavior.

The remainder of the paper proceeds as follows. Section 3.2 provides evidence of strategic uncertainty and coordination failure in Senegalese cooperatives. Section 3.3 discusses our framework and hypotheses in the context of existing literature. Section 3.4 details our study design and empirical strategy. Section 3.5 covers our main findings. Finally, Section 3.6 concludes.

3.2 Strategic uncertainty in Senegal

Groundnut production has long been the backbone of the Senegalese economy. According to Caswell (1984), groundnut processing contributed to 42% of all industrial output

[6] Similar findings have been documented by Cardenas and Carpenter (2005) and Turiansky (2017). This is also comparable to a series of papers on classroom experiments in *Journal of Economic Perspectives* in the late 1990s.

[7] E.g., Barr et al. (2010), Finan and Schechter (2012), Stoop et al. (2012), and Hoel et al. (2017).

and groundnuts represented 80% of all export revenues. For a long time, a dense network of state-controlled cooperatives was the exclusive interface between farmers and other actors. Each year, the price of groundnut was fixed at the national level and all marketed production was collected through cooperatives. Over time, however, revenues from the groundnut sector steadily declined and in the late 1990s, external and internal shocks led to privatization of all segments of the value chain. In 2010, the export-monopoly previously granted to Suneor (the principal end-buyer of groundnuts) was abolished, facilitating the entry of new potential buyers.

These reforms led to important changes in farmers' ability to coordinate within cooperatives. In this new system, group members became strategic complements since the cooperative could obtain better per-unit prices by aggregating production and accessing larger, more distant markets/processing plants, or by negotiating more strongly with local traders. However, farmers could also sell directly to local traders. So, this presents a tradeoff: While aggregation through the cooperative could lead to a higher unit price, receipt of such funds would usually be delayed relative to local traders who pay cash on delivery. Moreover, the price premium obtained from collective commercialization would also depend on the quantity that other members sell through the cooperative, which is uncertain when a given farmer is visited by the trader.

Bernard et al. (2014) document commercialization among close to 300 members of 27 groundnut cooperatives in Senegal whose main, stated objective is collective commercialization. A great majority of members sells individually to local traders in spot market transactions. In fact, only 11 of the 27 groups had sold collectively in 2011, and for those that did, only half of the members participated. This limited involvement contrasts with farmers' perceived potential for collective commercialization: While 79% of farmers cited problems when attempting to sell individually (e.g., lack of (1) transportation to reach more lucrative markets, (2) knowledge of current prices, and (3) production to negotiate better prices), practically all respondents were convinced that group commercialization could alleviate such constraints.

Bernard et al. (2014) also collect members' aversion to the strategic uncertainty arising from collective commercialization using a framed version of the instrument by Heinemann et al. (2009). Subjects were presented with the choice to sell to a trader at a certain price or sell through the group at a higher but uncertain price that depends on the number of other members selling through the group. The price offered by the trader was incrementally raised until subjects switched from selling through the cooperative to selling to the trader. The results suggest that aversion to strategic uncertainty is negatively correlated with individuals' effective sales through the cooperative. When asked about other group members' likely response to the above strategic uncertainty question, 67% of the sample believed that those members would by-pass sales through the group and sell individually to a trader for a lower payoff than they would. In short, reforms in the groundnut sector paved the way for issues of strategic uncertainty in Senegalese cooperatives.

3.3 Framework, literature, and hypotheses

3.3.1 Framework

Our framework is a critical-mass/threshold coordination game. There are $N \in \mathbb{N}$ players, each of whom j has a positive endowment $V_j \in \mathbb{N}$. All players simultaneously choose an amount $A_j \in [0, 1, ..., V_j]$ to send to the group (the equivalent of commercializing collectively) and keep the remainder $V_j - A_j$ (the equivalent of selling individually/via traders). Any quantity A_j sent to the group, earns a high return H if all players jointly send a quantity $A = \Sigma_j A_j \geq T$ to the group, where T represents some threshold. If they send a quantity $A < T$, A_j earns a low return L. Whatever a player chooses to keep individually, that is $V_j - A_j$, earns a medium return M, where $L < M < H$.[8]

Accordingly, player i's expected payoff can be expressed as $\Pi(A_i) = pA_iH + (1 - p)A_iL + (V_i - A_i)M$, where $p = P\,(A \geq T)$.[9] It is clear that equilibria in this game will be driven by player i's belief, p_i, about the likelihood of the threshold being surpassed, p. I.e., contributions to the group (how much the farmer chooses to commercialize collectively) will depend on the player's sense of strategic uncertainty. If s/he expects A to surpass T, $p_i = 1$, the player should send all of the endowment to the group, i.e., $A_i = V_i$. If s/he expects A not to surpass T, $p_i = 0$, the player should keep the full endowment, i.e., $A_i = 0$. If s/he expects a scenario in between these two extremes, the player should diversify by selling $A_i \in \{1, ..., V_i - 1\}$ through the cooperative.

As discussed in Section 3.2, farmer cooperatives seem to have been unable to commercialize collectively in this setting in recent years. Moreover, most members believe that other members are more likely to sell individually to traders than collectively through the cooperative. In short, empirical evidence suggests that $p = 0$ or close to it. So, a key question is whether and if so, how p can be increased, given farmers believe it would be beneficial to do so.

[8] As discussed in Sections 3.1-3.2, there are several reasons for $L < M < H$. This condition is also consistent with farmers' beliefs about the benefits to collective commercialization.

[9] Here, we abstract from the more general "utility" function and assume that H, M, L capture players' true payoffs. In the empirical analysis, we check/control for preferences.

3.3.2 Related literature and hypotheses

There is a relatively extensive theory and lab literature on coordination (in games) and the potential determinants thereof (see for example Crawford, 1998; Farrell, 1987; Goeree and Holt, 2002; Ochs, 1995; Van Huyck and Battalio, 2008, for some overviews). The literature that relates most closely to our setting is on (1) coordination games with Pareto-ranked equilibria, which cover both order-statistic games (e.g., Van Huyck and Battalio, 2008; Van Huyck et al., 1990) and stag-hunt games (e.g. Battalio et al., 2001; Charness and Grosskopf, 2004; Cooper et al., 1992) and (2) preplay communication/cheap talk (e.g., Avoyan and Ramos, 2017; Blume and Ortmann, 2007, and some of the previous references).[10]

Devetag and Ortmann (2007) and Cooper and Weber (2019) survey some of this literature and both point at a range of common factors that may enhance coordination. Among those are (1) smaller group sizes; (2) certain types of cheap talk/communication in certain types of games; and (3) lower attractiveness of the secure action relative to the risky one. As Feltovich and Grossman (2015) note, there has been little work jointly investigating the effect of group size and cheap talk. In fact, their study is one of the few exceptions.[11] It is also closest to ours, but there are a few key differences. First, their setting is a threshold public-good game with a dichotomous choice ("contribute" or "not"). Second, their communication takes the form of one subject/player sending a message that "Everyone should

[10] Kets and Sandroni (2017) address strategic uncertainty and cultural diversity.

[11] Balliet (2010) conducts a meta-analysis of 45 social-dilemma experiments and finds a positive interaction between group size and the effect of communication. However, this interaction is not robust to controlling for other relevant variables and removal of two outliers (see Feltovich and Grossman, 2015, for a more detailed discussion/analysis).

choose [action]". Third, their group sizes are smaller. Finally, we examine this interaction both in the "lab" and "field". Particularly in the latter case, there may be pre-existing experiences that impact the interaction between group size and cheap talk.

In order to state the main hypotheses that form the basis of the experimental design in Section 3.4, we further extend the framework discussed in Section 3.3.1. Suppose there is a pre-stage to the game in which all players simultaneously reveal an *intended* action towards the group A'_j (i.e., N-way communication) prior to taking the "true" action A_j. Player i will use intentions A'_j to inform her/his belief p_i about p only if s/he believes those intentions. In particular, player i will substitute A'_{-i} for A_{-i} in order to more precisely assess the likelihood of surpassing the threshold T. Accordingly, we state the following hypotheses.[12]

Hypothesis 1. As the group size N increases, player i decreases her/his contribution to the group A_i, since coordination failure is more likely. In other words, due to strategic uncertainty, smaller groups are more conducive to coordination.

Hypothesis 2. Player i will set $p_i = p = 1$ and thus $A_i = V_i$ if revealed intentions are such that $\Sigma_j A'_j = A' \geq T$. In other words, cheap talk can reduce coordination failure if aggregate intentions surpass the threshold.

Hypothesis 3. The effect of cheap talk (i.e., $\Sigma_j A'_j$ relative to T) will increase with the group size N. In other words, due to strategic uncertainty, cheap talk is more effective in

[12] While some of the literature on communication has suggested that players may seek to deceive others, particularly when sending non-costly and non-binding messages, there is no such incentive in this game. Unlike Prisoner's dilemma/public-good type games, a player cannot "free-ride" by sending a high intention A'_j and then taking a low action A_j. If a player truly believes that a high signal will cause others to increase their contributions to the group, this player should align her/his actual contribution with such intentions. This is the same argument that has been made by for example Crawford (1998) – recall the "reassurance" role of cheap talk in stag-hunt games.

larger groups. However, the direction of the effect will depend on whether or not aggregate intentions surpass the threshold. This hypothesis basically combines hypotheses 1 and 2.

3.4 Study design

As mentioned in Section 3.1, we are unable to exogenously manipulate two key parameters – the group size N and the threshold T – in the naturally-occurring environment. As such, group-size effects in the day-to-day setting could be capturing other (potentially unobserved) characteristics. For example, larger groups may (1) offer better services to farmers to begin with (e.g., input, credit, or negotiated prices); (2) be more likely to achieve a given threshold level, ceteris paribus; and (3) be associated with greater strategic uncertainty due to a greater number of "moving parts" (i.e., number of players). Accordingly, our design is based on a two-stage empirical strategy. First, we fully test our predictions through controlled, neutrally framed LFEs in which we are able to observe and manipulate N and T as well as other aspects such as external uncertainty and the premium for coordination (H). Then, we test our results in a naturally-occurring RCT where we do not manipulate the above parameters. The timing was as follows (Table C.1 summarizes the relationship between the LFE and RCT samples):

1. *LFEs:* From May to June 2013, randomly selected farmers from 28 cooperatives participated in variants of neutrally-framed, high-stakes, threshold coordination games (Section 3.4.1). A presurvey collecting basic (behavioral) characteristics was also administered at this time.

2. *RCT:* From November 2013 to February 2014, 79 farmer groups (including 26 of the 28 that were in the LFEs) participated in variants of intention-revelation treatments (Section 3.4.2). Intentions were revealed 2-4 weeks prior to the start of the commercialization period.

3. *Post-survey and administrative data:* In June 2014, a follow-up survey was administered to 10-12 randomly selected members of each cooperative (median size is 24, ranging from 4-91). Administrative data on all individual members' commercialization through the cooperative during the 2014 season were also collected from the group's books/records.

3.4.1 LFEs

The LFEs were based on two main treatments (see online appendix for detailed instructions): (1) a baseline coordination game (BCG) and (2) the same coordination game but with intentions revealed (i.e., communication/cheap talk) prior to play (CCG). These treatments were randomly assigned across experimental groups of subjects, which were created by randomly drawing members from existing cooperatives to form sets of players of size N equal to 10 or 20. All players in a given experimental group were members of the same cooperative. Table C.2 gives the sample distribution across the BCG and CCG. In the BCG, each player j had an endowment V_j of six chips. Each chip was worth 2000 West African francs/CFA (the equivalent of approximately 4 US dollars) if held individually. This was M. So, players were explained at the beginning of the game that they held an endowment of 12000 CFA. To mitigate windfall/house money effects, this endowment was framed as payment for the presurvey. The payoff for each chip sent to the experimental group was dependent on whether or not the threshold (T) was reached/surpassed. If $A \geq T$, each chip was worth 3000 CFA (the equivalent of H); if not, each chip was worth 500 CFA (the equivalent of L). So, each player had to decide how many of the six chips to send to the group (A_j) and how many to keep individually ($6 - A_j$), as shown in Figure 3.1.

Figure 3.1: LFEs – visual aid

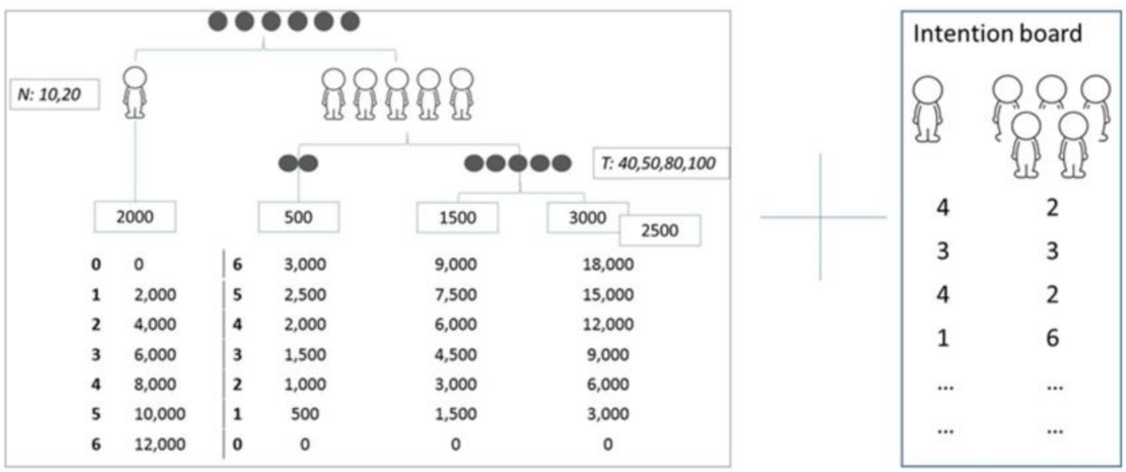

The CCG was identical to the BCG with one exception: prior to choosing and committing to A_j, each player was asked to reveal their intended action A'_j to the experimenter; that is, how much the player planned to send to the experimental group. This intention, which was confidentially revealed to all other players $-j$ alongside the aggregate intention on a board at the front of the room, indicated a given player's likely action. However, it was not a binding commitment and as such, other players did not know with full certainty that A_j would be the same as A'_j. It is in this sense that it was cheap talk.

Apart from the BCG and the CCG, we randomly varied N, T, H, and the presence of external uncertainty as separate treatments. We varied H and the presence of external uncertainty because prior literature has shown that attractiveness of the secure versus the risky action may impact coordination (recall Section 3.3.2). We also varied T because in a day-to-day context (in particular, the RCT) group size N is likely to be correlated with other characteristics such as the perceived likelihood of achieving the threshold. So, experimental variation in T enables us to better understand what mechanism is at play.

To see this, consider the following: If the level of T is fixed, it should be easier to coordinate in large groups than in small groups, since the former are more likely to reach the threshold. So, communication should be *less* informative in large (relative to small) groups for a fixed level of T. On the other hand, if communication is *more* informative in large groups, despite the level of T being fixed, this must be because of greater strategic uncertainty surrounding other players' actions. Similarly, if the average (i.e., per-player) number of chips needed to surpass the threshold is fixed across small and large groups, but communication matters differentially across N, this also suggests that strategic uncertainty is the main mechanism at play. Since we are unable to (1) experimentally control T in the RCT and (2) fully isolate these mechanisms in the day-to-day environment, we thus rely on the LFEs as a crucial part of our design in order to argue mechanisms.

Two primary aids were used when explaining the game. First, monetary payoffs were explained by displaying actual CFA bills on a board at the front of the room, also making "real stakes" more salient. Second, many hypothetical examples were used. For example, the experimenter and his assistant as well as pairs of subjects role-played through different scenarios. We also tested subject understanding by asking specific players to calculate such payoffs. A substantial part of the LFE sessions was dedicated to the instruction phase.

In the CCG, the exact same procedure as in the BCG was followed, except that prior to subjects making their actual decisions (A_j), the experimenter went around the room and asked players in private to reveal their intended actions (A'_j). Subjects were explained that this information would be collected by the experimenter and confidentially displayed in random order on a separate board at the front of the room. It was made clear that this was an intended, but non-binding action. Figure 3.1 also shows the logic behind the CCG. It was identical to the BCG, except for an additional board (right panel of Figure 3.1), which contained randomly ordered intentions A'_j and the aggregate A'.

The BCG and the CCG were implemented between sessions (and thus subjects), since introducing intentions mid-session would have complicated the protocol. The (experimental) group size, N, was fixed at either 10 or 20 during a session. So, N was varied across sessions (and subjects). The threshold, T, was 40 or 50 in 10-person groups and 40, 50, 80, or 100 in 20-person groups. T was varied randomly across rounds. H was either 3000 or 2500 CFA per chip and was varied randomly across rounds. Whether or not there was external uncertainty was implemented as follows. Subjects were explained that there was a 50 percent chance that due to bad luck H would be 1500 CFA per chip (instead of 2500 or 3000). This was varied across rounds by flipping a coin.[13]

[13] Each session comprised: (1) a presurvey collecting basic information (available upon request); (2) an introduction covering the purpose of the session and the fact that participants would be paid for decisions made during the session; (3) four rounds of decisions with no feedback, followed by debriefing; (4) a postsurvey collecting other information (available upon request); and (5) payment in private based on one randomly selected round. The sessions lasted 2.5-3 hours and average earnings were 9,500 CFA (approximately 20 US dollars), relative to a daily wage equivalent in this region of 5,000 CFA.

3.4.2 RCT

We worked with 79 groundnut-producing cooperatives – including 26 of the 28 with which we implemented the LFEs – comprising close to 2800 individuals. These cooperatives are part of two umbrella federations (i.e., conglomerates of cooperatives) in the main groundnut production zone of Senegal. From November to December 2013, two leaders of each farmer group attended a two-day training conducted by two development specialists. The training focused on the potential, pitfalls, and conduct of collective commercialization; in particular, strategies for identifying distant buyers, negotiating prices, and organizing transportation. Participants were instructed to conduct a meeting with all cooperative members upon returning to their village, to report the gist of what was covered during the training. Trainees were also provided with standardized booklets to keep records of each member's contribution to the group's sales in the upcoming commercialization season. A reward of 10,000 CFA was promised for filling in the booklets with all the requested information. All groups eventually received such reward.

After the training, during January 2014, enumerators went to the villages in order to elicit commercialization intentions from all cooperative members. Prior to doing so, they made sure that the leaders who had taken part in the training had held the briefing meeting. For each farmer group, all members who produced groundnuts for the 2014 commercialization season were asked how they intended to use their production. They had to split their anticipated harvest into (1) individual commercialization, (2) collective commercialization (via the cooperative), (3) inventories, and (4) other uses. They were told that the purpose of this survey was to better understand their decisions with regard to groundnut production. They were also informed that a subsequent group meeting would be held, where a message would be delivered to them. They were invited to attend that meeting. The 79 groups were allocated to one of four conditions, depending on the information that would be disclosed in the subsequent meeting (Table C.3):[14]

- In Condition A (the control group), members' intentions were not revealed. Enumerators said that a survey would be conducted after the end of commercialization. This was also announced in Conditions B-D.

- In Condition B, members' aggregate intentions were revealed.

- In Condition C, members' aggregate intentions as well as the distribution of intentions among members were revealed. I.e., how many members intended to contribute 100kg; how many intended to contribute 200kg, and so on. This was most comparable to the CCG in the LFEs.

- In Condition D, the same information as in C was revealed, but the distribution of intentions was disaggregated by ordinary members versus cooperative leaders (i.e., members who are part of the management committee). This treatment was inspired by the literature on leadership (e.g., Hermalin, 1998; Jack and Recalde, 2015; Potters et al., 2007).

[14]As discussed in some of the papers referenced in Section 3.3.2, messages have differential impacts depending on the "statistics"/content they convey. Hence, the logic for a variant of information treatments.

3.4.3 Sample description

An average of 10-12 randomly selected farmers per cooperative were surveyed before (time of eliciting intentions) and after commercialization. Table C.4 describes the characteristics of such individuals and compares means across groups that received no information about others' intentions and groups that did. Consistent with the context described in Section 3.2, our sample essentially comprises small farmers, with a total farm size of less then five hectares on average. The full sample cultivates groundnuts, but 58% also produce other crops. Groundnut production is also relatively small in magnitude, with the previous and current year's average expected harvest below 1.7 tons. 84% of the farmers indicate that they intend to sell some of their groundnut harvest through the group in the upcoming commercialization season. On average, intended sales through the group amounted to 60% of farmers' production.

Overall, Table C.4 does not show clear differences on pre-intervention characteristics across the two samples. There are, however, indications that generosity (as measured by a hypothetical dictator game) is higher on average for groups where no information was revealed. For a comparable groundnut-producing sample, Bernard et al. (2014) find that altruism is negatively related to group sales. So, in this case that might translate into an upward bias of the information effect. They find, however, that the associated coefficient is quite small in magnitude, in particular compared to the one associated with individuals' aversion to strategic uncertainty. Table C.4 also indicates that a greater number of farmers whose group was selected to participate in the LFEs were in the control group. All RCT-related estimations include these and other variables among the controls.[15]

3.4.4 Empirical strategy

With the LFE and RCT protocols being comparable, we rely on a common estimation strategy to assess the impact of others' intentions on the decision to contribute to the group. Individuals are indexed by subscript i and groups by subscript g – where "groups" refers to experimental sessions in the LFEs and day-to-day cooperatives in the RCT. Our basic estimation is as follows:

$$A_{ig} = \alpha + \beta T_g + \varepsilon_{ig} \tag{3.1}$$

where being exposed to others' intentions, which varies at the group level, is captured by a binary variable T_g (our proxy for A' in Section 3.3.1), and the associated parameter β measures its effect on the dependent variable of interest (more below). ε_{ig} is a composite error term defined as:[16] $\varepsilon_{ig} = \mu_g + \xi_{ig}$ where μ_g is a group-specific error and ξ_{ig} is the remaining idiosyncratic one. With group members' decisions to contribute to the group being strategic complements, we allow for within-group, individual errors to be correlated. Thus, all our standard errors are clustered at the group-level.[17]

While the same estimator is used for both the LFEs and the RCT, interpretation of the coefficients differs across the two. In LFE-related estimates, the β parameter is the treat-

[15] Table C.5 reports the same tests across Conditions A-D. The results are similar.

[16] In the estimating equations that follow, ε_{ig} will vary depending on the specification. But, for simplicity we keep the same notation.

[17] We rely on the cluster-correlated Huber-White covariance matrix method to compute cluster-robust standard errors. While our design includes a reasonably large number of clusters, we also computed standard errors based on more conservative randomized inference tests. Results are available upon request, and similar to those reported here.

ment effect on the treated. In fact, all individuals selected to participate in the game did participate and no one left before all rounds were completed (although that was allowed if they wanted to). In the RCT however, only 59% of group members attended the meeting in which aggregate intentions were revealed (last row of Table C.4). Without exogenous variation in attendance to the meeting, we estimate equation 3.1 on the full sample of farmers, regardless of whether or not they attended. The estimated β parameter therefore captures our "intent" to reveal or not aggregate intentions; i.e., an intent-to-treat effect.

Although we do not find evidence of major imbalances across conditions, we augment equation 3.1 with a vector of individual characteristics collected pre-intervention, X_{ig}, to both account for existing imbalances and enhance the precision of our estimates through reduced unexplained variance in our outcome variable. In LFE-related estimates, this vector of covariates includes age, sex, land size, education (dummy for going to a French school versus a Koranic or no schooling), and measures for risk, time (patience), and social preferences (generosity/altruism).[18] In RCT-related estimates, this includes age, sex, a dummy for whether the individual holds a leadership position in the group, the groundnut harvest (in kg), a dummy for whether the farmer produced other crops, land size, measures for risk, time, and social preferences, and participation in the LFEs. This leads to the following estimating equation:

$$A_{ig} = \alpha + \beta T_{ig} + X_{ig}^{'}\rho + \varepsilon_{ig} \tag{3.2}$$

[18] Generosity/social preference was elicited through a hypothetical dictator game (with a greater number indicating a more generous/altruistic individual). Risk preferences were elicited through a hypothetical Binswanger (1980)-style lottery (with a lower number indicating a more risk averse individual. Finally, patience was elicited through typical, hypothetical preference-over-time questions (with a one-day front-end delay and a higher number indicating a less patient individual).

Finally, in line with our framework, we test for heterogenous treatment effects with respect to group size by estimating equation 3.3, where S_g is a dummy for large groups ($N = 20$) in the LFE estimations and actual group size in the RCT estimations; β is the average effect of being exposed to aggregate intentions among all individuals in the sample; and δ captures the additional effect of this exposure for individuals in larger groups.

$$A_{ig} = \alpha + \beta T_g + \gamma S_g + \delta T_g \times S_g + X_{ig}^{'} \rho + \varepsilon_{ig} \tag{3.3}$$

For our main specifications, we estimate the effect of revealing aggregate intentions on both the extensive margin – with the dependent variable A_{ig} being a dummy for whether the individual contributed any amount to the group – and the intensive margin – where the dependent variable A_{ig} is the total amount that one contributed to the group. Both are consistent with the framework (and notation) discussed in Section 3.3.

In the RCT-related estimates, we also run several robustness checks. First, we test for significance of β and δ when the dependent variable is members' intentions. Because intentions were collected prior to any information revelation, we expect these parameters to be equal to 0. Second, one may be concerned about social desirability bias in individuals' reporting of collective groundnut sales. For both the intensive and extensive margin estimates, we thus also use effective individual sales from all group members, based on the group's own administrative data (recall the booklets/cooperative records discussed in Section 3.4.2); although such data do not include individual characteristics. Third, we also run some Tobit (to account for null reports) and inverse hyperbolic sine transformations specifications. Finally, we correct for multiple hypothesis testing (MHT) using the

Holm-Bonferroni method (see Tables 3.1 and 3.3).

3.5 Results

3.5.1 Cheap talk and coordination

Tables C.6 and 3.1 present our main results based on specifications 3.2 and 3.3. The first two columns are for the LFEs while the remaining ones are for the RCT. Our falsification test is presented in columns 3 and 4 where the dependent variable is one's intentions to contribute to the group, collected ahead of the intervention. In columns 5 and 6, we present the main results using data from individual surveys, while the last two columns rely on administrative data collected from the groups' books/records (there are no controls in this case).

In Table C.6 we assess the extensive margin – i.e., the effect of revealing others' intentions on an individual's decision to contribute part of one's resources (i.e., number of chips in the LFEs, kilograms/kgs of groundnuts in the RCT) or not to the group. We find no direct effect of revealing aggregate intentions, be it in the LFEs or in the RCT (columns 1 and 5); however, interacting with group size yields support for Hypothesis 3. In the LFEs, cheap talk (intentions) led to an 11 percentage point increase in the probability that one invests chips collectively, albeit limited to those sessions with groups of size 20 (as opposed to 10). The same is true in column 6, where revealing intentions increases the likelihood that one sold part of one's production through the cooperative by one percentage point for each additional group member. While insignificant, the direct effect of intentions is negative (-0.04 for LFEs and -0.08 for RCT), suggesting that the cheap-talk effect on increased coordination is limited to larger groups. In comparison, no such results are found in the falsification test (column 4). Further, we can fairly confidently rule out concerns of social desirability bias in the responses, since the results are very similar when using administrative data (column 8).

Results on the intensive margin are presented in Table 3.1 and show a similar pattern. While we do find a direct effect of cheap talk on the quantity of chips (column 1) and groundnuts (column 5) contributed to the group, this effect is only present in larger groups for both the LFEs (column 2) and the RCT (column 6). No such relationships are found in our placebo tests (columns 3 and 4) and the administrative data yield results comparable to those based on the individual survey data (columns 7 and 8).[19] We also correct for MHT (see adjusted p-values in square [..] brackets). The effects in the LFEs (columns 1 and 2) are no longer significant; however, the effects in the RCT persist (columns 5-8). The only exception is the direct effect of cheap talk when using administrative data, but recall that there are fewer (no) control variables in this specification, which also explains the lower R^2. Following Maniadis et al. (2014), we also calculate the post-study probability (PSP) associated with the main cheap-talk effect, i.e., the first and fifth columns of Table 3.1 for the LFEs and RCT respectively. Tables C.7 and C.8 illustrate that for low initial priors (i.e., below 0.2) the accuracy of the cheap-talk effect can be questioned. However, for moderate to high priors (i.e., above 0.2) we can be relatively confident that the main cheap-talk effect is accurate. [20]

[19] Results are robust to Tobit (Table C.9) and inverse hyperbolic sine (Table C.10).

[20] A separate issue is the presence of multiple observations per farmer for the LFEs, which may induce correlation among the residuals in Equation 3.2 and 3.3. We account for this by the use of standard errors clustered at the cooperative level. The fact that our main treatment is at the experimental group level, and thus does not vary over rounds for a farmer precludes the use of farmer fixed effects. We have, however, estimated the equations of interest using farmer random effects (see Table C.11), and the results are similar

Overall, Tables C.6 and 3.1 provide support for our predictions (in particular hypothesis 3): Cheap talk (intentions) can enhance coordination in situations of strategic uncertainty and this mainly seems to occur in larger groups, where coordination is more difficult to start with.[21] While consistent with our framework and prior work on the direct effect of cheap talk on coordination (recall theoretical and lab findings referenced in Sections 3.1 and 3.3.2), these results do not line up with Feltovich and Grossman (2015) who find that the beneficial effect of cheap talk decreases as the group size increases. However, as discussed previously, their findings are based on a threshold public-good game with messages that are different from ours. In addition, their group sizes are significantly smaller than ours.

As previously discussed, we are unable to exogenously vary group size in the RCT. Therefore, the effect in larger groups (particularly in the RCT) may be associated with other characteristics and thus, we should be cautious in interpreting the interaction effect. We address this concern in two main ways. First, we exploit the experimental design in order to separate the effect of group size from the potentially confounding threshold effect by using data from the LFEs. Table C.12 compares the group-size effect across three main specifications. Columns 1-2 use data from all LFE sessions, regardless of the threshold level required to achieve coordination. These are columns 1-2 from Table 3.1. Columns 3-4 use data from LFE sessions where the average (i.e., per-player) number of chips required to achieve the threshold is constant across group size. Finally, columns 5-6 use data from LFE sessions where the overall threshold is constant across group size. As

to not including these random effects.

[21] In the upper panel of Table C.13, we show that these results cannot be attributed to changes in individuals' risk, time (patience), and social preferences (generosity).

Table 3.1: Intensive margin: treatment effect on quantities contributed to the group (correcting for MHT)

Dependent variable	# Chips contributed		Contribution intentions (kg)		Quantity contributed (kg)		Quantity contributed (kg)	
Source	LFE		RCT - survey data		RCT - survey data		RCT - admin. data	
Treatment	0.40	-0.03	239.66	307.57	121.07	-78.93	129.47	-153.42
	(0.19)	(0.31)	(159.16)	(426.91)	(35.51)	(63.31)	(70.70)	(148.10)
	[.129]				[.006]		[.141]	
Size of Group	0.03	-0.31	-5.90	-3.72	7.18	0.75	-1.15	-9.30
	(0.24)	(0.27)	(5.57)	(10.15)	(1.85)	(1.28)	(2.18)	(3.53)
Treatment X Size		0.65		-2.44		7.19		8.83
		(0.38)		(11.62)		(2.23)		(3.57)
		[.142]				[.009]		[.062]
Control group mean	3.03	3.03	830.39	830.39	39.58	39.58	123.79	123.79
R^2	0.09	0.10	0.43	0.43	0.17	0.18	0.08	0.09
N	3,316	3,316	873	873	873	873	2,752	2,752

Notes: Standard errors are clustered at the group level. Holm-Bonferroni adjusted p-values are in square brackets. The correction for MHT is done for the direct impact of the treatment and the interaction with group size, across all specifications. Columns 1 and 2 use data from the LFEs. The dependent variable is the number of chips. The size in those specifications is a dummy variable that is 1 if the experimental session was run with 20 farmers and 0 if it was run with 10 farmers. Columns 3 and 4 use farmers' intended sales (of groundnuts) to the group, which were collected prior to the RCT interventions. Columns 5 and 6 use the self-reported (survey) collective commercialization. Columns 7 and 8 use administrative data on collective commercialization obtained from the cooperatives (recall booklets). In columns 3–8, the dependent variable is the quantity of groundnuts sold through the group and the size is the actual group size ranging from 4-91 (with median size being 24). Controls in the LFE regressions include age, sex, land size, education (dummy for going to a French school versus a Koranic or no schooling), generosity, risk aversion measures and time preferences elicited through hypothetical questions (see Table C.4), and dummies for the federation and whether the farmer is a leader. Controls in the RCT regressions with survey data include the same variables (except for education) and a dummy for whether the farmer produced crops other than groundnuts, all measured pre-intervention. A dummy for whether the individual indicated positive intentions to sell through the group is also added.

mentioned in Section 3.4.1, if communication is *more* informative in large groups, despite the average number of chips needed to surpass the threshold or the overall threshold being fixed across group size, strategic uncertainty is likely to be the main mechanism at play. Indeed, column 6 provides evidence that it is primarily about group size and its relation to strategic uncertainty. We also note that while the point estimate in column 4 is not statistically significant, it is positive and relatively large (in the order of 20 percent relative to the control group/CCG mean).

Second, Table C.14 compares group characteristics across above- and below-median day-to-day group size. We find limited evidence that these groups systematically differ from one another across size, except for being part of different umbrella federations, which we control for in all specifications. In short, we are reasonably confident that the group size effect is mainly associated with strategic uncertainty and less so with the threshold required to achieve coordination (which can be thought of as the amount required to fulfill contracts, e.g., when selling to large buyers in the day-to-day environment).

3.5.2 Collective commercialization and farmers' income

With the increased level of collective commercialization induced by the RCT intervention, we turn to its potential welfare effect. Specifically, we assess its impact on farmers' total income from the sale of groundnuts. This includes individual sales as well as those conducted through the cooperative. Since the intervention started shortly before the harvest, it did not affect farmers' production choices. Therefore, any effect of the intervention is mediated by a combination of (1) changes in the way farmers allocate their marketable surplus across individual versus collective sales and (2) changes in the price that may be obtained from collective sales due to greater aggregation.

The results are presented in Table 3.2. The dependent variable is the total revenue from selling groundnuts in thousands of CFA (1000 CFA was approximately equivalent to US$2 at the time of the study).[22] In column 1, we find no evidence of a direct effect of the intervention on farmers' income from groundnuts. In column 2 we extend the model to investigate whether the intervention differentially affected smaller versus larger farmers. We interact the treatment variable with the farmer's reported area of land that is dedicated to the production of groundnuts. While the direct effect of the intervention is large and positive, the coefficient associated with the interaction term is negative. In other words, larger farmers gain less from collective commercialization, in line with the existence of economies of scale in commercialization albeit with decreasing marginal returns. Finally, column 4 assesses the robustness of these findings by controlling for various individual and group characteristics. While the sign of the coefficients remain, they are no longer statistically significant. Nevertheless, the coefficient on the direct effect of the intervention does indicate that providing information on aggregate intentions led to a meaningful increase in groundnut revenues for farmers in the treatment group.

These results, combined with those in Section 3.5.1, point to the potential for a "coordination-based poverty trap" that is characterized by three key aspects. First, smaller (poorer) farmers have more to gain from collective commercialization. Second, collective commercialization requires a larger group of small farmers to surpass the threshold. Third, it

[22] Field work revealed that (1) quantities collected from farmers were quite noisy due to the lack of weighting scales for transactions that occur at the farm gate/local markets and (2) farmers tend to negotiate prices on the total quantity/amount. As such, we are unable to recover a per-unit price for those transactions conducted outside the group.

Table 3.2: Effect on revenues from groundnut sales

Dependent variable Source	Revenues from groundnut sales (CFA) RCT - survey data			
Treatment	15.57	185.11	7.91	61.05
	(73.25)	(56.22)	(51.46)	(39.07)
Size of land (ha)	3.48	45.62	0.85	14.28
	(3.55)	(20.68)	(1.16)	(16.38)
Treatment X Land size (ha)		-43.33		-13.70
		(20.83)		(16.47)
Control group mean	233.66			
R^2	0.01	0.06	0.34	0.34
N	868	868	868	868
Controls	No	No	Yes	Yes

Notes: Standard errors are clustered at the group level. The dependent variable in all the regressions is the total revenue from selling groundnuts in thousands of CFA (1000 CFA was approximately equivalent to US$2 at the time of the study). The last column adds as control variables age, sex, land size, generosity, risk aversion and time preference (see Table C.4), cooperative size, dummies for the federation, whether the farmer is a leader, and whether the farmer produced crops other than groundnuts, all measured pre-intervention, as well as a dummy for whether the farmer indicated positive intentions to sell through the group.

is more difficult to achieve coordination in larger groups. In short, it is precisely in larger groups consisting of more small farmers that coordination may be necessary but more difficult.

3.5.3 What information matters?

In order to assess whether the content of the information had an impact, we interact the aggregate intended actions (A') revealed by study participants with the main treatment; i.e., whether or not this information was conveyed to the group (Table C.15). Higher aggregate intentions lead to higher quantities sold to the group, but only in cooperatives where such information was revealed (column 2). First, this confirms that it is not just about "collecting intentions", but also "revealing them". Second, this finding supports the view that cheap talk helps farmers overcome coordination failure by shifting their beliefs.[23] One caveat to the analysis is that individual intentions and thus the aggregate, are endogenous, potentially reflecting other characteristics of the groups in question. The results are robust, however, to the inclusion of various individual and group characteristics (columns 3 and 4).

The RCT design also included variations on the type of information that was provided, which in turn allows us to explore which feature of aggregate intentions primarily enhances coordination (recall Table C.3). Table 3.3 presents the disaggregated impact of providing information on groundnut contributions to the cooperative, by informational treatment arm. Similarly to Table 3.1, we correct for MHT. The adjusted p-values in

[23] We did not explicitly elicit individual-level beliefs as that would have significantly complicated the implementation of the RCT.

Table 3.3: Effect on quantity sold by treatment condition (correcting for MHT)

Dependent variable Source	Contribution intentions (kg) RCT - survey data		Quantity contributed (kg) RCT - survey data		Quantity contributed (kg) RCT - admin. data	
B	224.66	384.10	76.54	-154.33	94.56	-349.51
	(158.64)	(391.49)	(38.01)	(92.63)	(83.92)	(178.26)
C	186.43	391.81	169.47	-141.87	117.05	-124.39
	(207.98)	(530.65)	(57.39)	(61.45)	(88.38)	(175.38)
D	207.89	72.79	121.45	46.96	177.40	-4.59
	(168.65)	(422.89)	(47.96)	(71.45)	(84.32)	(156.30)
Joint significance (p)			[.029]		[.226]	
Size of Group	-6.20	-4.60	7.39	1.26	-1.17	-8.74
	(5.46)	(9.78)	(1.81)	(1.29)	(2.19)	(3.43)
B X Size		-5.41		8.14		12.89
		(10.52)		(3.48)		(4.18)
C X Size		-7.69		11.40		7.70
		(14.01)		(2.28)		(3.60)
D X Size		4.30		3.01		6.28
		(12.43)		(2.06)		(3.73)
Joint significance (p)				[.000]		[.019]
Control group mean	830.39	830.39	39.58	39.58	123.79	123.79
R^2	0.44	0.44	0.18	0.21	0.09	0.10
N	873	873	873	873	2,752	2,752

Notes: Standard errors are clustered at the group level. Holm-Bonferroni adjusted p-values are in square brackets. The correction for MHT is done for the direct impact of the treatment and the interaction with group size, across all specifications. The first two columns test the experimental integrity using intentions of collective sales. The other columns use as the dependent variable the total sales, from survey (3 and 4) and administrative data (5 and 6) respectively. Size is the actual group size ranging from 4-91 (with median size being 24). Controls in the regression with survey data include age, sex, land size, generosity, risk aversion, and time preferences (see Table C.4), and dummies for the federation, whether the farmer is a leader, and whether the farmer produced crops other than groundnuts, all measured pre-intervention. A dummy for whether the farmer indicated positive intentions to sell through the group is also included for columns 3 to 6.

square [..] brackets suggests that the main results discussed below are robust, except for the direct effects estimated using administrative data (column 5). As before, our falsification tests in columns 1 and 2 show no relationship between one's intentions and the later provision of information. Results in column 3 indicate that provision of information increased group contributions in all three Conditions (B–D), although with significant differences in magnitude. In particular, while knowing others' aggregate intention leads to higher contributions (Condition B), the effect is more than two times larger in magnitude if one is also provided with the distribution of others' intentions (Condition C).

Two potential mechanisms may be at play here. First, one may trust a given aggregate more when it originates from a large number of individuals. For example, if farmers were to find out that most of the aggregate came from one individual's intended contribution, it would be reasonable to expect that this individual would be better off selling alone. So, revealing the distribution of intentions may help individuals refine their expectations. Second, individuals in the group may exhibit interdependent preferences, with group members' payoffs entering each others' utility functions, thus giving rise to social norms of equity and fairness (e.g., Manski, 2000; Sobel, 2005). While these effects may exist even in the absence of communication, they may be particularly salient when intentions and their distributions are revealed since they can be interpreted as signaling what other players consider "the right thing to do" (e.g., Bernheim, 1994). Examples of this type of norm and information signaling are Vesterlund (2003), Gächter et al. (2010), and Hill et al. (2012).

We find a smaller effect of revealing the distribution of intentions separately for leaders versus ordinary members (Condition D) in comparison to the overall distribution (Condition C). Leaders may be perceived as having superior information relative to ordinary members. In this case, that could mean that the leader for example has better knowledge of current market conditions.[24] So, it may be surprising that leaders' intentions lead to positive, but lower contributions to the group than the overall distribution (Condition C). To further unpack this finding, we conduct an analysis comparable to that in Table C.15, but now by informational treatment arm (see Table C.16). First, we assess whether the average amount that leaders or group ordinary members intend to sell matters. Second, we check whether leaders' intentions being greater than those of ordinary members matters. I.e., one reason why the impact of Condition D may be smaller than that of Condition C is because ordinary members expect leaders to contribute significantly high amounts to the group and upon learning that leaders' intentions do not match up with this, they reduce their contributions. Table C.16 suggests that the results are inconsistent with the above: Aggregate intentions, the overall distribution, and the distribution by leaders versus ordinary members all matter; however, the actual quantities revealed do not seem to matter. In other words, the findings do not support the idea that members in Condition D may have been deceived by the intentions provided/revealed by leaders.

3.5.4 Participation in games and day-to-day behavior

Our design further enables us to assess how participation in the LFE coordination games affects later behavior in the naturally-occurring RCT. After the LFEs, all participants were asked to provide feedback regarding the game they had played. The game was neutrally

[24] Bernard et al. (2014) find that 75% of members believed that leaders had better information about traders and/or prices in the region.

framed and there was no reference to collective marketing. Yet, many participants indicated that the coordination setting they had been faced with closely aligned with the situation they faced when deciding whether or not to sell their groundnut harvest through the cooperative. Recall that players were randomly selected from the full list of members for their group, but that each session only gathered players who were members of the same cooperative.[25] Considering the above as well as a related literature on transfer in coordination games (e.g., Cooper and Van Huyck, 2016), we hypothesize that experiencing coordination games with members of one's group may affect one's expectations vis-à-vis group members in subsequent, naturally-occurring, collaborative contexts.

We test for this in Table 3.4 where, in addition to the RCT treatment, we assess the impact of one's participation in the LFEs. We find no such effect in columns 1 and 2, where the dependent variable is one's intentions to commercialize collectively. In columns 3 and 4, however, we find a positive effect of participating in the LFEs, on top of the treatment effects that were previously identified. Importantly, there is no effect at the group level (i.e., whether the cooperative was selected for the LFEs), but there is a clear effect for whether the individual was (randomly) selected for participation in the LFEs. In fact, the results hold even when restricting our sample to only those cooperatives that were selected for the LFEs. Furthermore, this effect is independent of the treatment: In column 7, we restrict our sample to those cooperatives selected for the LFEs, but that belonged to the control group in the RCT (where no intentions were revealed). The effects are large in

[25]This was mainly driven by the fact that cooperatives are typically defined at the village level. With villages sometimes being far apart, we chose to (1) organize each session in the village school and (2) avoid further complications such as communicating decisions by phone (e.g., Hill et al., 2012). In Table C.17 we report tests for equality of characteristics between LFE participants and non-participants within groups that were selected for the LFEs. Similarly, we provide in Table C.18 tests for equality between the LFE cooperatives and those that were only in the RCTs.

magnitude, with participation in the LFEs leading to an additional 74 kgs of groundnuts being contributed to the cooperative during the (naturally-occurring) commercialization season. [26]

In summary, while the LFEs are a useful diagnostic tool to test our hypotheses in a highly controlled environment, they seem to improve farmers' understanding of the strategic complementarities that exist, which in turn impacts collective commercialization.

3.6 Conclusion

This paper tests selected drivers of coordination among members of farmer cooperatives (groundnut producers) in rural Senegal. Consistent with our framework and prior literature, we find the following. First, cheap talk increases coordination (specifically, sales through the cooperative), but primarily in larger groups. Second, there is suggestive evidence that revelation of intentions leads to higher revenues from groundnut sales for smaller farmers. Third, participation in the "lab" impacts subsequent behavior in the "field".

Given the simplicity of this intervention, one might wonder why these cooperatives, and

[26] Tables and C.17 and C.18 show that some of the characteristics differ across cooperatives in the LFEs and those that were in the RCTs only, and similarly between those who attended the games and those who did not, within LFE cooperatives. More importantly, the omnibus test performed in table C.18 failed to reject that those characteristics jointly are significantly associated with being in an LFE cooperative. Such characteristics, if they were collected prior to the RCT, are controlled for in the specifications in the main tables, e.g., Table 3.4. One additional characteristic that differs across these groups is the attendance at the intention revelation meeting. This is potentially problematic, as participation in the meeting is the main mechanism through which the intervention affects behavior, although there is nothing that prevents cooperatives to reflect on, and share thoughts on what was discussed at the meeting. In that way, information could reach those who were not in the meeting as well. However, this variable was only collected long after the meeting, i.e.in the post-commercialization survey. For this reason, the variable on the meeting attendance is not included in the regressions of which the results are presented in this article. Regressions that include this variable show, however, little difference with respect to the effect of participation in the LFE on the behavior in the RCT.

Table 3.4: Participation in LFEs and behavior in RCT

Dependent variable	Contribution intentions (kg)		Quantity contributed (kg)		Quantity contributed (kg)		
Sample	All Cooperatives				Game Cooperatives		Control Cooperatives
Farmer was in LFEs	-136.32	-181.08	68.60	95.83	73.27	92.20	74.17
	(367.59)	(174.88)	(63.71)	(43.26)	(66.81)	(43.83)	(33.96)
Group was in LFEs	15.19	151.03	110.59	49.44			-3.93
	(412.55)	(162.18)	(68.91)	(53.55)			(17.96)
Treatment	151.17	204.15	130.40	122.53	204.01	140.04	
	(264.67)	(152.32)	(38.22)	(36.01)	(87.41)	(66.63)	
LFEs non participants mean	996.49	996.49	88.40	88.40	181.46	181.46	10.75
R^2	0.01	0.44	0.05	0.18	0.04	0.25	0.23
N	873	873	882	873	279	279	194
Controls	No	Yes	No	Yes	No	Yes	Yes

Notes: Standard errors are clustered at the group level. In the first two columns, intentions on collective commercialization (in kg) are used as the dependent variable. For all other regressions, the dependent variable is the actual quantity sold through the cooperative. All regressions include as controls age, sex, land size, generosity, risk aversion, and time preferences (see Table C.4), cooperative size, and dummies for the federation, whether the farmer is a leader, and whether the farmer produced crops other than groundnuts, all measured pre-intervention.

arguably others, have not already attempted this potential solution to the coordination problem. One possible explanation is that the study was carried out by "others" in an environment that is characterized by important interpersonal ties. So, participants may have been more willing to reveal and believe intended actions than had the process been implemented by familiar peers who could have had "hidden agendas". Future research can explore related issues further.

Chapter A: Appendix for chapter 1

Figure A.1: Evolution of outcomes over time

Figure A.2: Distribution of the proxy before the policy

Proxy for experience in 2013Q3, just before the policy

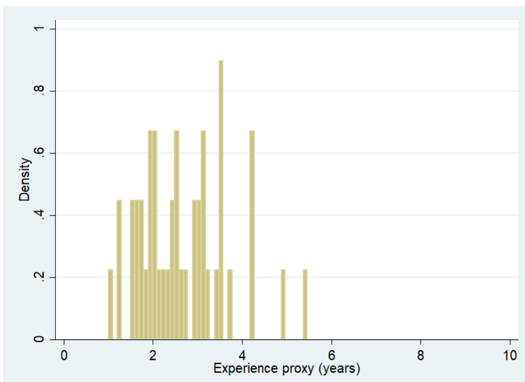

Eligible cohorts

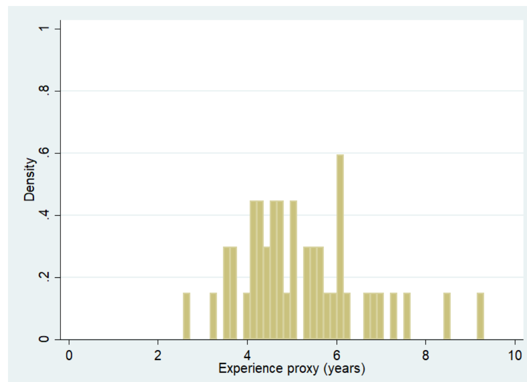

Ineligible cohorts

Proxy for experience in 2011-2013Q3, just before the policy

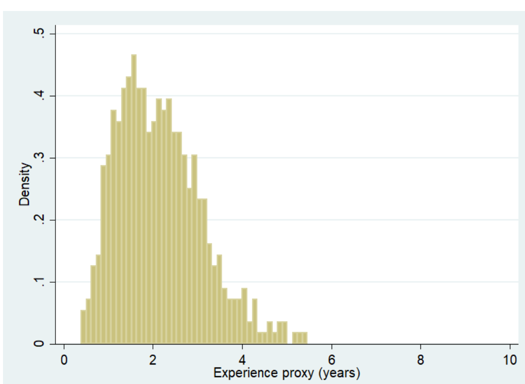

Eligible cohorts

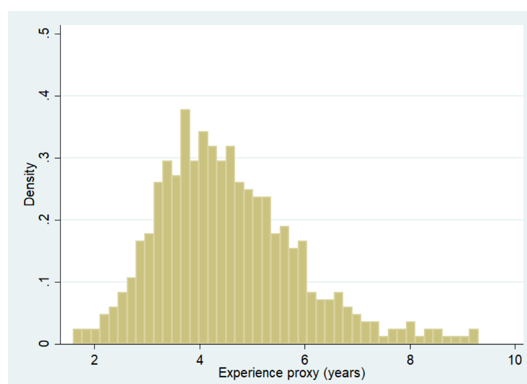

Ineligible cohorts

Figure A.3: Evolution of outcomes over time for women

Figure A.4: Evolution of outcomes over time for men

Table A.1: Cohorts and eligibility

	Young cohort					Old cohort				
Age in 2013	25	26	27	28	29	30	31	32	33	34
Year of birth	88	87	86	85	84	83	82	81	80	79
Year										
2013Q1-2013Q3										
2013Q3-2014	X	X	X	X	S					

Notes: S denotes that some of the individuals are eligible. X denotes that all the individuals are eligible.

Table A.2: Detail summary statistics of outcomes of interest and control variables

Periods	Before policy			After policy			
Cohort	Ineligible	Eligible	Diff (t-stat)	Ineligible	Eligible	Diff (t-stat)	DID (t-stat)
Private formal employment	32.22	27.20	-5.02	33.10	29.40	-3.21	0.67
			(-7.94)			(-3.48)	(0.67)
Private employment	43.32	35.79	-7.53	43.94	38.94	-4.62	1.16
			(-12.21)			(-4.50)	(1.11)
Employment	55.19	43.90	-11.29	57.59	48.95	-7.76	1.84
			(-18.17)			(-8.22)	(1.91)
Participation	63.34	54.50	-8.85	65.15	58.14	-6.39	0.70
			(-14.78)			(-7.02)	(0.74)
Gender is female	48.86	49.97	1.11	48.76	48.84	0.14	-1.14
			(1.96)			(0.14)	(-1.26)
African	82.27	82.55	0.28	82.15	82.67	0.31	-0.36
			(0.52)			(0.41)	(-0.43)
No schooling	1.57	0.88	-0.69	1.63	1.26	-0.34	0.39
			(-4.82)			(-1.86)	(1.87)
Primary or less	9.57	7.87	-1.70	9.04	7.13	-1.90	-0.22
			(-4.81)			(-4.03)	(-0.40)
Secondary	42.10	42.56	0.46	41.20	41.27	-0.19	-0.48
			(0.68)			(-0.21)	(-0.45)
Matric	31.36	35.57	4.21	31.49	35.10	3.81	-0.33
			(6.86)			(4.11)	(-0.32)
Post-secondary	14.46	12.42	-2.04	15.79	14.39	-1.41	0.36
			(-2.99)			(-1.73)	(0.45)
Gauteng	28.62	26.38	-2.24	29.07	26.55	-2.06	-0.36
			(-2.92)			(-2.04)	(-0.37)
Single	57.37	74.67	17.30	53.92	70.25	15.29	-0.91
			(26.64)			(14.20)	(-0.90)
Living with a partner	14.44	11.87	-2.57	14.93	12.60	-2.48	-0.51
			(-5.26)			(-3.23)	(-0.63)
Married	26.61	12.93	-13.68	29.37	16.49	-11.77	1.49
			(-21.39)			(-13.46)	(1.74)
HH has an employed member	54.04	59.51	5.47	53.33	58.65	4.92	-0.40
			(9.27)			(5.26)	(-0.38)
HH has a pensioner	19.92	21.48	1.57	20.10	20.60	0.26	-0.69
			(3.21)			(0.37)	(-0.84)

Table A.3: Effects on unemployment

	(1)	(2)	(3)	(4)
δ_1	-2.490	-1.887	-0.075	0.234
	(0.885)***	(0.838)**	(0.923)	(0.847)
Experience proxy			-13.737	-6.247
			(1.255)***	(0.668)***
Experience proxy square			0.824	0.567
			(0.093)***	(0.056)***
R^2	0.01	0.15	0.05	0.16
N	32,861	32,861	32,861	32,861
Controls	N	Y	N	Y
Average young before policy	19.4			

Notes: Data are from the QLFS 2011-2014. Sample: individuals born in 1979-1988 from the incoming rotation group. δ_1 is the difference-in-difference estimator (see equation 1.1). Standard errors in brackets are clustered at the stratum level. Regressions are weighted using the sample weights from the QLFS and estimated using a linear probability model. Controls are race (Blacks, Colored, Asian/Indians and Whites), gender, education (no education, some primary, primary completed, some secondary, secondary completed, post-secondary), marital status, the presence of an individual older than 65 and the presence of another individual who is working, all included in the regression equation as dummy variables.

Table A.4: Robustness checks

	Main results	Young workers eligible for 2 years	Interacting education with time	Interacting education with cohort	Different experience functions
	(1)	(2)	(3)	(4)	(5)
A: Employment					
δ_1	-0.735	-1.248	-0.835	-0.735	-0.694
	(0.893)	(1.067)	(0.896)	(0.893)	(0.908)
R^2	0.15	0.15	0.15	0.15	0.15
B: Private employment					
δ_1	-0.765	-0.172	-0.868	-0.777	-0.633
	(0.979)	(1.224)	(0.978)	(0.978)	(0.981)
R^2	0.13	0.12	0.13	0.13	0.13
C: Private formal employment					
δ_1	-0.782	-0.501	-0.862	-0.785	-0.530
	(0.927)	(1.119)	(0.928)	(0.926)	(0.956)
R^2	0.14	0.14	0.14	0.14	0.14
N	56,453	40,596	56,453	56,453	56,453

Notes: Data are from the QLFS 2011-2014. Sample: individuals born in 1979-1988, from the incoming rotation group. Column 1 presents the main results, as obtained in table 1.2, column 4. Column 2 restricts the analysis on workers born between 1986-1988 and between 1981-1983, allowing to test the effects on young workers eligible for two years of subsidies. Column 3 presents results of specifications that include interaction terms between education and year dummies, and column 4 includes interaction terms between education and a dummy for eligible cohorts. Column 5 allows for different effects of experience for eligible and ineligible cohorts. δ_1 is the difference-in-difference estimator (see Equation 1.1). Standard errors in brackets are clustered at the stratum level. Regressions are weighted using the sample weights from the QLFS and estimated using a linear probability model. Controls include race (Blacks, Colored, Indians, Whites), gender, marital status, the presence of an individual older than 65 and the presence of another individual who is working, all included in the regression equation as dummy variables.

Table A.5: Testing the sensitivity of the results to the functional form of experience proxy

	(1)	(2)	(3)	(4)	(5)
A: Overall employment					
δ_1	1.350	1.847	-0.637	-0.846	-0.472
	(0.887)	(0.892)**	(0.913)	(0.905)	(0.916)
Experience proxy		2.923	8.439	10.750	
		(0.356)***	(0.871)***	(1.523)***	
Experience proxy square			-0.602	-1.193	
			(0.076)***	(0.358)***	
Experience proxy cubic				0.042	
				(0.026)	
R^2	0.14	0.15	0.15	0.15	0.15
B: Private employment					
δ_1	0.743	1.168	-0.520	-0.657	-0.552
	(0.970)	(0.972)	(0.994)	(0.979)	(0.977)
Experience proxy		2.503	6.253	7.765	
		(0.322)***	(0.791)***	(1.427)***	
Experience proxy square			-0.409	-0.796	
			(0.070)***	(0.342)**	
Experience proxy cubic				0.028	
				(0.025)	
R^2	0.12	0.12	0.13	0.13	0.12
C: Private formal employment					
δ_1	0.387	0.884	-0.762	-0.826	-0.885
	(0.921)	(0.925)	(0.939)	(0.941)	(0.940)
Experience proxy		2.927	6.582	7.294	
		(0.301)***	(0.705)***	(1.252)***	
Experience proxy square			-0.399	-0.581	
			(0.065)***	(0.307)*	
Experience proxy cubic				0.013	
				(0.023)	
R^2	0.13	0.13	0.13	0.13	0.13
N	56,453	56,453	56,453	56,453	56,453

Notes: Data are from the QLFS 2011-2014. Sample: individuals from the incoming rotation group, born in 1979-1988. δ_1 is the difference-in-difference estimator (see Equation 1.1). This table tests for the sensitivity for the functional form of the experience proxy. In column 5, deciles of the distribution of the experience proxy are used, to account for a non-parametric functional form.

Table A.6: Heterogeneous effects by gender

	Overall employment		Private employment		Private formal employment	
	(1)	(2)	(3)	(4)	(5)	(6)
A: Men						
δ_1	0.838	0.425	1.382	1.025	1.795	1.729
	(1.403)	(1.292)	(1.378)	(1.300)	(1.352)	(1.311)
Experience proxy	11.615	7.852	8.872	4.944	9.664	3.548
	(1.341)***	(1.015)***	(1.332)***	(1.070)***	(1.280)***	(1.067)***
Experience proxy square	-0.684	-0.595	-0.451	-0.363	-0.456	-0.265
	(0.113)***	(0.094)***	(0.116)***	(0.096)***	(0.119)***	(0.100)***
Controls	N	Y	N	Y	N	Y
R^2	0.04	0.17	0.02	0.13	0.02	0.13
N	26,836	26,836	26,836	26,836	26,836	26,836
B: Women						
δ_1	-1.509	-1.857	-1.960	-2.516	-3.031	-3.379
	(1.455)	(1.326)	(1.509)	(1.385)*	(1.429)**	(1.267)***
Experience proxy	12.398	9.015	9.886	6.565	9.855	5.119
	(1.232)***	(0.979)***	(1.062)***	(0.855)***	(1.036)***	(0.769)***
Experience proxy square	-0.724	-0.659	-0.532	-0.529	-0.508	-0.437
	(0.105)***	(0.097)***	(0.095)***	(0.083)***	(0.094)***	(0.074)***
Controls	N	Y	N	Y	N	Y
R^2	0.03	0.14	0.02	0.10	0.02	0.14
N	29,617	29,617	29,617	29,617	29,617	29,617

Notes: Data are from the QLFS 2011-2014. Sample: individuals in the incoming rotation group, born in 1979-1988. δ_1 is the difference-in-difference estimator (see Equation 1.1). Standard errors in brackets are clustered at the stratum level. Regressions are weighted using the sample weights attached to the QLFS and estimated using a linear probability model. Controls include race, education, marital status, the presence of an individual older than 65 and the presence of another individual who is working, all included in the regression equation as dummy variables.

Table A.7: Heterogeneous effects by race

	Overall employment		Private employment		Private formal employment	
	(1)	(2)	(3)	(4)	(5)	(6)
A: Blacks						
δ_1	-0.490	-0.925	-0.483	-1.004	-0.891	-1.133
	(1.099)	(1.023)	(1.193)	(1.123)	(1.117)	(1.046)
Experience proxy	10.093	7.966	7.288	5.241	7.614	3.798
	(1.212)***	(0.892)***	(1.028)***	(0.778)***	(0.887)***	(0.681)***
Experience proxy square	-0.580	-0.535	-0.384	-0.374	-0.374	-0.297
	(0.110)***	(0.097)***	(0.096)***	(0.083)***	(0.086)***	(0.073)***
Controls	N	Y	N	Y	N	Y
R^2	0.03	0.13	0.02	0.10	0.02	0.10
N	45,957	45,957	45,957	45,957	45,957	45,957
B: Whites						
δ_1	0.393	-0.200	-0.061	-0.672	2.293	1.565
	(4.180)	(3.811)	(4.062)	(4.101)	(4.083)	(4.006)
Experience proxy	9.848	9.835	9.134	9.409	6.598	6.739
	(2.210)***	(2.372)***	(2.531)***	(2.666)***	(2.528)**	(2.771)**
Experience proxy square	-0.765	-0.760	-0.654	-0.682	-0.426	-0.436
	(0.205)***	(0.203)***	(0.220)***	(0.220)***	(0.221)*	(0.231)*
Controls	N	Y	N	Y	N	Y
R^2	0.02	0.12	0.01	0.08	0.01	0.07
N	3,221	3,221	3,221	3,221	3,221	3,221

Notes: Data from the QLFS 2011-2014. Sample: individuals born in 1979-1988, from the incoming rotation group. δ_1 is the difference-in-difference estimator (see equation 1.1). Standard errors in brackets are clustered at the stratum level. Regressions are weighted using the sample from the QLFS and estimated using a linear probability model. Controls include gender, education (no education, some primary, primary completed, some secondary, secondary completed, post-secondary), marital status, the presence of an individual older than 65 and the presence of another individual who is working, all included in the regression equation as dummy variables.

Table A.8: Heterogeneous effects by level of education

	Overall employment		Private employment		Private formal employment	
	(1)	(2)	(3)	(4)	(5)	(6)
A: Individuals with less than a matric						
δ_1	-0.209	-0.769	-0.018	-0.812	-0.515	-0.954
	(1.454)	(1.359)	(1.425)	(1.300)	(1.194)	(1.118)
Experience proxy	7.957	6.409	5.987	4.203	6.076	2.421
	(1.260)***	(0.988)***	(1.125)***	(0.893)***	(0.827)***	(0.736)***
Experience proxy square	-0.403	-0.420	-0.242	-0.282	-0.236	-0.180
	(0.109)***	(0.096)***	(0.101)**	(0.086)***	(0.083)***	(0.075)**
Controls	N	Y	N	Y	N	Y
R^2	0.02	0.11	0.01	0.13	0.02	0.10
N	30,020	30,020	30,020	30,020	30,020	30,020
B: Individuals with a matric or more						
δ_1	-0.454	-0.852	-0.451	-0.905	-0.531	-0.850
	(1.450)	(1.321)	(1.597)	(1.490)	(1.563)	(1.456)
Experience proxy	14.879	10.790	12.002	7.478	12.098	6.397
	(1.242)***	(0.984)***	(1.291)***	(1.003)***	(1.325)***	(1.028)***
Experience proxy square	-1.000	-0.839	-0.726	-0.582	-0.697	-0.512
	(0.108)***	(0.091)***	(0.117)***	(0.096)***	(0.121)***	(0.098)***
Controls	N	Y	N	Y	N	Y
R^2	0.04	0.13	0.03	0.11	0.02	0.11
N	25,997	25,997	25,997	25,997	25,997	25,997

Notes: Data from the QLFS 2011-2014. Sample: individuals born in 1979-1988, from the incoming rotation group. δ_1 is the difference-in-difference estimator (see equation 1.1). Standard errors in brackets are clustered at the stratum level. Regressions are weighted using the sample from the QLFS and estimated using a linear probability model. Controls include race (Blacks, Colored, Asians/Indians, Whites), gender, marital status, the presence of an individual older than 65 and the presence of another individual who is working, all included in the regression equation as dummy variables.

Table A.9: Heterogeneous effects by industry

Main industry	Trade	Community and social	Financial	Manufac-turing	Construction
δ_1	-0.593	0.705	0.337	-0.619	-0.688
	(0.668)	(0.628)	(0.531)	(0.489)	(0.447)
R^2	0.02	0.08	0.06	0.03	0.04
N	56,453	56,453	56,453	56,453	56,453
Eligible cohorts, before policy	11.5	7.4	6.5	5.5	3.8
	Private households	Transport	Agriculture	Mining	
δ_1	-0.021	0.095	0.307	-0.031	
	(0.352)	(0.370)	(0.307)	(0.249)	
R^2	0.03	0.02	0.04	0.05	
N	56,453	56,453	56,453	56,453	
Eligible cohorts, before policy	2.3	2.5	2.0	1.0	

Notes: Data are from the QLFS 2011-2014. Sample: individuals from the incoming rotation group, born in 1979-1988. δ_1 is the difference-in-difference estimator (see Equation 1.1). For each column, the dependent variable is a dummy variable which is 1 if when an individual is employed in the mentioned industry, and zero otherwise. Regressions are estimated using a linear probability model weighted using the sample from the QLFS. Standard errors in brackets are clustered at the stratum level. Controls are race (Blacks, Colored, Asians/Indians and Whites), gender, education (no education, some primary, primary completed, some secondary, secondary completed and post-secondary), marital status, the presence of an individual older than 65 and the presence of another individual who is working, all included in the regression equation as dummy variables.

Table A.10: Effects and local labor market conditions

	Employment (1)	Private (2)	Priv. formal (3)	Participation (4)
A: Main effects				
Young cohorts * $post$	-0.735	-0.765	-0.782	-0.952
	(0.893)	(0.979)	(0.927)	(0.926)
R^2	0.15	0.13	0.14	0.09
B: Adding unemployment rate				
Young cohorts * $post$	-0.756	-0.783	-0.811	-0.918
	(0.894)	(0.980)	(0.927)	(0.928)
Unemp. rate (35+)	-0.229	-0.203	-0.321	0.370
	(0.230)	(0.226)	(0.190)*	(0.253)
R^2	0.15	0.13	0.14	0.09
N	56,453	56,453	56,453	56,453
C: Interacting with unemployment rate				
Young cohorts * $post$	-1.054	-2.291	0.034	-1.421
	(1.835)	(2.051)	(2.050)	(1.820)
Unemp. rate (35+)	-0.442	-0.364	-0.445	0.213
	(0.266)*	(0.246)	(0.208)**	(0.294)
Young cohorts * $post$ * Unemp. rate (35+)	0.037	0.128	-0.120	0.015
	(0.178)	(0.190)	(0.179)	(0.184)
R^2	0.15	0.13	0.14	0.09
N	56,453	56,453	56,453	56,453

Notes: Data are from the QLFS 2011-2014. Sample: individuals from the incoming rotation group, born in 1979-1988. This table tests for heterogeneous effects by local market conditions, proxied by the unemployment rate of individuals aged 35 and above. Panel A presents the main results (as in Table 1.2). Panel B presents the results when the unemployment rate of individuals aged 35 and above is added. Panel C presents the results when the variable of interest, the interaction of young cohorts and the dummy variable for periods after the start of the ETI, are interacted.

Table A.11: Impact of the ETI on youth employment: main results, sensitivity to alternative clustering

	Employment		Private		Private formal		Participation	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
δ_1	-0.206	-0.735	-0.137	-0.765	-0.485	-0.782	-0.502	-0.952
	(1.057)	(0.809)	(1.047)	(0.847)	(1.207)	(0.896)	(1.054)	(0.919)
Experience proxy	12.120	8.549	9.524	5.854	9.860	4.390	5.601	6.330
	(1.029)***	(0.721)***	(1.050)***	(0.673)***	(1.387)***	(0.668)***	(0.975)***	(0.851)***
Experience proxy square	-0.716	-0.622	-0.507	-0.441	-0.492	-0.347	-0.377	-0.412
	(0.094)***	(0.069)***	(0.100)***	(0.064)***	(0.134)***	(0.062)***	(0.085)***	(0.070)***
R^2	0.03	0.15	0.02	0.13	0.02	0.14	0.01	0.09
N	56,453	56,453	56,453	56,453	56,453	56,453	56,453	56,453
Controls	N	Y	Y	Y	Y	Y	Y	Y
P-value with wild bootstrap	0.871	0.439	0.896	0.615	0.690	0.463	0.704	0.393

Notes: Data are from the QLFS 2011-2014. Sample: individuals from the incoming rotation group, born in 1979-1988. δ_1 is the difference-in-difference estimator (see Equation 1.1). Standard errors in brackets are clustered at the province * cohort level, to take into consideration treatment occurring at the cohort level, possible correlation within provincial labor markets, and the construction of the experience proxy at the province * cohort level. The last line tests for the significance of the difference-in-difference estimators, using a wild bootstrap with clustering only at the cohort level.

Chapter B: Appendix for chapter 2

Figure B.1: Income distribution among migrants and non-migrants

All Africans

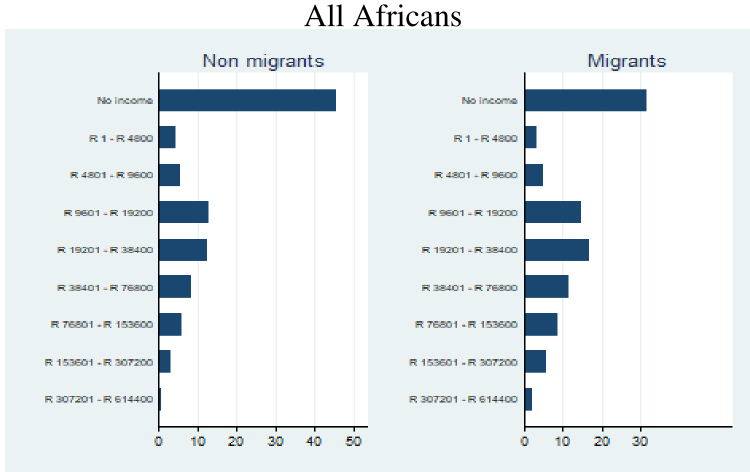

Employed Africans only

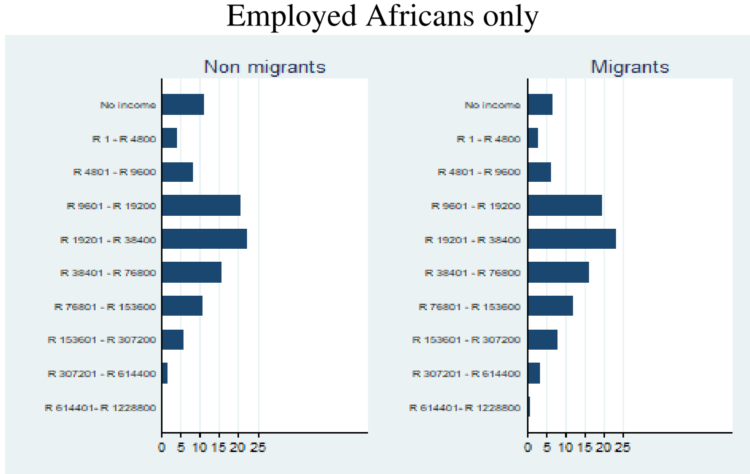

Notes: Data: South Africa 2011 census. Sample: all African internal migrants aged 15 to 64 in 2011. Averages for established migrants are computed for those aged 25 or above. Income is self-reported from all sources.

Figure B.2: Employment rates of migrants and non-migrants of same origin and language group

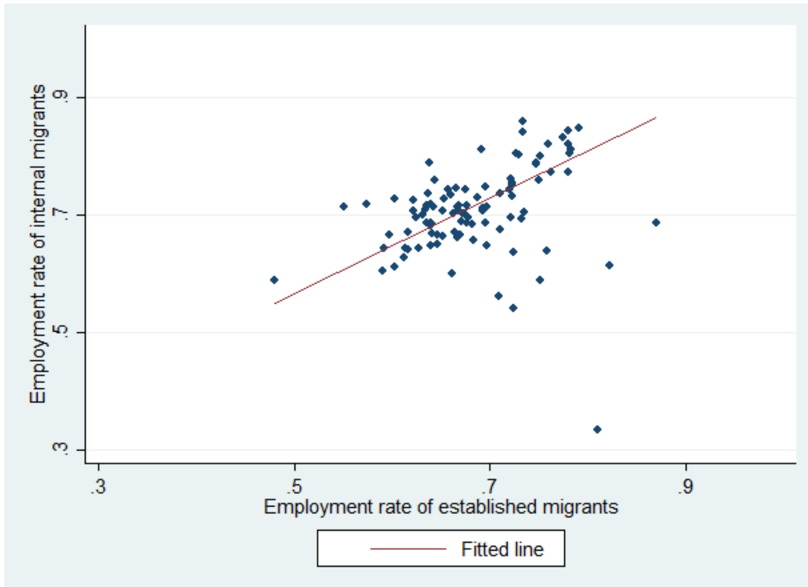

Notes: Data: South Africa 2011 census. Sample: all African internal migrants aged 15 to 64 in 2011. Averages for established migrants are computed for those aged 25 or above.

Table B.1: Province of birth and internal migration incidence

	Whole population	Africans
Western cape	5.74	6.05
Eastern cape	13.33	13.87
Northern cape	12.89	12.53
Free state	11.49	9.87
Kwazulu-natal	9.78	9.13
North west	12.20	11.19
Gauteng	10.35	7.13
Mpumalanga	10.57	9.75
Limpopo	13.90	13.68

Data: South Africa 2011 census. Sample: individuals aged 15 to 64 in 2011. Averages for established migrants are computed for those aged 25 or above.

Figure B.3: Occupational choices of migrants and non-migrants of same origin and language group

Notes: Data: South Africa 2011 census. Sample: all African internal migrants aged 15 to 64 in 2011. Averages for established migrants are computed for those aged 25 or above.

Table B.2: Language and internal migration incidence

	Whole population	Africans
Afrikaans	12.07	10.80
English	12.60	14.76
IsiNdebele	10.01	9.97
IsiXhosa	12.63	12.62
IsiZulu	8.59	8.58
Sepedi	12.60	12.60
Sesotho	9.55	9.52
Setswana	10.75	10.70
SiSwati	8.11	8.09
Tshivenda	13.28	13.28
Xitsonga	11.69	11.65

Data: South Africa 2011 census. Sample: individuals aged 15 to 64 in 2011. Averages for established migrants are computed for those aged 25 or above.

Table B.3: Province of birth and current province of residence

Province of birth	Current province of residence								
	WC	EC	NC	FS	KN	NW	GA	MP	LI
Western cape	43.05	12.86	1.65	1.65	4.75	2.82	28.4	2.89	1.93
Eastern cape	28.28	26.55	0.55	2.25	10.99	6.08	21.73	2.08	1.5
Northern cape	5.72	2.1	48.46	6.17	2.1	12.13	19.29	2.63	1.4
Free state	2.5	1.41	2.35	37.89	2.35	9.4	38.18	4.41	1.51
Kwazulu-natal	1.44	1.56	0.21	1.06	53.98	1.45	34.7	4.76	0.84
North west	0.9	0.28	7.56	2.13	0.56	44.92	38.06	1.77	3.83
Gauteng	3.54	1.8	0.6	2.72	2.51	7.76	70.39	6.32	4.35
Mpumalanga	0.96	0.6	0.37	0.95	1.61	4.43	44.13	40.75	6.2
Limpopo	0.6	0.42	0.28	0.72	0.42	4.89	61.16	6.84	24.67
Total	8.36	7.17	2.07	4.25	12.8	8.5	41.93	7.65	7.27

Notes: The provinces are : WC : Western cape, EC : Eastern cape, NC : Northern cape, FS : Free state, KN : Kwazulu-natal NW : North west, GA : Gauteng, MP : Mpumalanga, LI : Limpopo.

Table B.4: Employment and ethnic capital, results at the ethnic group - municipality level

	(1)	(2)	(3)	(4)	(5)
Migrant own characteristics (dummy variables)					
No education or some primary	0.185		0.156		0.059
	(0.042)***		(0.035)***		(0.023)**
Post-secondary	0.351		0.366		0.306
	(0.023)***		(0.024)***		(0.020)***
Migrant is male	0.270		0.255		0.224
	(0.019)***		(0.018)***		(0.016)***
Age 25-44	0.188		0.186		0.212
	(0.031)***		(0.033)***		(0.020)***
Age 45-65	0.197		0.189		0.269
	(0.045)***		(0.046)***		(0.029)***
Ethnic group average characteristics (share)					
Ethnic group share with primary or no education		0.153	0.125	0.067	0.060
		(0.041)***	(0.032)***	(0.024)***	(0.021)***
Ethnic group share with post-secondary education		0.344	0.117	0.215	0.106
		(0.046)***	(0.028)***	(0.029)***	(0.024)***
Ethnic group share of male		0.118	0.043	0.104	0.051
		(0.033)***	(0.027)	(0.020)***	(0.018)***
Ethnic group share 45-64		0.018	-0.059	0.004	-0.020
		(0.034)	(0.026)**	(0.020)	(0.019)
R^2	0.20	0.06	0.22	0.42	0.52
N	5,542	4,140	4,140	4,135	4,135
Significance of ethnic group characteristics		16.56	6.27	19.43	7.13
P-value of joint significance test		0.00	0.00	0.00	0.00
Ethnic group fixed effects	No	No	No	Yes	Yes
Local municipality fixed effects	No	No	No	Yes	Yes

Notes: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Data: South Africa 2011 census. Sample: all Black South African internal migrants aged 15 to 64 in 2011, aggregated values at the ethnic group - municipality level. Averages for established migrants are computed for those aged 25 or above. The table presents the OLS estimates of the regression of the share of employment in each group on individual characteristics and the characteristics of established migrants of the same ethnic group. Specification 4 and 5 include and fixed effects for the ethnic group and municipality. The F-tests presented are for the joint significance of the characteristics of established migrants.

Table B.5: Employment and ethnic capital, by specific and multiple language proficiency

	All	No	Two	Main local language	
		English	languages	No	Yes
	(1)	(2)	(3)	(4)	(5)
Migrant own characteristics (dummy variables)					
No education or some primary	-0.031	-0.018	-0.031	-0.022	-0.036
	(0.005)***	(0.006)***	(0.008)***	(0.008)***	(0.007)***
Post-secondary	0.201	0.198	0.209	0.175	0.218
	(0.005)***	(0.007)***	(0.009)***	(0.006)***	(0.007)***
Person is male	0.167	0.206	0.209	0.179	0.157
	(0.006)***	(0.006)***	(0.007)***	(0.009)***	(0.007)***
Age 25-44	0.212	0.202	0.195	0.225	0.204
	(0.005)***	(0.006)***	(0.008)***	(0.007)***	(0.006)***
Age 45-65	0.285	0.282	0.267	0.308	0.271
	(0.007)***	(0.009)***	(0.010)***	(0.009)***	(0.011)***
Ethnic group average characteristics					
Ethnic group share with primary or no education	0.072	0.082	0.114	0.051	0.056
	(0.021)***	(0.025)***	(0.039)***	(0.022)**	(0.032)*
Ethnic group share with post-secondary education	0.136	0.146	0.101	0.073	0.113
	(0.024)***	(0.033)***	(0.049)**	(0.026)***	(0.044)**
Ethnic group share of male	0.069	0.089	0.090	0.006	0.072
	(0.018)***	(0.023)***	(0.032)***	(0.019)	(0.028)***
Ethnic group share 45-64	-0.020	-0.028	-0.022	0.005	-0.060
	(0.019)	(0.023)	(0.036)	(0.019)	(0.031)*
R^2	0.13	0.14	0.14	0.15	0.12
N	121,793	63,425	35,575	49,052	72,725
Significance of ethnic group characteristics	11.84	8.80	3.86	2.87	3.93
P-value of joint significance test	0.00	0.00	0.00	0.02	0.00
Ethnic group fixed effects	Yes	Yes	Yes	Yes	Yes
Local municipality fixed effects	Yes	Yes	Yes	Yes	Yes

Notes: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. The table presents the estimates from a linear probability model explaining the probability of being employed by individual characteristics and the characteristics of established migrants of the same ethnic group. The sample used is all African internal migrants aged 15 to 64 that moved to another province and to a non-neighboring municipality. The established migrants are all aged 25 or above. Column 1 presents the main results, as in Table 2.6. Column 2 restricts the sample on migrants who speak English as first or second language. Column 3 restricts the sample on individuals who report that they speak two migrants. Column 4 restricts the sample to migrants who do not speak the most spoken language in their municipality of residence, while column 5 restricts the sample to migrants who speak the language in their municipality of residence. Standard errors are clustered at the ethnic group x municipality level. The F-tests presented are for the joint significance of the characteristics of established migrants.

Table B.6: Employment and ethnic capital, excluding migrants moving to neighboring local municipalities or within same province

	(1)	(2)	(3)
Migrant own characteristics (dummy variables)			
No education or some primary	-0.009	-0.018	-0.026
	(0.010)	(0.008)**	(0.008)***
Post-secondary	0.198	0.196	0.188
	(0.007)***	(0.007)***	(0.006)***
Migrant is male	0.187	0.184	0.183
	(0.007)***	(0.008)***	(0.008)***
Age 25-44	0.209	0.209	0.210
	(0.007)***	(0.007)***	(0.006)***
Age 45-65	0.278	0.280	0.284
	(0.012)***	(0.011)***	(0.010)***
Ethnic group share (fraction)			
Ethnic group share with primary or no education		0.178	0.044
		(0.055)***	(0.023)*
Ethnic group share with post-secondary education		0.218	0.119
		(0.038)***	(0.025)***
Ethnic group share of male		0.059	0.036
		(0.034)*	(0.019)*
Ethnic group share 45-64		-0.064	-0.007
		(0.041)	(0.020)
R^2	0.11	0.11	0.14
N	74,812	73,092	73,089
Significance of ethnic group characteristics (F-stat)		9.28	6.68
P-value of joint significance test		0.00	0.00
Ethnic group fixed effects	No	No	Y
Local municipality fixed effects	No	No	Y

Notes: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. The table presents the estimates from a linear probability model explaining the probability of being employed by individual characteristics and the characteristics of established migrants of the same ethnic group. The sample used is all African internal migrants aged 15 to 64 that moved to another province and to a non-neighboring municipality. The established migrants are all aged 25 or above. Specification 1 contains only individual characteristics, whereas 2 and 3 introduce ethnic group characteristics and fixed effects for the ethnic group and the municipality. Standard errors are clustered at the ethnic group x municipality level. The F-tests presented are for the joint significance of the characteristics of established migrants.

Table B.7: Employment and ethnic capital: results for migrants aged 25-64

	(1)	(2)	(3)
Migrant own characteristics (dummy variables)			
No education or some primary	-0.026	-0.030	-0.035
	(0.006)***	(0.006)***	(0.006)***
Post-secondary	0.210	0.209	0.204
	(0.005)***	(0.005)***	(0.005)***
Age 45-65	0.072	0.073	0.074
	(0.006)***	(0.005)***	(0.005)***
Person is male	0.166	0.163	0.161
	(0.006)***	(0.006)***	(0.006)***
Ethnic group share (fraction)			
Ethnic group share with primary or no education		0.115	0.046
		(0.032)***	(0.020)**
Ethnic group share with post-secondary education		0.194	0.132
		(0.034)***	(0.022)***
Ethnic group share of male		0.081	0.063
		(0.027)***	(0.017)***
Ethnic group share 45-64		-0.026	-0.018
		(0.026)	(0.018)
R^2	0.07	0.07	0.09
N	94,966	93,247	93,244
Significance of ethnic group characteristics (F-stat)		10.75	12.66
P-value of joint significance test		0.00	0.00
Ethnic group fixed effects	No	No	Y
Local municipality fixed effects	No	No	Y

Notes: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. The table presents the estimates from a linear probability model explaining the probability of being employed by individual characteristics and the characteristics of established migrants of the same ethnic group characteristics. The sample used is all African internal migrants aged 25 to 64. The established migrants are all aged 25 or above. Specification 1 contains only individual characteristics, whereas 2 and 3 introduce ethnic group characteristics and fixed effects for the ethnic group and the municipality. The F-tests presented are for the joint significance of the characteristics of established migrants.

Table B.8: Employment and ethnic capital: results for the province of Gauteng

	(1)	(2)	(3)
Migrant own characteristics (dummy variables)			
No education or some primary	-0.025	-0.026	-0.027
	(0.010)**	(0.010)***	(0.010)***
Post-secondary	0.203	0.194	0.194
	(0.009)***	(0.008)***	(0.008)***
Person is male	0.153	0.156	0.157
	(0.009)***	(0.009)***	(0.009)***
Age 25-44	0.242	0.240	0.239
	(0.007)***	(0.007)***	(0.007)***
Age 45-65	0.348	0.340	0.338
	(0.009)***	(0.009)***	(0.009)***
Ethnic group share (fraction)			
Ethnic group share with primary or no education		0.141	0.078
		(0.052)***	(0.072)
Ethnic group share with post-secondary education		0.340	0.197
		(0.084)***	(0.082)**
Ethnic group share of male		-0.055	0.035
		(0.050)	(0.052)
Ethnic group share 45-64		0.193	0.063
		(0.057)***	(0.053)
R^2	0.12	0.12	0.13
N	51,957	51,867	51,867
Significance of ethnic group characteristics		10.57	1.87
P-value of joint significance test		0.00	0.11
Ethnic group fixed effects	No	No	Y
Local municipality fixed effects	No	No	Y

Notes: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. The table presents the estimates from a linear probability model explaining the probability of being employed by individual characteristics and the characteristics of established migrants of the same ethnic group characteristics. The sample used is all African internal migrants aged 15 to 64 that moved to Gauteng. Specification 1 contains only individual characteristics, whereas 2 and 3 introduce ethnic group characteristics and fixed effects for the ethnic group and the municipality. Standard errors are clustered at the ethnic group x municipality level. The F-tests presented are for the joint significance of the characteristics of established migrants.

Table B.9: Level of education and occupational choice

Grouped occupation	No school	Some primary	Completed	Some seco	Grade 12	Higher
Manager	2.23	5.53	2.35	22.38	37.4	29.45
Professional	1.41	3.69	1.48	14.09	24.11	54.5
Technician	2.03	4.57	2.02	20.77	37.88	31.86
Clerk	1.87	4.64	2.28	24.34	48.56	17.86
Sales and services	3.2	7.7	3.67	34.33	39.83	10.91
Skilled agriculture	8.11	12.93	5.25	31.51	28.94	12.96
Craft and related traders	4.6	11.06	5.5	39.58	29.58	9.34
Plant and machine operators	4.71	11.96	5.19	37.86	31.01	8.9
Elementary occupations	7.36	14.5	6.38	39.9	25.45	6.16
Domestic workers	7.81	16.82	6.97	40.3	22.27	5.56

Table B.10: First-stage of IV-estimates

Ethnic group share with (2011)	Primary	Post-secondary	Male	45-65
Ethnic group share with (in 2001)				
Primary or no education	0.178	-0.050	0.039	-0.003
	(0.002)***	(0.002)***	(0.003)***	(0.003)
Post-secondary education	0.005	0.168	0.018	0.087
	(0.004)	(0.003)***	(0.004)***	(0.004)***
Gender: male	0.045	0.027	0.202	-0.010
	(0.002)***	(0.002)***	(0.003)***	(0.002)***
Age 45-65	-0.020	0.013	-0.024	0.082
	(0.003)***	(0.002)***	(0.003)***	(0.003)***
R^2	0.66	0.47	0.54	0.36
N	99,410	99,410	99,410	99,410
F-statistic for excluded instruments	1658.98	1422.31	1774.41	364.65
Language X origin fixed effects	Y	Y	Y	Y
Residence fixed effects	Y	Y	Y	Y

Notes: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. The table presents the first stage results for the IV estimates presented in Table . The dependent variables are the 2011 average characteristics of established migrants of the same ethnic group at the place of residence. The instruments are the value of these averages for 2001.

Chapter C: Appendix for chapter 3

Table C.1: Sample distribution across RCT and LFEs

	# of Cooperatives		# of farmers in survey		# farmers in admin data	
Was in LFEs?	No	Yes	No	Yes	No	Yes
RCT Condition A	10	7	114	80	320	207
RCT Condition B	15	6	173	68	640	201
RCT Condition C	14	6	160	59	360	276
RCT Condition D	14	7	156	72	418	331
Total	53	26	603	279	1738	1015

Condition A is the control group in the RCT

Conditions B, C, and D are treatment variations (see Section 3.4.2 or Table C.3)

Table C.2: Distribution of Baseline and Communication Coordination Games

Variable	BCG	CCG	Total
# Sessions	28	28	56
# Rounds	112	110	222
# Players	429	410	839
# Observations	1716	1600	3316

Table C.3: RCT treatments

RCT Condition	What is revealed?		
A	–	–	–
B	Aggregate intentions	–	–
C	Aggregate intentions	Distribution	–
D	Aggregate intentions	Distribution	Leader vs. Member

In all conditions, there was a training; intentions were collected; and a survey was conducted.

Table C.4: Balance tests with respect to the binary treatment - farmer level

	N	All	No communication	Communication	P-value of diff.
Age	882	46.27	45.81	46.40	0.58
Sex (1=male; 0=female)	873	0.67	0.70	0.67	0.43
Leader (1=yes, 0=no)	873	0.18	0.22	0.18	0.19
Size of land (ha)	873	4.25	3.89	4.36	0.69
Risk (1 to 5)	873	2.80	2.81	2.80	0.96
Generosity (1 to 7)	873	2.92	3.19	2.84	0.00
Patience (1 to 5)	873	2.52	2.50	2.52	0.87
Federation (1=CCPA, 0=FEGPAB)	873	0.52	0.47	0.54	0.13
PO exposed to lablike exp.: 1=yes; 0=no	882	0.32	0.41	0.29	0.00
2013 harvest (kg)	873	1,680.85	1,868.76	1,627.16	0.32
Expected 2014 harvest (kg)	873	1,628.10	1,593.37	1,638.02	0.86
Intended to coll. com. : 1=yes, 0=no	873	0.84	0.80	0.85	0.10
Intentions coll. com. (kg)	873	966.43	830.39	1,005.30	0.25
Intentions indiv. com. (kg)	873	148.05	200.95	132.93	0.34
Intentions consumption (kg)	873	141.75	207.62	122.93	0.09
Intentions seeds stock (kg)	873	303.12	265.67	313.81	0.49
Farmed other crops : 1=yes, 0=no	873	0.57	0.59	0.57	0.55
Attended int. revelation meeting: 1=yes; 0=no	882	0.59	0.56	0.59	0.44

Notes: The p-values in the last column are reported based on two-tailed t-tests and the standard errors are clustered at the group level. All variables were measured before the intervention, that is before revealing intentions. Generosity/social preference was elicited through a hypothetical dictator game (with a greater number indicating a more generous/altruistic individual); risk through a hypothetical Binswanger (1980)-style lottery (with a lower number indicating a more risk averse individual); and patience through typical, hypothetical preference-over-time questions (with a one-day front-end delay and a higher number indicating a less patient individual).

Table C.5: Balance tests with respect to the treatment conditions - farmer level

	N	All	A	B	C	D	P-value of diff.
Age	882	46.27	45.81	46.90	47.98	44.53	0.04
Sex (1=male; 0=female)	873	0.67	0.70	0.64	0.71	0.65	0.35
Leader (1=yes, 0=no)	873	0.18	0.22	0.18	0.17	0.17	0.60
Size of land (ha)	873	4.25	3.89	5.53	3.67	3.76	0.45
Risk (1 to 5)	873	2.80	2.81	2.72	2.85	2.85	0.79
Generosity (1 to 7)	873	2.92	3.19	2.88	2.64	3.00	0.00
Patience (1 to 5)	873	2.52	2.50	2.54	2.70	2.32	0.17
Federation (1=CCPA, 0=FEGPAB)	873	0.52	0.47	0.53	0.54	0.53	0.50
Group exposed to lablike exp.: 1=yes; 0=no	882	0.32	0.41	0.28	0.27	0.33	0.01
2013 harvest (kg)	873	1,680.85	1,868.76	1,395.65	1,855.39	1,653.70	0.29
Expected 2014 harvest (kg)	873	1,628.10	1,593.37	1,748.09	1,670.35	1,484.56	0.82
Intended to coll. com. : 1=yes, 0=no	873	0.84	0.80	0.87	0.84	0.85	0.34
Intentions coll. com. (kg)	873	966.43	830.39	966.20	1,068.47	985.15	0.64
Intentions indiv. com. (kg)	873	148.05	200.95	205.02	123.58	62.95	0.26
Farmed other crops : 1=yes, 0=no	873	0.57	0.59	0.61	0.57	0.53	0.29
Attended int. revelation meeting: 1=yes; 0=no	882	0.59	0.56	0.56	0.60	0.61	0.70

Notes: The p-values in last column are obtained by running a one-way ANOVA test using STATA, with standard errors clustered at the group level.

Table C.6: Extensive margin: treatment effect on the probability of collective action

Dependent variable Source	Contributed to group LFE		Intended to contribute to group RCT - survey data		Contributed to group RCT - survey data		Contributed to group RCT - admin. data	
	1	2	3	4	5	6	7	8
Treatment	0.03	-0.04	0.06	0.05	0.07	-0.08	0.15	-0.21
	(0.03)	(0.04)	(0.03)	(0.06)	(0.05)	(0.09)	(0.10)	(0.21)
Size of Group	0.03	-0.02	0.00	0.00	0.01	0.00	-0.00	-0.01
	(0.03)	(0.04)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Treatment X Size		0.11		0.00		0.01		0.01
		(0.05)		(0.00)		(0.00)		(0.01)
Control group mean	0.89	0.89	0.80	0.80	0.13	0.13	0.27	0.27
R^2	0.05	0.06	0.10	0.10	0.16	0.16	0.08	0.09
N	3,316	3,316	873	873	873	873	2,752	2,752

Notes: Standard errors are clustered at the group level. Columns 1 and 2 use data from the LFEs. The dependent variable is the number of chips. The size in those specifications is a dummy variable that is 1 if the experimental session was run with 20 farmers and 0 if it was run with 10 farmers. Columns 3 and 4 use farmers' intended sales (of groundnuts) to the group, which were collected prior to the RCT interventions. Columns 5 and 6 use the self-reported (survey) collective commercialization. Columns 7 and 8 use administrative data on collective commercialization obtained from the cooperatives (recall booklets). In columns 3–8, size is the actual group size ranging from 4-91 (with median size being 24). Controls in the LFE regressions include age, sex, land size, education (dummy for going to a French school versus a Koranic or no schooling), generosity, risk aversion measures and time preferences elicited through hypothetical questions (see Table C.4), and dummies for the federation and whether the farmer is a leader. Controls in the RCT regressions with survey data include the same variables (except for education) and a dummy for whether the farmer produced crops other than groundnuts, all measured pre-intervention.

Table C.7: PSP estimates: LFE cheap-talk effect in Table 3.1 - column 1 ($\beta = 0.66$)

π	$k=1$	$k=2$	$k=5$	$k=10$	$k=15$
0.01	0.118	0.084	0.043	0.025	0.018
0.02	0.212	0.156	0.082	0.048	0.037
0.05	0.410	0.323	0.188	0.116	0.089
0.1	**0.595**	**0.502**	0.328	0.217	0.172
0.2	**0.767**	**0.694**	**0.524**	0.384	0.318
0.3	**0.850**	**0.795**	**0.653**	**0.516**	0.444
0.35	**0.877**	**0.830**	**0.703**	**0.573**	**0.501**
0.55	**0.942**	**0.917**	**0.843**	**0.753**	**0.695**

Note: $PSP = \frac{(1-\beta^k)\pi}{(1-\beta^k)\pi+[1-(1-\alpha)^k](1-\pi)}$, where:
typical power $= 1-\beta$, $\alpha = 0.05$, $\pi =$ a prior
k represents the number of independent researchers working on specific associations in a field (i.e., competition)

Table C.8: PSP estimates: RCT cheap-talk effect in Table 3.1 - column 5 ($\beta = 0.99$)

π	$k=1$	$k=2$	$k=5$	$k=10$	$k=15$
0.01	0.167	0.094	0.043	0.025	0.018
0.02	0.288	0.173	0.083	0.048	0.037
0.05	**0.510**	0.351	0.189	0.116	0.089
0.1	**0.688**	**0.533**	0.329	0.217	0.172
0.2	**0.832**	**0.719**	**0.525**	0.384	0.318
0.3	**0.895**	**0.815**	**0.655**	**0.516**	0.444
0.35	**0.914**	**0.847**	**0.704**	**0.573**	**0.501**
0.55	**0.960**	**0.926**	**0.844**	**0.753**	**0.695**

Note: See Table C.7 for definition of PSP.

Table C.9: Effect on collective commercialization - Tobit model

Dependent variable Source	Contributed to group RCT - survey data			
Treatment	278.60	-622.40	460.23	-88.53
	(295.31)	(550.79)	(258.33)	(512.10)
Size of Group	27.50	-2.87	25.79	7.97
	(6.20)	(15.59)	(6.50)	(15.38)
Treatment X Size		31.70		18.94
		(16.94)		(16.57)
Control group mean	39.58			
Pseudo R^2	0.03	0.03	0.04	0.04
N	882	873	873	873
Control	No	No	Yes	Yes

Notes: The table shows the estimations of the main equation of interest (collective commercialization on treatment) using a Tobit model for the null quantities reported. The dependent variable is the quantity of groundnuts sold through the group. Controls include age, sex, land size, generosity, risk aversion measures and time preferences elicited through hypothetical questions (see Table C.4), and dummies for the federation, whether the farmer is a leader, and whether the farmer produced crops other than groundnuts, all measured pre-intervention. A dummy for whether the individual indicated positive intentions to sell through the group is also added.

Table C.10: Effect on collective commercialization - inverse hyperbolic sine transformation

Dependent variable	Contributed to group			
Source	RCT - survey data			
Treatment	0.42	-1.13	0.64	-0.52
	(0.39)	(0.68)	(0.34)	(0.63)
Size of Group	0.05	-0.00	0.05	0.01
	(0.01)	(0.02)	(0.01)	(0.02)
Treatment X Size		0.06		0.04
		(0.02)		(0.02)
Control group mean	39.58			
R^2	0.10	0.11	0.17	0.17
N	882	873	873	873
Control	No	No	Yes	Yes

Notes: Standard errors are clustered at the group level. The dependent variable is the quantity of groundnuts sold through the group, transformed by the inverse hyperbolic sine. Controls include age, sex, land size, generosity, risk aversion measures and time preferences elicited through hypothetical questions (see Table C.4), and dummies for the federation, whether the farmer is a leader, and whether the farmer produced crops other than groundnuts, all measured pre-intervention. A dummy for whether the individual indicated positive intentions to sell through the group is also added.

Table C.11: Intensive margin results for LFEs: testing sensitivity to farmer random effects

	(1)	(2)	(3)	(4)
Treatment	0.40	-0.03	0.40	-0.03
	(0.19)	(0.31)	(0.19)	(0.31)
Size of PO	0.03	-0.31	0.03	-0.30
	(0.24)	(0.27)	(0.24)	(0.27)
Treatment X Size		0.65		0.64
		(0.38)		(0.38)
R^2	0.09	0.10		
N	3,316	3,316	3,316	3,316
Farmer random effects	No	No	Yes	Yes

Notes: Standard errors are clustered at the group level. The table reproduces the results in Table 3.1 and explores the sensitivity to the inclusion of farmer random effects. All estimations use data from the LFEs. The dependent variable is the number of chips. The size in the specifications is a dummy variable that is 1 if the experimental session was run with 20 farmers and 0 if it was run with 10 farmers. Columns 3 and 4 include a farmer random effect. Controls in the LFE regressions include age, sex, land size, education (dummy for going to a French school versus a Koranic or no schooling), generosity, risk aversion measures and time preferences elicited through hypothetical questions (see Table C.4), and dummies for the federation and whether the farmer is a leader.

Table C.12: Isolating the effect of group size from that of the threshold (LFEs)

Dependent variable Source (recall Section 3.4.1)	# Chips contributed LFE					
Treatment	0.40	-0.03	0.22	-0.03	0.36	-0.02
	(0.19)	(0.31)	(0.21)	(0.30)	(0.24)	(0.29)
Size of Group	0.03	-0.31	-0.27	-0.53	-0.00	-0.38
	(0.24)	(0.27)	(0.21)	(0.35)	(0.23)	(0.25)
Treatment X Size		0.65		0.53		0.77
		(0.38)		(0.41)		(0.44)
Control group mean	3.03		3.09		3.09	
R^2	0.09	0.10	0.08	0.09	0.11	0.12
N	3,316	3,316	2,116	2,116	2,320	2,320
Fixed threshold?	No		Yes - Average		Yes - Overall	

Notes: Standard errors are clustered at the group level. All columns use data from the LFEs. The dependent variable is the number of chips. The size in all specifications is a dummy variable that is 1 if the experimental session was run with 20 farmers and 0 if it was run with 10 farmers. Columns 1 and 2 use data from all LFEs regardless of the threshold required for coordination. Columns 3 and 4 only use data from the LFEs where the average (per-player) number of chips required to reach the threshold was the same across large and small groups. Columns 5 and 6 only use data from the LFEs where the overall threshold was fixed across small and large groups. Controls in the regressions include age, sex, land size, education (dummy for going to a French school versus a Koranic or no schooling), generosity, risk aversion measures and time preferences elicited through hypothetical questions (see Table C.4), and dummies for the federation and whether the farmer is a leader.

Table C.13: Robustness check: effect of the treatment on patience, generosity and risk

Dependent variable	Patience		Generosity		Risk	
Source			RCT - survey			
Binary treatment						
Treatment	-0.15	-0.14	0.08	0.09	0.17	0.40
	(0.11)	(0.28)	(0.10)	(0.20)	(0.11)	(0.33)
Size of Group	-0.00	-0.00	0.00	0.00	-0.00	0.00
	(0.00)	(0.01)	(0.00)	(0.01)	(0.00)	(0.01)
Treatment X Size		-0.00		-0.00		-0.01
		(0.01)		(0.01)		(0.01)
R^2	0.07	0.07	0.08	0.08	0.04	0.04
N	873	873	873	873	873	873
Treatment arms						
B	-0.12	-0.04	0.11	0.16	0.18	0.44
	(0.13)	(0.32)	(0.12)	(0.26)	(0.13)	(0.35)
C	-0.29	-0.22	0.04	-0.03	0.08	0.30
	(0.14)	(0.32)	(0.12)	(0.25)	(0.16)	(0.37)
D	-0.04	-0.15	0.08	0.16	0.23	0.48
	(0.14)	(0.32)	(0.13)	(0.24)	(0.14)	(0.38)
Size of Group	-0.00	-0.00	0.00	0.00	-0.01	0.00
	(0.00)	(0.01)	(0.00)	(0.01)	(0.00)	(0.01)
B X Size		-0.00		-0.00		-0.01
		(0.01)		(0.01)		(0.01)
C X Size		-0.00		0.00		-0.01
		(0.01)		(0.01)		(0.01)
D X Size		0.00		-0.00		-0.01
		(0.01)		(0.01)		(0.01)
Control group mean	1.84	1.84	3.10	3.10	2.15	2.15
R^2	0.07	0.07	0.08	0.08	0.04	0.04
N	873	873	873	873	873	873

Notes: Standard errors are clustered at the group level. See Table C.4 for definitions of generosity/social preferences, risk preferences, and patience/time preferences. All dependent variables were measured after the intervention and in the same way they were measured before the intervention. Their values pre-intervention are used as controls in the regressions. Additional controls include age, sex, land size, and dummies for the federation, whether the farmer is a leader, and whether the farmer produced crops other than groundnuts, all measured pre-intervention.

Table C.14: Comparing large and small cooperatives

	N	All	Small Cooperatives	Large Cooperatives	P-value of diff.
Age	882	46.27	46.59	45.99	0.60
Sex (1=male; 0=female)	873	0.67	0.71	0.64	0.37
Leader (1=yes, 0=no)	873	0.18	0.18	0.19	0.83
Size of land (ha)	873	4.25	4.13	4.36	0.81
Risk (1 to 5)	873	2.80	2.74	2.86	0.37
Generosity (1 to 7)	873	2.92	2.89	2.94	0.77
Patience (1 to 5)	873	2.52	2.56	2.48	0.61
Federation (1=CCPA, 0=FEGPAB)	873	0.52	0.33	0.69	0.00
Coop exposed to lablike exp.: 1=yes; 0=no	882	0.32	0.27	0.36	0.39
2013 harvest (kg)	873	1,680.85	1,780.21	1,592.46	0.54
Expected 2014 harvest (kg)	873	1,628.10	1,663.21	1,596.86	0.87
Intended to coll. com. : 1=yes, 0=no	873	0.84	0.81	0.87	0.08
Intentions coll. com. (kg)	873	966.43	972.34	961.17	0.96
Intentions indiv. com. (kg)	873	148.05	170.02	128.50	0.58
Farmed other crops : 1=yes, 0=no	873	0.57	0.59	0.56	0.55
Attended int. revelation meeting: 1=yes; 0=no	882	0.59	0.62	0.55	0.11

Notes: The p-values in the last column are reported from two-tailed t-tests and the standard errors are clustered at the group level. Cooperatives are divided at the median size (24 members) into large and small. All the variables were measured before the intervention, that is before revealing intentions. See Table C.4 for definitions of generosity/social preferences, risk preferences, and patience/time preferences.

Table C.15: Collective commercialization and aggregate intentions

Dependent variable Source	Contributed to group RCT - survey data			
Treatment	91.90	-16.87	121.66	50.55
	(37.79)	(36.80)	(34.92)	(35.13)
Aggregate intentions (tonnes)	3.41	-0.49	0.58	-1.83
	(1.08)	(0.51)	(1.46)	(1.29)
Treatment X Aggregate intentions		4.59		2.96
		(1.42)		(1.44)
Control group mean	39.58			
R^2	0.08	0.10	0.17	0.18
N	873	873	873	873
Controls?	No	No	Yes	Yes

Notes: Standard errors are clustered at the group level. The dependent variable is the quantity of groundnuts sold through the group. Aggregate intentions (in tonnes) are obtained as the sum of the individual intentions reported by farmers. Controls include age, sex, land size, generosity, risk aversion measures and time preferences elicited through hypothetical questions (see Table C.4), cooperative size, and dummies for the federation, whether the farmer is a leader, and whether the farmer produced crops other than groundnuts, all measured pre-intervention. A dummy for whether the individual indicated positive intentions to sell through the group is also added.

Table C.16: Collective commercialization and leader's intentions

Dependent variable	Contributed to group			
Source	RCT - survey data			
B	73.39	108.56	57.59	88.06
	(48.97)	(43.60)	(44.72)	(40.47)
C	133.66	169.72	148.73	181.64
	(56.57)	(54.83)	(60.85)	(58.72)
D	125.63	144.68	70.84	117.12
	(63.79)	(59.99)	(43.04)	(45.77)
Average intentions leader	0.06	0.07		
	(0.04)	(0.04)		
Average intentions member	-0.04	-0.06		
	(0.03)	(0.03)		
D X Average intentions leader	0.06	0.03		
	(0.07)	(0.08)		
D X Average intentions member	-0.10	-0.06		
	(0.07)	(0.08)		
Leader's intentions -/- member's ≥ 0			46.66	49.67
			(49.42)	(45.38)
D X difference ≥ 0			43.34	6.62
			(72.45)	(71.19)
Size of Group	7.14	6.97	7.52	7.19
	(2.03)	(2.18)	(1.96)	(1.92)
Control group mean	39.58			
R^2	0.15	0.19	0.14	0.18
N	868	868	868	868
Control	No	Yes	No	Yes

Notes: Standard errors are clustered at the group level. The dependent variable is the quantity of groundnuts sold through the group. Average intentions are used as an explanatory variable. This information was only revealed in the RCT Group D. Controls include age, sex, land size, generosity, risk aversion measures and time preferences elicited through hypothetical questions (see Table C.4), and dummies for the federation, whether the farmer is a leader, and whether the farmer produced crops other than groundnuts, all measured pre-intervention. A dummy for whether the individual indicated positive intentions to sell through the group is also added.

Table C.17: Comparing farmers attending games to farmers not attending, within LFE cooperatives

			Farmers in games		
	N	All	No	Yes	P-value of diff.
Age	279	45.05	46.85	44.44	0.07
Sex (1=male; 0=female)	279	0.65	0.61	0.66	0.45
Leader (1=yes, 0=no)	279	0.19	0.10	0.23	0.00
Size of land (ha)	279	3.96	3.36	4.17	0.26
Risk (1 to 5)	279	2.96	2.79	3.02	0.19
Generosity (1 to 7)	279	3.16	3.01	3.21	0.30
Patience (1 to 5)	279	2.44	2.46	2.43	0.87
Federation (1=CCPA, 0=FEGPAB)	279	0.40	0.52	0.36	0.14
2013 harvest (kg)	279	1,807.42	1,537.59	1,899.52	0.40
Expected 2014 harvest (kg)	279	1,432.58	1,585.80	1,380.28	0.65
Intended to coll. com. : 1=yes, 0=no	279	0.86	0.86	0.87	0.91
Intentions coll. com. (kg)	279	916.64	1,052.35	870.32	0.61
Intentions indiv. com. (kg)	279	71.65	54.79	77.40	0.56
Farmed other crops : 1=yes, 0=no	279	0.66	0.68	0.66	0.80
Attended int. revelation meeting: 1=yes; 0=no	279	0.72	0.59	0.77	0.00

Notes: The p-values in the last column are reported from two-tailed t-tests, and the standard errors are clustered at group level. Only cooperatives that took part in the LFEs are considered here. All members did not take part, and those who took part were chosen randomly. The table compares those two categories of farmers, within cooperatives exposed to LFEs. All the variables were measured before the intervention, except for the attendance at the intention revelation meeting, that is before revealing intentions. See Table C.4 for definitions of generosity/social preferences, risk preferences, and patience/time preferences.

Table C.18: Comparing farmers in LFE cooperatives to farmers in non-LFE cooperatives

		LFE cooperatives				
	N	All	No	Yes	P-value of diff.	Omnibus test
	(1)	(2)	(3)	(4)	(5)	(6)
Age	882	46.27	46.84	45.05	0.16	-0.00
						(0.00)
Sex (1=male; 0=female)	873	0.67	0.68	0.65	0.63	-0.03
						(0.07)
Leader (1=yes, 0=no)	873	0.18	0.18	0.19	0.52	-0.04
						(0.03)
Size of land (ha)	873	4.25	4.39	3.96	0.60	0.00
						(0.00)
Risk (1 to 5)	873	2.80	2.73	2.96	0.09	0.02
						(0.01)
Generosity (1 to 7)	873	2.92	2.80	3.16	0.04	0.04
						(0.02)
Patience (1 to 5) [pre-intervention]	873	2.52	2.56	2.44	0.45	-0.02
						(0.01)
Federation (1=CCPA, 0=FEGPAB)	873	0.52	0.58	0.40	0.13	-0.14
						(0.10)
2013 harvest (kg)	873	1,680.85	1,621.40	1,807.42	0.58	0.02
						(0.01)
Expected 2014 harvest (kg)	873	1,628.10	1,719.93	1,432.58	0.45	-0.02
						(0.01)
Intended to coll. com. : 1=yes, 0=no	873	0.84	0.83	0.86	0.33	0.06
						(0.04)
Intentions coll. com. (kg)	873	966.43	989.81	916.64	0.78	0.01
						(0.02)
Intentions indiv. com. (kg)	873	148.05	183.93	71.65	0.05	-0.02
						(0.02)
Farmed other crops : 1=yes, 0=no	873	0.57	0.53	0.66	0.02	0.10
						(0.05)
Attended int. revelation meeting: 1=yes; 0=no	882	0.59	0.52	0.72	0.00	0.17
						(0.04)
Omnibus test F-stat						4.04
Omnibus test P-value						0.00

Notes: The p-values in the last column are reported from two-tailed t-tests, and the standard errors for the tests are clustered at group level. Column 6 presents the results of an omnibus test, which regresses the dummy for being in a games cooperative against all the variables on which balance is tested. The table here compares farmers in the LFE cooperative sample to those who later took part only in the RCT. All the variables, except for the attendance at the intention revelation meeting, were measured before the intervention, i.e., before revealing intentions. See Table C.4 for definitions of generosity/social preferences, risk preferences, and patience/time preferences.

Bibliography

Abel, M. (2019). Unintended Labor Supply Effects of Cash Transfer Programs: New Evidence from South Africa's Pension. *Journal of African Economies* (April).

Abel, M., R. Burger, and P. Piraino (2019). The Value of Reference Letters: Experimental Evidence from South Africa. *American Economic Journal: Applied Economics forthcoming*.

Andersson, F., M. Garca-Prez, J. Haltiwanger, K. McCue, and S. Sanders (2014). Workplace concentration of immigrants. *Demography 51*, 2281–2306.

Ardington, C., A. Case, and V. Hosegood (2009). Labor Supply Responses to Large Social Transfers: Longitudinal Evidence from South Africa. *American Economic Journal: Applied Economics 1*(1), 22–48.

Ariza, J. and P. Cedano (2017). Empleo y Formalizacin Laboral Juvenil en Colombia. Una Evaluacin de la Ley del Primer Empleo. *Revista de Economia Laboral 14*(1).

Ashraf, N., X. Giné, and D. Karlan (2009). Finding missing markets (and a disturbing epilogue): Evidence from an export crop adoption and marketing intervention in Kenya. *American Journal of Agricultural Economics 91*(4), 973–990.

Avoyan, A. and J. Ramos (2017). A road to efficiency through communication and commitment. Working paper.

Balkan, B., Y. Soner, and K. Semih (2016). Evaluating the Impact of the Post-2008 Employment Subsidy Program in Turkey. *IZA Discussion Paper No. 9993*(June).

Balliet, D. (2010). Communication and cooperation in social dilemmas: A meta-analytic review. *Journal of Conflict Resolution 54*(1), 39–57.

Banerjee, A., R. Banerji, J. Berry, E. Duflo, H. Kannan, S. Mukerji, M. Shotland, and M. Walton (2017, November). From proof of concept to scalable policies: Challenges and solutions, with an application. *Journal of Economic Perspectives 31*(4), 73–102.

Banerjee, A., S. Galiani, J. Levinsohn, Z. McLaren, and I. Woolard (2008). Why has unemployment risen in the New South Africa? *Economics of Transition 16*(4), 715–740.

Barr, A., A. Zeitlin, et al. (2010). Dictator games in the lab and in nature: External validity tested and investigated in Ugandan primary schools. Technical report, Centre for the Study of African Economies, University of Oxford.

Barrett, C. B. (2008). Smallholder market participation: Concepts and evidence from eastern and southern Africa. *Food Policy 33*(4), 299 – 317.

Battalio, R., L. Samuelson, and J. Van Huyck (2001). Optimization incentives and coordination failure in laboratory stag hunt games. *Econometrica 69*(3), 749–764.

Bayer, P., S. L. Ross, and G. Topa (2008). Place of work and place of residence: Informal hiring networks and labor market outcomes. *Journal of Political Economy 116*(6), 1150–1196.

Beaman, L. and J. Magruder (2012). Who Gets the Job Referral? Evidence from a Social Networks Experiment. *American Economic Review 102*(7), 3574–93.

Beaman, L. A. (2011, 08). Social Networks and the Dynamics of Labour Market Outcomes: Evidence from Refugees Resettled in the U.S. *The Review of Economic Studies 79*(1), 128–161.

Bell, B., R. Blundell, and J. V. Reenen (1999). Getting the unemployed back to work: The role of targeted wage subsidies. *International Tax and Public Finance 6*(3), 339–360.

Benjamin, P., H. Bhorat, and H. Cheadle (2010). The cost of "doing busines" and labour regulation: The case of South Africa. *Inernational Labour Review 149*, 73–91.

Bergquist, L. F. (2017). Pass-through, competition, and entry in agricultural markets: Experimental evidence from Kenya. *UC Berkeley Mimeograph*.

Bernard, T., M.-H. Collion, A. de Janvry, P. Rondot, and E. Sadoulet (2008). Do village organizations make a difference in african rural development? a study for senegal and burkina faso. *World Development 36*(11), 2188 – 2204.

Bernard, T., L. Sène, A. C. G. Viceisza, and F. Wouterse (2014). Leaders needed: Experimental evidence from rural producer organizations in Senegal. Working paper.

Bernard, T., A. S. Taffesse, and E. Gabre-Madhin (2008). Impact of cooperatives on smallholders' commercialization behavior: evidence from Ethiopia. *Agricultural Economics 39*(2), 147–161.

Bernheim, B. D. (1994). A theory of conformity. *Journal of Political Economy 102*(5), 841–877.

Bertrand, M. and B. Crepon (2019). Teaching Labor Laws: Evidence From a Randomized Control Trial in South Africa. *Working Paper*.

Bertrand, M., S. Mullainathan, and D. Miller (2003). Public Policy and Extended Families: Evidence from Pensions in South Africa. *The World Bank Economic Review 17*(1), 27–50.

Betcherman, G., N. M. Daysal, and C. Pagés (2010). Do employment subsidies work? Evidence from regionally targeted subsidies in Turkey. *Labour Economics 17*(4), 710–722.

Betcherman, G., K. Olivas, and A. Dar (2004). Impacts of Active Labour Market Programs: New evidence form evaluations with particular attention to developing and transition countries. *World Bank Social Protection Discussion Paper Series* (No. 0402), 1 – 94.

Bhorat, H., R. Kanbur, and B. Stanwix (2014). Estimating the Impact of Minimum Wages on Employment, Wages, and Non-Wage Benefits: The Case of Agriculture in South Africa. *American Journal of Agricultural Economics 96*(5), 1402–1419.

Binswanger, H. P. (1980). Attitudes toward risk: Experimental measurement in rural india. *American Journal of Agricultural Economics 62*(3), 395.

Blattman, C. and L. Ralston (2015). Generating employment in poor and fragile states: Evidence from labor market and entrepreneurship programs. *Working Paper*.

Blume, A. and A. Ortmann (2007). The effects of costless pre-play communication: Experimental evidence from games with pareto-ranked equilibria. *Journal of Economic Theory 132*(1), 274 – 290.

Borjas, G. J. (1992). Ethnic Capital and Intergenerational Mobility. *The Quarterly Journal of Economics 107*(1), 123–150.

Borjas, G. J. (1999). *Heaven's Door - Immigration Policy and the American Economy*. Princeton University Press.

Brandts, J. and D. J. Cooper (2006, June). A change would do you good .... an experimental study on how to overcome coordination failure in organizations. *American Economic Review 96*(3), 669–693.

Bravo, D. and T. Rau (2013). Effects of Large-scale Youth Employment Subsidies: Evidence from a Regression Discontinuity Design. *Working Paper*.

Broecke, S. (2013). Tackling graduate unemployment in North Africa through employment subsidies: A look at the SIVP programme in Tunisia. *IZA Journal of Labor Policy 2*(1), 1–19.

Bruhn, M. (2016). Can wage subsidies boost employment in the wake of an economic crisis? evidence from mexico. *World Bank Policy Research Working Paper 7607*, 33.

Burns, J., S. Godlonton, and K. Malcolm (2010). Social networks, employment and worker discouragement: Evidence from south africa. *Labor Economics 68*(2), 235260.

Calvo-Armengol, A. and M. O. Jackson (2004, June). The effects of social networks on employment and inequality. *American Economic Review 94*(3), 426–454.

Camerer, C. (2003). *Behavioral Game Theory*. Princeton, NJ: Princeton University Press.

Camerer, C. F. (2015). The promise and success of lab-field generalizability in experimental economics: A critical reply to Levitt and List. In G. Frechette and A. Schotter (Eds.), *Handbook of Experimental Economic Methodology*. Oxford University Press.

Card, D., J. Kluve, and A. Weber (2015). What Works? A Meta Analysis of Recent Active Labor Market Program Evaluations. *NBER Working Paper 21431*.

Card, D. and A. Krueger (1994). Minimum wages and employment: A case study of the fast-food industry in new jersey and pennsylvania. *American Economic Review 84*(4), 772–93.

Cardenas, J. C. and J. P. Carpenter (2005). *Three Themes on Field Experiments and Economic Development*, pp. 71–123.

Carranza, E., R. Garlick, K. Orkin, and N. Rankin (2019). Job Search and Hiring with Two-sided Limited Information about Workseekers' Skills. *Working Paper*.

Casaburi, L. and R. Macchiavello (2019, February). Demand and supply of infrequent payments as a commitment device: Evidence from Kenya. *American Economic Review 109*(2), 523–55.

Casaburi, L. and T. Reed (2017). Competition in agricultural markets: An experimental approach. Working paper.

Caswell, N. (1984). Autopsie de l'ONCAD. *Politique Africaine 14*(June), 39–73.

Charness, G. and B. Grosskopf (2004). What makes cheap talk effective? experimental evidence. *Economics Letters 83*, 383–389.

Chauang, Y. and L. Shecther (2015). Social networks in developing countries. *Annual Review of Resource Economics*, 51–72.

Chiswick, B. R. and P. W. Miller (2005). Do enclaves matter in immigrant adjustment. *City and Community*, 5–35.

Cooper, D. J. and J. Van Huyck (2016). Coordination and transfer. Working paper.

Cooper, D. J. and R. A. Weber (2019). Recent advances in experimental coordination games. Working paper.

Cooper, R., D. V. DeJong, R. Forsythe, and T. W. Ross (1992). Communication in coordination games. *Quarterly Journal of Economics 107*(2), 731–771.

Cramer, C. (2010). Unemployment and participation in violence. *World Development Report 2011 Background Paper*.

Crawford, V. P. (1998). A survey of experiments on communication via cheap talk. *Journal of Economic Theory 78*, 286–298.

Damm, A. P. (2009). Ethnic enclaves and immigrant labor market outcomes: Quasiexperimental evidence. *Journal of Labor Economics 27*(2), 281–314.

de Mel, S., D. McKenzie, and C. Woodruff (2019). Labor Drops: Experimental Evidence on the Return to Additional Labor in Microenterprises. *American Economic Journal: Applied Economics 11*(1), 202–35.

De Weerdt, J. and S. Dercon (2006). Risk-sharing networks and insurance against illness. *Journal of Development Economics 81*(2), 337 – 356.

Deumert, A., B. Inder, and P. Maitra (2005). Language, Informal Networks and Social Protection: Evidence from a Sample of Migrants in Cape Town, South Africa. *Global Social Policy 5*(3), 303–328.

Deumert, A. and N. Mabandla (2009). I-Dollar EYI One!Ethnolinguistic Fractionalisation, Communication Networks and Economic ParticipationLessons from Cape Town, South Africa. *The Journal of Development Studies 45*(3), 412–440.

Devetag, G. and A. Ortmann (2007). When and why? a critical survey on coordination failure in the laboratory. *Experimental Economics 10*(3), 331 – 344.

Dillon, B. and C. Dambro (2017). How competitive are crop markets in Sub-Saharan Africa? *American Journal of Agricultural Economics 99*(5), 1344–1361.

Dinkelman, T. and V. Ranchhod (2012). Evidence on the impact of minimum wage laws in an informal sector: Domestic workers in South Africa. *Journal of Development Economics 99*(1), 27–45.

Ebrahim, A., M. Leibbrandt, and V. Ranchhod (2017). The effects of the Employment Tax Incentive on South African employment. *Wider working paper* (January), 24.

Edin, P.-A., P. Fredriksson, and O. slund (2003, 02). Ethnic enclaves and the economic success of immigrantsevidence from a natural experiment. *The Quarterly Journal of Economics 118*(1), 329–357.

Fafchamps, M. and F. Gubert (2007, May). Risk sharing and network formation. *American Economic Review 97*(2), 75–79.

Fafchamps, M. and R. V. Hill (2005). Selling at the farmgate or traveling to market. *American Journal of Agricultural Economics 87*(3), 717–734.

Fafchamps, M. and S. Lund (2003). Risk-sharing networks in rural Philippines. *Journal of Development Economics 71*(2), 261 – 287.

Fafchamps, M. and A. Moradi (2015). Referral and job performance: Evidence from the ghana colonial army. *Economic Development and Cultural Change 63*(4), 715–751.

Farrell, J. (1987). Cheap talk, coordination, and entry. *Rand Journal of Economics 18*, 34–39.

Feltovich, N. and P. J. Grossman (2015). How does the effect of pre-play suggestions vary with group size? Experimental evidence from a threshold public-good game. *European Economic Review 79*(C), 263–280.

Filmer, D. and L. Fox (2014). Youth employment in sub-saharan africa. *Africa Development Series, World Bank*.

Finan, F. and L. Schechter (2012). Vote-buying and reciprocity. *Econometrica 80*(2), 863–881.

Gächter, S., D. Nosenzo, E. Renner, and M. Sefton (2010). Sequential vs. simultaneous contributions to public goods: Experimental evidence. *Journal of Public Economics 94*(7), 515–522.

Galasso, E., M. Ravallion, and A. Salvia (2004). Assisting the Transition from Workfare to Work : A Randomized Experiment . *Industrial and Labor Relations Review 58*(1), 128–142.

Garlick, J., M. Leibbrandt, and J. Levinsohn (2016). Individual migration and households incomes. *NBER Working Paper N 22326*.

Gelderblom, D. and W. Adams (2006). Migration in South and Southern Africa: dynamics and determinant, pp. 227–248. Cape Town: HSRC Publishers.

Goeree, J. and C. Holt (2002). Coordination games. Volume 2 of *Encyclopedia of Cognitive Science*, pp. 204?208. London: Nature Publishing Group, McMillan.

Gorry, A. (2013). Minimum wages and youth unemployment. *European Economic Review 64*, 57 – 75.

Granovetter, M. (2005, March). The impact of social structure on economic outcomes. *Journal of Economic Perspectives 19*(1), 33–50.

Groh, M., N. Krishnan, D. McKenzie, and T. Vishwanath (2016). Do Wage Subsidies Provide A Stepping-Stone To Employment For Recent College Graduates? Evidence From A Randomized Experiment In Jordan. *Review of Economics and Statistics 98*(3), 488–502.

Guerra, J. A. and M. Mohnen (2016). Occupational choice and social interactions: A study of victorian london. *Working Paper*.

Hardy, M. and J. McCasland (2017). Are Small Firms Labor Constrained? Experimental Evidence from Ghana.

Harrison, G. W. and J. A. List (2004). Field experiments. *Journal of Economic Literature 42*(4), 1009–1055.

Heath, R. (2018). Why do firms hire using referrals? evidence from bangladeshi garment factories. *Journal of Political Economy 126*(4), 1691–1746.

Heckman, J. J. and T. E. Macurdy (1985). A simultaneous equations linear probability model. *The Canadian Journal of Economics / Revue canadienne d'Economique 18*, 28–37.

Heinemann, H., R. Nagel, and P. Ockenfels (2009). Measuring strategic uncertainty in coordination games. *The Review of Economic Studies 76*(1), 181–221.

Hermalin, B. E. (1998). Toward an economic theory of leadership: Leading by example. *American Economic Review*, 1188–1206.

Hill, R., E. Maruyama, M. Olapade, and M. Frölich (2014). Strengthening producer organizations to increase market access of smallholder farmers in Uganda. Working paper.

Hill, R. V., E. Maruyama, and A. Viceisza (2012). Breaking the norm: An empirical investigation into the unraveling of good behavior. *Journal of Development Economics 99*(1), 150–162.

Hoel, J. B., M. Hidrobo, T. Bernard, and M. Ashour (2017). Productive inefficiency in dairy farming and cooperation between spouses: Evidence from Senegal. IFPRI discussion paper 1698, Washington, DC: IFPRI.

Hoff, K. (2001). Beyond Rosenstein-Rodan: The modern theory of coordination problems in development. In *Proceedings of the Annual World Bank Conference on Development Economics, 2000*, pp. 145–188. Washington DC: World Bank.

Hofmeyr, A. (2010). Social networks and ethnic niches: An econometric analysis of the manufacturing sector in south africa. *South African Journal of Economics 78*(1), 107–130.

ILO (2018). Women and Men in the Informal Economy: A Statistical Picture - Third Edition .

Ioannides, Y. M. and L. Datcher Loury (2004, December). Job information networks, neighborhood effects, and inequality. *Journal of Economic Literature 42*(4), 1056–1093.

Jack, B. K. and M. P. Recalde (2015). Leadership and the voluntary provision of public goods: Field evidence from Bolivia. *Journal of Public Economics 122*, 80–93.

Jackson, M. O. (2011). An overview of social networks and economic applications. Volume 1 of *Handbook of Social Economics*, pp. 511 – 585. North-Holland.

Kaldor, N. (1936). Wage subsidies as a remedy for unemployment. *Journal of Political Economy 44*(6), 721–742.

Katz, L. F. (1996). Wage subsidies for the disavantaged. *NBER Working Paper 5679*.

Kerr, A. and M. Wittenberg (2019). Earnings and employment microdata in south africa. *WIDER Working Paper* (47).

Kets, W. and A. Sandroni (2017). A theory of strategic uncertainty and cultural diversity. Working paper.

Key, N., E. Sadoulet, and A. D. Janvry (2000). Transactions costs and agricultural household supply response. *American Journal of Agricultural Economics 82*(2), 245–259.

Kingdon, G. G. and J. Knight (2004). Unemployment in South Africa: The Nature of the Beast. *World Development 32*(3), 391408.

Kluve, J., S. Puerto, D. Robalino, J. M. Romero, F. Rother, J. Stöterau, F. Weidenkaff, and M. Witte (2017). Interventions to improve the labour market outcomes of youth: a systematic review. *A Campbell Systematic Review: International Development Coordinating Group*.

Kok, P., M. O'Donovan, O. Bouare, and J. van Zyl (2003). *Post-Apartheid Patterns of Internal Migration in South Africa*. HSRC Publishers.

Lall, S. V., H. Selod, and Z. Shalizi (2006). *Rural-Urban Migration In Developing Countries : A Survey Of Theoretical Predictions And Empirical Findings*. The World Bank.

Layard, P. R. G. and S. J. Nickell (1980). The case for subsidising extra jobs. *The Economic Journal 90*(357), 51–73.

Levinsohn, J. (2007). Two policies to alleviate unemployment in South Africa. *Center for International Development at Harvard, Working Paper* (August), 1–23.

Levinsohn, J. and T. Pugatch (2014). Prospective analysis of a wage subsidy for Cape Town youth. *Journal of Development Economics 108*, 169–183.

Levinsohn, J., N. Rankin, G. Roberts, and V. Schöer (2014). Wage Subsidies and Youth Employment in South Africa : Evidence from a Randomised Control Trial. *Stellenbosh Economic Working Papers 02/14*.

Levitt, S. D. and J. A. List (2007). What do laboratory experiments measuring social preferences reveal about the real world? *The Journal of Economic Perspectives 21*, 153–174.

Madhavan, S. and L. B. Landau (2011). Bridges to nowhere: Hosts, migrants, and the chimera of social capital in three african cities. *Population and Development Review 37*(3), 473–497.

Magruder, J. R. (2012). High Unemployment Yet Few Small Firms: The Role of Centralized Bargaining in South Africa. *American Economic Journal: Applied Economics 4*, 138–166.

Maniadis, Z., F. Tufano, and J. A. List (2014, January). One swallow doesn't make a summer: New evidence on anchoring effects. *American Economic Review 104*(1), 277–90.

Manski, C. F. (2000). Economic analysis of social interactions. *The Journal of Economic Perspectives 14*(3), 115–136.

McKenzie, D. J. (2017). How effective are active labor market policies in developing countries ? a critical review of recent evidence. *World Bank Group, Policy Research Working Paper 8011*(March), 1–32.

Moeletsi, T. (2017). South Africa's Youth Unemployment and the Employment Tax Incentive: An Empirical Re-evaluation. pp. 1–31.

Montgomery, J. D. (1991). Social networks and labor-market outcomes: Toward an economic analysis. *American Economic Review 81*, 140818.

Munshi, K. (2003). Networks in the modern economy: Mexican migrants in the u.s. labor market. *Quarterly Journal of Economics 118*, 54997.

Munshi, K. (2014). Community networks and the process of development. *Journal of Economic Perspectives 28*, 4976.

National Treasury (2016). Employment Tax Incentive Descriptive Report. pp. 1–28.

Neumark, D., I. J. Salas, and W. Wascher (2014). Revisiting the Minimum Wage-Employment Debate: Throwing Out the Baby with the Bathwater? *Industrial Labor Relations Review 67*(3), 608–648.

Neumark, D. and W. Wascher (2004). Minimum wages, labor market institutions, and youth employment: A cross-national analysis. *Industrial and Labor Relations Review 57*(2), 223–248.

Ochs, J. (1995). Chapter 3: Coordination problems. Handbook of Experimental Economics. Princeton, NJ: Princeton University Press.

Patel, K. and F. Vella (2013). Immigrant networks and their implications for occupational choice and wages. *The Review of Economics and Statistics 98*, 12491277.

Posel, D. and D. Casale (2003). What Has Been Happening to Internal Labour Migration in South Africa, 1993-1999?

Posel, D., J. A. Fairburn, and F. Lund (2006). Labour migration and households: A reconsideration of the effects of the social pension on labour supply in South Africa. *Economic Modelling 23*, 836–853.

Potters, J., M. Sefton, and L. Vesterlund (2007). Leading-by-example and signaling in voluntary contribution games: an experimental study. *Economic Theory*, 169–182.

Ragasa, C. and J. Golan (2014). The role of rural producer organizations for agricultural service provision in fragile states. *Agricultural Economics 45*(5), 537–553.

Ranchhod, V. and A. Finn (2016). Estimating the Short Run Effects of South Africa's Employment Tax Incentive on Youth Employment Probabilities using A Difference-in-Differences Approach. *South African Journal of Economics 84*(2), 199–216.

Reed, H. E. (2009). Networks in the new democracy: Internal migration and social networks in late twentieth century south africa. *PhD dissertation*.

Rosenstein-Rodan, P. N. (1943). Problems of industrialisation of eastern and south-eastern europe. *The Economic Journal 53*(210/211), 202–211.

Sobel, J. (2005). Interdependent preferences and reciprocity. *Journal of economic literature 43*(2), 392–436.

Solon, G., S. J. Haider, and J. M. Wooldridge (2015). What are we weighting for? *The Journal of Human Resources 50*(2), 301–312.

Standing, G., J. Sender, and J. Weeks (1996). *Restructuring the Labour Market: The South African Challenge: An ILO Country Review*. International Labour Organization, Geneva, Switzerland.

Stapleton, C. (2015). The Migrant Network Effect: An empirical analysis of rural-to-urban migration in South Africa. *Economic Research Southern Africa Working Paper 504*.

Statistics South Africa (2008). Quarterly Labour Force Survey Quarter 1, 2008.

Stoop, J., C. N. Noussair, and D. van Soest (2012). From the lab to the field: Cooperation among fishermen. *Journal of Political Economy 120*(6), 1027–1056.

Sulla, V. and P. Zikhali (2018). Overcoming Poverty and Inequality in South Africa : An Assessment of Drivers, Constraints and Opportunities. *World Bank report*.

Townsend, R. M. (1994). Risk and insurance in village india. *Econometrica 62*(3), 539–591.

Turiansky, A. (2017). Collective action in games as in life: Experimental evidence from canal cleaning in Haiti. Mathematica Policy Research Report.

Udry, C. (1994). Risk and insurance in a rural credit market: An empirical investigation in northern nigeria. *The Review of Economic Studies 61*, 495–526.

Van Huyck, J. and R. Battalio (2008). Chapter 51 coordination failure in market statistic games. Volume 1 of *Handbook of Experimental Economics Results*, pp. 454 – 462. Elsevier.

Van Huyck, J. B., R. C. Battalio, and R. O. Beil (1990). Tacit coordination games, strategic uncertainty, and coordination failure. *American Economic Review 80*(1), 234–48.

Verick, S. (2012). Giving up job search during a recession: The impact of the global financial crisis on the South African labour market. *Journal of African Economies 21*(3), 373–408.

Vesterlund, L. (2003). The informational value of sequential fundraising. *Journal of Public Economics 87*(3), 627–657.

Viceisza, A. C. G. (2016). Creating a lab in the field: Economics experiments for policy-making. *Journal of Economic Surveys 30*(5), 835–854.

Voors, M., T. Turley, A. Kontoleon, E. Bulte, and J. A. List (2012). Exploring whether behavior in context-free experiments is predictive of behavior in the field: Evidence from lab and field experiments in rural Sierra Leone. *Economics Letters 114*(3), 308–311.

Wolff, J. and G. Stephan (2013). Subsidized work before and after the German Hartz reforms: design of major schemes, evaluation results and lessons learnt. *IZA Journal of Labor Policy 16*(2).

World Bank (2018). The little data book on financial inclusion.

Wydick, B. (2007). *Games in Economic Development*. Cambridge, UK: Cambridge University Press.

Zellner, A. and T. H. Lee (1965). Joint estimation of relationships involving discrete random variables. *Econometrica 33*, 382–394.