

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Labor Productivity and Employment Consequences in East Africa

Permalink

<https://escholarship.org/uc/item/9hp9x7zw>

Author

Hjort, Jonas Krabbe

Publication Date

2012

Peer reviewed|Thesis/dissertation

Labor Productivity and Employment Consequences in East Africa

by

Jonas Krabbe Hjort

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Economics

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:
Professor Edward Miguel, Chair
Professor David Card
Professor Ernesto Dal Bo
Professor Stefano DellaVigna

Spring 2012

Labor Productivity and Employment Consequences in East Africa

Copyright 2012
by
Jonas Krabbe Hjort

Abstract

Labor Productivity and Employment Consequences in East Africa

by

Jonas Krabbe Hjort
Doctor of Philosophy in Economics
University of California, Berkeley
Professor Edward Miguel, Chair

Economic development rarely happens in the absence of large-scale job creation. The scarcity of research on formal employment in Africa in the field of development economics is thus noteworthy. Part of the explanation is that, although steady employment represents an overarching aspiration for many Africans—often preferred, for example, over self-employment or small-scale farming—formal jobs were until recently relatively uncommon on the continent. Variation that can be exploited in statistical analysis is thus hard to come by. Another reason is that few African countries systematically record detailed employment data for large samples of workers. Researchers are therefore typically compelled to collect their own data.

Rapid urbanization and sustained economic growth—including in more labor-intensive sectors—has, however, begun to increase the availability of formal jobs in some parts of Africa, simultaneously enhancing the importance of employment research and the ability of researchers to carry out such research. Focusing on both causes and consequences of formal employment in East Africa, this dissertation examines the effect of ethnic diversity—a characteristic of many African societies—on worker productivity in the Kenyan context, as well as the impact within the household of a parent gaining employment in the Ethiopian context. Knowledge about the factors that constrain labor productivity and the consequences for households once jobs appear is necessary for effective policymaking and a goal for researchers.

I explore both issues in the context of a sector that has been particularly successful in Africa in recent decades: floriculture. A rapid expansion of the sector began in the 1980s; Kenya, for example, is now the third-largest exporter of flowers in the world and supplies approximately 31 percent of flowers imported into Europe (African Business, 2011). Neighboring Ethiopia, with its lower labor costs and abundant land, has more recently been taking market share from other African countries. Agribusiness as a whole is expected to see significant growth in Africa in the coming decades and flower farms account for a notable proportion of formal jobs in Kenya and Ethiopia—such farms are of interest to researchers in their own right. Because the workforce on flower farms often resembles a microcosm of

the labor force as a whole, they also represent a meaningful case study from which broader lessons can be learned.

Two types of data are used in this dissertation: surveys of flower farm workers and applicants (ethnicity, time use, etc) and the output records of a flower farm in Kenya. The farm recorded individual and team output for pay purposes.

The first chapter of this dissertation explores the influence of ethnic diversity on labor productivity in a team production setting. Ethnic diversity has long been known to constrain economic development, but the direct effect on output remains largely unexplored. In Kenya, the land- and water-abundant areas where flower farms are located have experienced immigration from other parts of the country, yielding ethnic diversity in the farms' workforces. I study teams of "packing plant" workers at a large flower farm. Working in teams of three, the workers pack flowers and prepare them for shipping.

I show that ethnically diverse teams are less productive than homogeneous teams. Although an inability to socially sanction non-coethnics may also play a role (see Miguel and Gugerty, 2004 and Habyarimana, Humphreys, Posner and Weinstein, 2007), the primary reason appears to be preference-driven: workers upstream in the triangular production chain lower total output and their own pay by skewing their supply of intermediate flowers toward coethnic downstream workers.

I then go on to analyze the firm's response and the change in the magnitude of the ethnic diversity effect during a period of increased ethnic conflict in Kenya, illuminating how the response of output to diversity is likely to vary across time and space. I find that the productivity loss from ethnic diversity in teams varies with the political environment (see also Posner, 2004). It appears that, in high-cost environments firms are forced to adopt second-best policies to limit discrimination distortions.

Overall chapter 1 shows that inter-ethnic rivalries lower allocative efficiency and productivity in Kenyan floriculture, and highlights the likely consequences for firm behavior and employment growth in the private sector in Africa. The implications for policy and future research are potentially wide-ranging. Most African countries are ethnically diverse and cross-ethnic joint production will increase as urbanization brings together larger groups of workers in cities. Modernization of the economy typically entails greater specialization which also increases the scope for distortions due to ethnic discrimination in production chains.

In the second part of my dissertation, which consists of two separate articles, I focus on the consequences (rather than the causes) of employment. I analyze the effects within the household of a parent gaining employment in rural Ethiopia. Taking advantage of a unique situation in the labor market for farm-workers in Ethiopia at the time, I worked with five flower farms that agreed to randomize fall 2008 hiring due to significant excess demand for jobs and a perceived inability to screen applicants.

In chapter 2, I analyze the impact on children's lives, focusing primarily on time use. Mother's employment has been argued to especially benefit children, but there is little existing evidence to back up such claims. I therefore analyze the effect of mother's and father's employment separately.

The results show that mother’s and father’s employment affects sons and daughters very differently. Daughters spend significantly less time in school when mothers work because they are expected to take over house-work tasks. Daughters’ time use is unaffected by father’s employment, while sons spend significantly more time in school when either parent works. It appears that both the reconfiguration of a parent’s time use implied by employment and the associated increase in income affect children’s time use. Daughters’ human capital accumulation suffers from the greater time requirements of “female” house-work in Ethiopia.

In chapter 3, I analyze the impact of female employment on domestic violence, which is believed to respond to large shifts in spouses’ relative incomes in poor countries. Contrary to the predictions of standard economic models of the household, I find a significant increase in domestic violence when women get employed. The reason appears to be that men in rural Ethiopia attempt to restore their dominance in the household through violence when their relative economic standing is weakened.

In combination chapters 2 and 3 give a rather bleak picture of the influence of female employment on the position of women and girls in poor countries. It is important to recognize that this dissertation focuses on the effects of employment in the short-term, however. In the longer term gender norms may respond to employment, in which case the longer term impact could differ from the deleterious effects observed here. Rather than suggesting that female employment should not be encouraged, the evidence presented thus highlights that theory and employment policy should take traditional gender roles seriously.

In combination, the three chapters of this dissertation highlight that features of society that particularly characterize Africa—such as ethnic diversity in the workforce and time-consuming house-work—interact in first-order order ways with the causes and consequences of employment. We must thus study Africa directly rather than rely on evidence from rich countries when shaping policy.

Beyond seeking to address the substantive issues raised, it is my hope that this dissertation illustrates how direct, micro-level output data can be used to advance research on the determinants of productivity in poor countries, and how a labor market situation often found in developing countries with small formal sectors allows randomized evaluations of an otherwise hard-to-analyze “treatment”—employment itself.

This dissertation is dedicated to my teachers.

Acknowledgments

I am very grateful for invaluable guidance and support from Edward Miguel, David Card, Stefano DellaVigna, Fred Finan and Ernesto Dal Bo. I also thank George Akerlof, Miguel Almunia, Gustavo Bobonis, Oystein Daljord, Willa Friedman, Francois Gerard, Markus Goldstein, Bryan Graham, Pat Kline, Jamie McCasland, Enrico Moretti, Matthew Rabin, Gautam Rao, Ana Rocca, Michel Serafinelli, Lowell Taylor, and numerous other friends, colleagues and seminar participants for comments and suggestions. This work was supported, in part, by fellowships and grants from the Christian Michelsen Institute, the Hewlett Foundation/IEE Dissertation Fellowship, the Institute of Business and Economic Research and the Center for Effective Global Action at U.C. Berkeley, and the Russel Sage Foundation.

Contents

Abstract	1
Acknowledgements	ii
1 Ethnic Divisions and Production in Firms	1
2 Parents' Employment and Child Outcomes	76
3 Backlash: Female Employment and Domestic Violence	106

Chapter 1

Ethnic Divisions and Production in Firms

Abstract

A body of literature suggests that ethnic heterogeneity limits economic growth. This paper provides microeconomic evidence on the direct effect of ethnic divisions on productivity. In team production at a plant in Kenya, an upstream worker supplies and distributes flowers to two downstream workers who assemble them into bunches. The plant uses an essentially random rotation process to assign workers to positions, leading to three types of teams: (a) ethnically homogeneous teams, and teams in which (b) one or (c) both downstream workers belong to a tribe in rivalry with the upstream worker's tribe. I find strong evidence that upstream workers undersupply non-coethnic downstream workers (vertical discrimination) and shift flowers from non-coethnic to coethnic downstream workers (horizontal discrimination), at the cost of lower own pay and total output. A period of ethnic conflict following Kenya's 2007 election led to a sharp increase in discrimination, which did not decay in the nine months after conflict ended. In response, the plant began paying the two downstream workers for their combined output (team pay). This led to a modest output reduction in (a) and (c) teams – as predicted by standard incentive models – but an *increase* in output in (b) teams, and overall. Workers' behavior before conflict, during conflict, and under team pay is predicted by a model of taste-based discrimination. My findings suggest that inter-ethnic rivalries lower allocative efficiency in the private sector, that the economic costs of ethnic diversity vary with the political environment, and that in high-cost environments firms are forced to adopt “second best” policies to limit discrimination distortions.

1. Introduction

There is evidence to suggest that ethnic heterogeneity may impede economic growth. A negative influence on decision-making in the public sphere has been documented: public goods provision is lower and macroeconomic policies of lower quality in ethnically fragmented societies (Easterly and Levine, 1997; Alesina and Spolaore, 1997; La Ferrara, 2003; Miguel, 2004). The possibility of an additional *direct* effect on productivity in the private sector has long been recognized, however. Individuals of different ethnicities may have different skill-sets and therefore complement each other in production, but it is also possible that workers of the same ethnic background collaborate more effectively (Lang, 1986; Lazear, 1998). Evidence from poor countries on the productivity effects of ethnic diversity is largely absent.

This paper provides novel microeconomic evidence on the productivity effects of ethnic divisions. I identify a negative effect of ethnic diversity on output in the context of joint production at a large plant in Kenya where workers were quasi-randomly assigned to teams. I then begin to address how output responds to increased conflict between ethnic groups, how firms respond to lower productivity in diverse teams, and how workplace behavior responds to policies implemented by firms to limit ethnic diversity distortions. A model of taste-based discrimination at work explains my findings across these dimensions.

I study a sample of 924 workers working in teams at a plant in Kenya. The workers package flowers and prepare them for shipping: productivity is observed and measured by daily individual output. The effects of ethnic divisions are of particular importance in the Kenyan context. Tribal competition for political power and economic resources has been a defining character of Kenyan society since independence (Ndegwa, 1997; Oyugi, 1997; Barkan, 2004). Workers at the flower plant are almost equally drawn from two historically antagonistic ethnic blocs - the Kikuyu (and allied tribes) and the Luo (and allied tribes).

Production takes place in triangular packing units. One upstream “supplier” supplies and arranges roses that are then passed on to two downstream “processors” who assemble the flowers into bunches, as illustrated in figure 1a. The output of each of the two processors is observed. During the first period of the sample, processors were paid a piece rate based on own output and suppliers a piece rate based on total team output. Inefficiently low supply of roses to downstream workers of the rival ethnic group was thus costly for suppliers.

I show that the plant’s system of assigning workers to positions through a rotation process generates quasi-random variation in team composition. A worker’s past productivity and observable characteristics are orthogonal to those of other workers in her assigned team. The productivity effect of ethnic diversity can thus be identified by comparing the output of teams of different compositions.

Two natural experiments during the time period for which I have data allow me to go further. During the second period of the sample, in early 2008, contentious presidential election results led to political and violent conflict between the Kikuyu and Luo ethnic groups, but production at the plant continued as usual. In the third period of the sample, starting six weeks after conflict began, the plant implemented a new pay system in which processors were paid for their combined output (“team pay”). By taking advantage of the

three periods observed, I identify (a) the source of productivity effects of ethnic diversity in the context of plant production in Kenya; (b) how the economic costs of ethnic diversity vary with the political and social environment; and (c) how managers responded to ethnic diversity distortions at the plant, and how workplace behavior changed as a consequence of the policies implemented in response.

I model ethnic diversity effects as arising from a “taste for discrimination” among upstream workers: suppliers attach a potentially differential weight to coethnics’ and non-coethnics’ utility, a formulation that follows [Becker \(1974\)](#), [Charness and Rabin \(2002\)](#) and others. The model predicts that discriminatory suppliers in mixed teams will “misallocate” flowers both vertically - undersupplying downstream workers of the other ethnic group - and horizontally - shifting flowers from non-coethnic to coethnic downstream workers.¹ The impact of horizontal misallocation on total output will depend on the relative productivity of favored and non-favored downstream workers. If conflict led to a decrease in non-coethnics’ utility-weight, a differential fall in mixed teams’ output in early 2008 is predicted. Under team pay, a positive output effect of a reduction in horizontal misallocation is expected to offset negative freeriding effects, in teams in which the two processors are of different ethnic groups. The reason is that suppliers can no longer influence the relative pay of the two processors through relative supply under team pay.

Quasi-random assignment led to teams of three different ethnicity configurations. About a quarter of observed teams are ethnically homogeneous, another quarter are “vertically mixed” teams in which both processors are of a different ethnic group than the supplier, and about half are “horizontally mixed” teams in which (only) one processor is of a different ethnic group than the supplier. The ethnicity configurations are displayed in figure 1b. I test the model’s predictions by comparing the average output of teams of different ethnicity configurations within and across the three sample periods.

In the first main result of the paper, I find that vertically mixed teams were eight percent less productive and horizontally mixed teams five percent less productive than homogeneous teams during the first period of the sample. The output gap between vertically mixed and homogeneous teams points to vertical discrimination: it appears that upstream workers are willing to accept lower own pay in order to lower the pay of non-coethnic co-workers. About 86 percent of the output gap between horizontally mixed and homogeneous teams is due to vertical misallocation and 14 percent due to horizontal misallocation. Because Kikuyu and Luo workers are of similar productivity on average, horizontal misallocation has little impact on total output. But the distribution of output across downstream workers is affected: in horizontally mixed teams, processors of the supplier’s ethnic group earn 27 percent more than processors of the other ethnic group.

In the second main result of the paper, I find that the output gap between homogeneous and diverse teams nearly doubled when conflict between the Kikuyu and Luo political blocs began in early 2008. The reason appears to an increase in workers’ taste for ethnic

¹Unless otherwise specified, I use “coethnic” to indicate a processor of the supplier’s tribal bloc, and “non-coethnic” to indicate a processor who is not of the supplier’s tribal bloc. I also use “upstream worker” and “supplier” synonymously, and “downstream worker” and “processor” synonymously.

discrimination. I estimate a decrease of approximately 35 percent in the utility-weight of non-coethnic co-workers when conflict began, through a reduced form approach. As also predicted by the model, there was a small but significant increase in the output of processors of the supplier's ethnic group in horizontally mixed teams in early 2008. A back-of-the-envelope calculation suggest that the decrease in productivity in mixed teams may have cost the farm half a million dollars in annual profit, had it not responded. It is clear from these results that the economic costs of ethnic diversity vary with the political environment.

In the third main result of the paper, I find that the introduction of team pay for processors six weeks into the conflict period led to an increase in output in horizontally mixed team, returning the difference in output between homogeneous and horizontally mixed teams to pre-conflict levels. The increase was likely due to a reduction in horizontal misallocation: a 32 percent output gap between coethnic and non-coethnic processors in horizontally mixed teams was eliminated when team pay was introduced, as predicted by the model. As a result, overall output increased, even though there was a modest decrease in output in homogeneous and vertically mixed teams. These results indicate that that firms are forced to adopt "second best" policies to limit the distortionary effects of ethnic diversity in the workforce when taste for discrimination is high enough. Figure 2 illustrates the evolution of output in teams of different ethnicity configurations during each of the three sample periods observed.

This paper's findings have important implications for theory and policy. Distortionary, taste-based discrimination in production appears to be the primary explanation behind my results. Theories of non-taste-based ethnic diversity effects are unlikely to simultaneously explain a differential fall in mixed teams' output during conflict and equalization of downstream workers' output under team pay. Distinguishing between different channels through which ethnic diversity may affect productivity is important. Higher output in homogeneous teams may be *efficient* if due to technological differences across diverse and homogeneous teams. But discriminatory preferences should lead to distortionary misallocation of resources in most joint production situations in which individuals influence the output and income of others. Interacting economically with individuals of other ethnic backgrounds is hard to avoid when urbanization and economic modernization brings larger groups of workers together, and large multiplier effects are associated with misallocation of intermediate goods (Jones, 2011). The contribution of taste-based discrimination in production to the lower incomes observed in diverse countries may thus be sizable.

The findings of this paper also suggest that relatively brief episodes of conflict can have a long-lasting impact on distortionary attitudes towards individuals of other groups. I find no reversion in ethnic discrimination in the nine months after conflict ended. It appears that the economic costs of ethnic diversity vary with the political environment because social preferences are affected by conflict, forcing firms to adjust their policies in conflictual environments. Entirely removing incentives to discriminate through contractual design is difficult, however. At the plant, biased upstream workers continued to derive less benefit from flowers supplied to pairs of processors that included non-coethnics under team pay. As a consequence, it appears, output in vertically mixed teams was 15 percent lower than in homogeneous teams after team pay was introduced.

This paper contributes to and ties together several areas of research. Its results are to my knowledge the first to carefully identify and explain a negative effect of ethnic diversity on productivity in the private sector, perhaps because well-measured, micro-level output data from poor countries is rarely available.² By showing that a taste for ethnic discrimination can lower output by leading to misallocation of intermediate goods, I also contribute to the literature on workplace favoritism initiated by [Becker \(1957\)](#) and the recent literature on social preferences at work ([Bandiera, Barankay, and Rasul, 2005](#); [Mas and Moretti, 2009](#)). The difference between the findings of [Bandiera, Barankay, and Rasul \(2005\)](#) in the U.K. and my findings in Kenya are particularly interesting. The authors find that “upstream” supervisors at a fruit farm in the U.K., in their allocation of own effort and in their assignment of “downstream” workers to rows with different amounts of fruit, discriminate against workers to whom they are not socially connected *only when doing so is costless to the supervisor*. In contrast, this paper documents an upstream willingness to pay to lower the incomes of non-favored downstream workers, to my knowledge the first paper to do so in data on consequential choices made every day. Ethnic antagonism may be of greater importance to workers in Kenya than social (dis)connections are to workers in the U.K. [Burgess, Jedwab, Miguel, Morjaria, and i Miquel \(2011\)](#) and [La Ferrara \(2002\)](#) show that Africans belonging to a different ethnic group than “upstream” decisionmakers have less access to economic resources in other contexts, suggesting that distortionary discrimination may be a common phenomenon in Africa.

If individuals have discriminatory *preferences*, output is likely to be lower in diverse production units in most production situations in which co-workers affect each other’s income. I begin to address how the productivity effects of ethnic diversity are likely to vary across time and space by studying how workplace discrimination responds to increased ethnic conflict in society, and how firms respond to distortionary discrimination. I follow an innovative paper by [Krueger and Mas \(2004\)](#) in exploring worker behavior during conflict, but my focus is on a poor country characterized by frequent, ethnic tensions. I follow [Ksoll, Macchiavello, and Morjaria \(2010\)](#) in studying Kenyan flower farms during the political crisis of 2008, but focus on the effect of conflict on distortionary attitudes towards non-coethnics. As such, this paper also adds to an emerging literature investigating how social preferences are shaped ([Bauer, Cassar, and Chytilová, 2011](#); [Jakiela, Miguel, and te Velde, 2011](#)).

How firms respond to distortions due to ethnic diversity and how to optimally organize production in the presence of discriminatory attitudes is an exciting venue for future research. [Prendergast and Topel \(1997\)](#) provides a theoretical analysis of the influence of favoritism on optimal compensation and extent of authority for managers. In studying the motivation behind the introduction of team pay at the plant, this paper is particularly re-

²There is a literature on the effects of demographic diversity in production in rich countries, although it consists primarily of theoretical work and descriptive empirical studies. [Lazear \(1998\)](#) provides an interesting theoretical discussion of the potential costs and benefits of diversity in joint production situations. [Hamilton, Nickerson, and Owan \(2005\)](#) analyzes the effects of diversity in joint production in a setting in which workers selected into teams as a factory in California switched from individual to joint production. See [Alesina and Ferrara \(2005\)](#) for a survey of the literature.

lated to [La Ferrara \(2002\)](#) who shows that ethnically diverse cooperatives are more likely to adopt group-pay. I also investigate why the plant chose not to segregate Kikuyu and Luo workers.

Finally, there are interesting connections between this paper’s results on within-firm misallocation and the literature in macroeconomics on across-firm misallocation of capital and intermediate goods in poor countries ([Banerjee and Moll, 2009](#); [Hsieh and Klenow, 2009](#)). First, some of the distortionary policies studied by macroeconomists may exist in part as a means for politicians to skew the distribution of resources towards their own ethnic groups and thus ultimately arise from biased preferences upstream. Second, firms whose output suffers from internal misallocation due to ethnic diversity distortions may survive due to macro-level misallocation of capital. [Jones \(2011\)](#) points out that to understand development we need to understand both why misallocation occurs and the intermediate goods and linkages through which its effects are amplified.

The paper is organized as follows. In section 2, I describe the setting and the organization of production at the plant, outline the data used, and test for systematic assignment to teams. The model of upstream discrimination is presented in section 3, and its predictions for the three sample periods observed tested in section 4. Section 5 explores the extent to which other ethnic diversity mechanisms may explain my results. Section 6 investigates the response of distortionary attitudes towards non-coethnics to conflict in more depth, and section 7 the plant’s response to discrimination. Section 8 concludes.

2. The Setting

2.1 Ethnic diversity and floriculture in Kenya

Ethnic divisions have influenced Kenyan society and politics since independence and contributed to periodical violence. The country’s biggest tribe, the Kikuyu, was favored by Kenya’s British colonizers, a fact that has had long-lasting influence on tribal relations. The Kikuyu has also been the most economically successful and politically influential tribe during most periods of the post-independence era. Although the relationships between different tribes have varied over time, the other major tribes have typically defined themselves politically in opposition to the Kikuyu. In recent years the opposition has been led by the second biggest - and persistently politically active - tribe, the Luo. Most Kenyan tribes have aligned themselves with one of the two associated camps. I therefore categorize workers according to the tribal coalition (“ethnic group”) to which their tribe is seen to belong - the “Kikuyu” (and associated tribes) and the “Luo” (and associated tribes).³

An interesting case study in the context of ethnic divisions is Kenya’s vibrant floriculture sector, which brings together large numbers of workers of different backgrounds. A rapid expansion of the sector began in the 1980s; Kenya is now the third-largest exporter of flowers in the world and supplies approximately 31 percent of flowers imported into Europe ([Noury, 2011](#)). Around 50,000 Kenyans are employed in floriculture, and 500,000 in associated in-

³I designate individuals of the Kikuyu, Embu, Meru, Kamba, Maasai and Kisii tribes as “Kikuyu” and those of the Luo, Luhya and Kalenjin tribes as “Luo”, but focusing on individual tribes instead gives similar results - see section 4.

dustries. Flower farms are part of the fastest growing sub-sector of the Kenyan economy (Kenya Flower Council, 2011). Production takes place on large farms that typically sell their product through auctions in The Netherlands. Most flower farm employees work either in greenhouses (growing and harvesting) or packing plants (packing and preparing flowers for sale).

On some farms, including the one I focus on, workers reside on farm property in gated communities. Such farms essentially constitute a miniature society - complete with schools, health clinics and other amenities - in which groups of individuals from different ethnic backgrounds live and work together. Flower farm jobs are considered relatively desirable.

2.2 Organization of production at the plant

The sample farm primarily produces roses. Plant workers are roughly equally divided across three halls. Packing takes place in three-person teams, as depicted in figure 1a. One upstream “supplier” supplies two downstream “processors” working on separate tables. The supplier brings flowers arriving from the greenhouses to her worktable and throws out poor quality flowers. She then sorts flowers of different lengths/types into piles that are placed on the worktable of one of the processors. The processors remove leaves, cut flowers down to the right size, and finally create bunches that are labeled with the worker’s ID number. Nearly all workers are observed in both positions (supplier and processor).

My primary data source is records of daily processor output from 2007 and 2008. There are 924 packing plant workers in total. The quantities produced were recorded on paper by the farm for remuneration purposes and subsequently converted to electronic format by the research team. A survey provides additional information about workers’ experience, ethnicity, birthplace and other background information. Summary statistics are in table 1. 59 percent of workers are female and 46 percent Kikuyu. The average worker is 35 years old and has five years of tenure at the factory. These figures are similar for Kikuyu and Luo workers.

On average, workers are observed working for 22 days followed by two leave days. When a worker takes leave, another worker returning from leave joins the two remaining workers. Teams are observed for 10 consecutive days on average, but because there is substantial variation in the length of individual work spells, the same is true for team spells. The length of work spells is statistically unrelated to characteristics of workers and teams. 28280 different teams are observed during the sample period. Individual workers are observed on 90 different teams on average.

Suppliers are paid a piece rate w per rose finalized by the processors supplied throughout the sample period. In 2007, the first year of the sample period, each rose finalized by a processor earned her a piece rate $2w$. Workers thus earn the same when working as a supplier and as a processor on average.⁴ In February 2008 the factory began paying the two processors based on their combined output, which led to a change in suppliers’ incentives that I exploit in section 4.

⁴Workers were additionally paid a small fixed component.

2.3 Assignment to teams at the plant

Identification of the productivity effects of ethnic diversity is complicated by the fact that individuals typically sort into joint production, or are assigned to production units so as to maximize productivity. Any third factor that influences both a team’s productivity and its ethnicity configuration will induce spurious correlation between team output and diversity.

The plant I study is ideal for analyzing the impact of ethnic diversity on productivity because of its position rotation system. The supervisors described the system as follows. Workers returning from leave were assigned to open positions in the order in which they arrived at the plant in the morning. Supervisors would start in one corner of a packing hall and work their way through open positions row by row. A priori it is difficult to see how such an assignment system could lead to systematic correlation between the characteristics of the workers in a team.

The team ethnicity configuration classification I use is depicted in figure 1b. With 46.10 percent Kikuyu and 53.90 percent Luo workers, 25.46 percent of teams should be ethnically homogeneous, 24.85 percent vertically mixed and 49.69 percent horizontally mixed, if assignment was random. The percentages observed in the data are 25.64/49.61/24.66 ($p = 0.85$) during the pre-conflict period, 27.38/48.35/24.26 ($p = 0.44$) during the conflict period and 25.32/49.26/25.42 ($p = 0.68$) during the team pay period.⁵ Appendix figure 1 displays the distribution of co-workers’ tribe (and other characteristics) across Kikuyu and Luo suppliers, during each of the three periods. It is clear that workers are not assigned to, or sort into, teams based on ethnicity.

A possible concern is that the underlying (individual or joint) productivity of workers that end up in homogeneous teams may nevertheless differ from that of workers in diverse teams, for reasons unrelated to ethnicity itself. Suppose that individuals are equally productive in homogeneous and diverse teams but prefer interacting with coethnics, as in [Becker \(1957\)](#). In that case it may for example be that supervisors assign well-liked, high-productivity workers to desirable homogeneous teams. Appendix figure 2 displays the distribution of workers’ gender, years of education and years of experience across homogeneous, horizontally mixed and vertically mixed teams, during each of the three sample periods. The distributions are essentially identical. A formal test of quasi-random assignment is in table 2. The matrices in the table display the characteristics, $\text{tribe} \times \text{gender} \times \text{past productivity}$, of one worker in the row dimension, and those of another worker in the team in the column dimension. The proportion of teams observed in a given cell is shown, as well as the proportion expected under the null hypothesis of independence between the row worker’s characteristics and the column worker’s characteristics. Because the worker rotation system leads to complex temporal correlation in team composition and output, the assumptions required for validity of Pearson’s chi-square tests would be violated if all data was used. A periodical “snapshot” of data is thus used in the table: team compositions on the first day of every month.⁶ For

⁵The pre-conflict period is 2007. The conflict period is here considered the first six weeks of 2008, when processors were paid individually. The team pay period is the remainder of 2008 (see section 4).

⁶The tests are insignificant if data from other days is used instead. Note that the table uses three, binary worker characteristics in order to avoid small cell sizes and enable a visual presentation of the results. The

the same reason, productivity is measured by a worker's average output in month $t - 2$. The chi-square tests give no indication of systematic team assignment in any of the three sample periods.

In the context of the plant I study, quasi-random assignment is less surprising than one might think. Supervisors had little incentive to attempt to optimize team assignment,⁷ and little ability to do so given their limited knowledge of worker characteristics and the plant's leave and rotation system.⁸ Managers appeared to be unaware of systematic differences in output across teams of different ethnicity configurations during the first year of the sample period, their limited attention to the packing plant perhaps due to labor costs making up a relatively low proportion of flower farms' total costs (EDRI, 2008).

To alleviate any remaining concerns about systematic team assignment, individual fixed effects are used in the main regressions of the paper.

3. Discrimination: Theoretical Framework

3.1 Set-up

In the context of a joint production situation in Kenya in which workers perform standardized, repetitive tasks, it is reasonable to expect non-positive effects of ethnic diversity in teams. The simple, triangular structure of production at the plant also suggests that for example technological diversity effects - better communication in homogeneous teams, say - and informational diversity effects - e.g. downward-biased beliefs about non-coethnics' productivity - may have limited influence on output. Becker (1957) points out that the presence of other mechanisms is not necessary to explain potentially higher output in homogeneous production units if individuals have discriminatory tastes.

In this section, I present a simple framework in which the supply of intermediate flowers is skewed towards downstream workers of the supplier's ethnicity, and output thereby lowered, if suppliers have discriminatory preferences. The model's predictions are tested in the next section; in section 5 I consider the ability of non-taste-based ethnic diversity effects to explain the results.

Let production take place in teams consisting of one supplier and two processors, the supplier being paid w per rose produced by the team and each processor $2w$ per rose produced by the processor in question. Let processor output depend on supplier effort and ability, e_{sp}

Supplier - Processor 2 matrix is not displayed because the two processor positions are "interchangeable", but the chi-square statistics are insignificant also for that pair of workers.

⁷Supervisors were rarely, if ever, promoted, and their pay did not depend on performance.

⁸Team rotation was unavoidable given the system of irregularly timed leave. The payroll department's representatives, who managed the leave system, explained that the system's flexibility reflected a demand from union representatives and management inertia. Having their families on-site and being able to take leave when needed apparently made infrequent leave acceptable to plant workers. Supervisors found out who was on duty on a given day as team assignment was taking place. An attempt at optimizing assignment by supervisors would thus (i) need to be accomplished in "real time", (ii) be constrained by the available workers returning from leave on a given day, and (iii) be further complicated by the fact that supervisors had limited knowledge of specific workers' characteristics. The reason is that management attempted to attract supervisors that were not socially connected to the rank and file, and low pay relative to the outside options of those considered qualified for supervisor jobs led to high turn-over.

and α_s , and on processor effort and ability, e_p and α_p , through a concave output function displaying decreasing returns to scale, $q_p = f(e_{sp}, \alpha_s, e_p, \alpha_p)$. Worker i 's costs of production are given by an increasing and convex function of her total effort, $d(\sum e_i)$. Assume that the supplier and processors choose effort simultaneously.⁹ Finally, assume that the supplier attaches weight θ_p to the utility of processor p , where θ_p can be either positive or negative.¹⁰ Suppliers with a different weight for coethnics and non-coethnics have discriminatory preferences.

A processor thus maximizes her utility of pay minus her cost of effort:

$$\text{Max}_{e_p} 2wf(e_{sp}, \alpha_s, e_p, \alpha_p) - d(e_p)$$

and the supplier her utility of pay minus her cost of effort plus the additional utility (or disutility) she derives from the well-being of processor 1 and processor 2:

$$\begin{aligned} \text{Max}_{e_{s1}, e_{s2}} w(f(e_{s1}, \alpha_s, e_1, \alpha_1) + f(e_{s2}, \alpha_s, e_2, \alpha_2)) - d(e_{s1} + e_{s2}) \\ + \theta_1 (2wf(e_{s1}, \alpha_s, e_1, \alpha_1) - d(e_1)) + \theta_2 (2wf(e_{s2}, \alpha_s, e_2, \alpha_2) - d(e_2)) \end{aligned}$$

If we abstract from the supplier's cost of effort for purposes of illustration, the analogy between the specification here and [Becker \(1957\)](#)'s specification of a taste for discrimination is clear. The supplier derives $w(1 + 2\theta_1)$ benefit from a unit of q_1 produced. If θ_1 is negative, the supplier is willing to pay out-of-pocket to lower the utility of processor 1. $2\theta_1 w$ is then effectively a Becker-style "discrimination coefficient". However, even if the supplier derives positive utility from ceteris paribus improvements in processor 1's well-being, she may be willing to accept lower own income in order to lower the income of processor 1 relative to processor 2 if $0 < \theta_1 < \theta_2$.

A full model with output a Cobb-Douglas function of its arguments is developed in the theoretical appendix, and the propositions it implies shown (proofs are in the online theoretical appendix). Here I lay out the intuition of the framework in the appendix and discuss its predictions for each of the three sample periods observed.

3.2 Pre-conflict period

Let $\theta_i = \theta_C$ if processor i is of the supplier's ethnic group, and $\theta_i = \theta_{NC}$ if not. Processors are then observed in four different positions: in homogeneous teams (H), in vertically mixed teams (VM), and in horizontally mixed teams in which the processor in question may

⁹In reality, supply and processing decisions take place continuously throughout the work-day. A processor working fast relative to the supplier will at times be held up waiting for more roses, whereas the work-table of a processor working slowly will be overflowing with flowers. Early on, the supplier and processor likely react to each others' speeds; after a while an equilibrium work speed may be reached. When the time unit of the data to which the model must be compared is a whole work-day, the process of adjustment and re-adjustment to the speed(s) of the other worker(s) can sensibly be approximated by assuming simultaneous moves.

¹⁰This formulation follows [Becker \(1974\)](#), [Charness and Rabin \(2002\)](#) and others.

(HM, C) or may not (HM, NC) be of the supplier's ethnic group. From a team perspective there are three types of ethnicity configurations, as illustrated in figure 1b. For conciseness, I do not distinguish between the two specific ethnic groups here; homogeneous teams may for example be either Kikuyu-Kikuyu-Kikuyu or Luo-Luo-Luo. When informative, I highlight the additional cases to be considered if ability or taste for discrimination differs across the two ethnic groups.

The model predicts that processor output is increasing in own ability and the ability of the supplier, but decreasing in the ability of the other processor, in equilibrium. A processor's output is also increasing in the weight the supplier attaches to her utility, but decreasing in the weight of the other processor. The reason is that the upstream worker, in making her supply decisions, considers not only her direct utility from pay, but also the indirect benefits she derive from the output of each of the two processors due to her weight on their utility. If the supplier has discriminatory preferences ($\theta_C > \theta_{NC}$), the model predicts that processor output is higher (a) when working with a coethnic supplier, and (b) when working with another processor who is not of the supplier's ethnicity: $q_{HM,C} > q_H > q_{VM} > q_{HM,NC}$.

Consider now the impact of upstream favoritism on total team output. Biased suppliers are predicted to discriminate both "horizontally" and "vertically" in mixed teams. Vertical discrimination occurs when an upstream worker undersupplies a processor of the other ethnic group - irrespective of her supply to the other processor in the team - because the returns to effort devoted to supplying non-coethnics are lower. Horizontal discrimination occurs when biased suppliers additionally "shift" roses from non-coethnic to coethnic processors, in which case the relative supply to the two processors deviates from that based on their relative productivities. Vertical and horizontal misallocation of roses is predicted to lower team output, so that output is higher homogeneous than in mixed teams: $Q_H > Q_{VM}$ and $Q_H > Q_{HM}$.¹¹

Total supply will be lower in vertically mixed teams than in horizontally mixed teams because the degree of vertical discrimination in teams is increasing in the number of non-coethnic downstream workers. But horizontal misallocation is predicted to occur only in horizontally mixed teams. The impact of horizontal misallocation on the average output of horizontally mixed teams will depend on (a) the ethnic make-up of the population of workers and (b) the relative productivity of individuals of different ethnic groups. If on average flowers are shifted towards comparatively unproductive workers when the two processors are of different ethnic groups, output in horizontally mixed teams may be lower than in vertically mixed teams. Otherwise output is expected to be lowest in vertically mixed teams.

The framework also predicts that a processor's output will benefit more from higher supplier ability (a) when working with a coethnic supplier, and (b) when working with another processor who is not of the supplier's ethnicity. The reason is that biased high-

¹¹Note that horizontal misallocation occurs in this framework because the supplier's cost of effort function is convex in the sum of effort devoted to supplying the two processors. If instead the cost of effort devoted to one processor was separable from the cost of effort devoted to the other processor, horizontal misallocation would not occur. The assumption made here would appear more reasonable.

ability suppliers allocate more of their additional capacity to supplying coethnic processors. Thus: $\partial q_{HM,C}/\partial\alpha_s > \partial q_H/\partial\alpha_s > \partial q_{VM}/\partial\alpha_s > \partial q_{HM,NC}/\partial\alpha_s$.

3.3 Conflict period

It is possible that the period of ethnic conflict in Kenya in early 2008 led to a change in attitudes towards co-workers of the other ethnic group, which I model as a change in θ_{NC} . If θ_{NC} falls, the output of the processor of the supplier's ethnicity in horizontally mixed teams is expected to increase - the relative benefits of supplying such processors go up in that case. A decrease in the output of non-coethnic processors is expected if θ_{NC} decreases. The fall in output will be greatest for non-coethnic processors in horizontally mixed teams because the *relative* benefits of supplying a non-coethnic processor also decrease when the other processor is of the supplier's ethnicity. Unless conflict changes θ_C , output in homogeneous teams should be unaffected by conflict.¹² Thus, $\partial q_{HM,C}/\partial\theta_{NC} < 0 = q_H/\partial\theta_{NC} < \partial q_{VM}/\partial\theta_{NC} < \partial q_{HM,NC}/\partial\theta_{NC}$.

3.4 Team pay period

Six weeks into the conflict period the plant began paying processors for their combined output. Under such a pay system, processor 1's utility from pay is $w(q_1 + q_2)$, rather than $2wq_1$. Because a processor's pay thus partly depends on the effort of the other processor in the team, freeriding is expected, which will have a negative influence on output in all teams.

The supplier's pay system did not change, but due to her social preferences the supplier's problem changes and becomes:

$$\begin{aligned} \text{Max}_{e_{s1}, e_{s2}} \quad & w(f(e_{s1}, \alpha_s, e_1, \alpha_1) + f(e_{s2}, \alpha_s, e_2, \alpha_2)) - d(e_{s1} + e_{s2}) \\ & + (\theta_1 + \theta_2)w(f(e_{s1}, \alpha_s, e_1, \alpha_1) + f(e_{s2}, \alpha_s, e_2, \alpha_2)) - \theta_1 d(e_1) - \theta_2 d(e_2) \end{aligned}$$

In scenarios in which the two downstream workers are of the same ethnic group - homogeneous and vertically mixed teams - the supplier's problem reduces to the same problem she faced under individual pay. In such teams, equilibrium production is thus expected to fall under team pay due to processor freeriding: $Q_H^{TP} < Q_H$ and $Q_{VM}^{TP} < Q_{VM}$.

Because the two processors in a team are paid the same under team pay, the supplier is unable to increase her own utility by "shifting" flowers from less to more favored processors. The average output of coethnic and non-coethnic processors in horizontally mixed teams is thus expected to be equal under team pay, even if suppliers have discriminatory preferences: $q_{HM,C}^{TP} = q_{HM,NC}^{TP}$. The impact of team pay on total output in horizontally mixed teams will depend on the relative magnitude of the positive effect of eliminating horizontal misallocation and the negative effect of processors freeriding on each other: $Q_{HM}^{TP} \gtrless Q_{HM}$.¹³

¹²The results of Eifert, Miguel, and Posner (2010) indicate that individuals' weight on coethnics' utility may increase during periods of ethnically-based political conflict. In that case, output in homogeneous teams is expected to increase in early 2008. Otherwise, no change in the output of homogeneous teams is expected.

¹³It is not the case in this framework that more is supplied to non-coethnic processors in horizontally mixed teams under team pay. This is because assuming simultaneous moves means that the supplier does not take processors' cost of effort into account when making her supply decisions.

Because biased suppliers’ incentive for vertical discrimination remains under team pay, output in homogeneous teams is expected to continue to exceed that in vertically mixed teams, if suppliers have discriminatory preferences: $Q_H^{TP} > Q_{VM}^{TP}$.

Figure 3 and table 3 summarize the predictions to be tested. The table also highlights where the predictions of non-taste-based models of ethnic diversity effects differ. In the next section I interpret the results in light of the model presented here; in section 5 I discuss the ability of non-taste-based mechanisms to explain the results.

4. The Effect of Ethnic Diversity on Productivity

4.1 Investigating the shape of the production function

The framework presented in the previous section predicts that processor output is increasing in processor and supplier ability, but decreasing in the ability of the other processor in the team. In order to correctly interpret observed ethnic diversity effects, it is useful to investigate the shape of the production function. Proxies for workers’ ability as processor and as supplier are needed. I follow an approach comparable to that in [Mas and Moretti \(2009\)](#). Individual ability proxies are first estimated controlling for co-workers’ identities. Focusing on homogeneous teams, processor p ’s output $q_{p,d}$ is regressed on indicator variables for processor p being worker i , supplier s worker k , and other processor o worker j , on date d :

$$q_{p,d} = \alpha_i^{p'} D_{i,d}^p + \beta_j^o D_{j,d}^o + \alpha_k^{s'} D_{k,d}^s + \varepsilon_{p,d}$$

where $D_{i,d}^p = 1$ if $p = i$ on date d . $D_{j,d}^o$ and $D_{k,d}^s$ are defined analogously. $\hat{\alpha}_i^p$ then provides an estimate of i ’s “permanent productivity” as processor and $\hat{\alpha}_i^s$ as supplier.¹⁴

Focusing on homogeneous teams during the first year of the sample period, figure 4 non-parametrically depicts how average processor output varies with (a) processor permanent productivity (across the x-axis), (b) supplier permanent productivity (across the plot lines),

¹⁴Two limitations of this approach should be noted. (1) Ability proxies would ideally be estimated on, say, one half of the data, and then used in second-stage analysis using outcome data from the other half of the data. But the two-stage approach yields inconsistent and downward-biased estimates of the effect of one worker’s ability on another worker’s output in the second stage when T is fixed ([Arcidiacono, Foster, Goodpaster, and Kinsler, 2011](#)). This is likely unproblematic in [Mas and Moretti \(2009\)](#) because their data has a very large number of observations per worker over time, but the dataset used here, while also large, is significantly smaller than the one in [Mas and Moretti \(2009\)](#). Because a large T is important in two-stage approaches, I estimate the ability proxy using the whole period of data observed. (2) If the exact approach in [Mas and Moretti \(2009\)](#) was followed, $q_{p,d}$ would be regressed on $D_{i,d}^p$ and team dummies. However, in the current setting a team is defined as a specific worker in the supplier position and two other workers in the processor positions. The [Mas and Moretti \(2009\)](#) approach therefore provides no natural way to estimate supplier ability proxies (and only two processors share a given team dummy). I therefore use additive, individual fixed effects. Because it is not clear that additive fixed effects provide consistent estimates of structural ability parameters in a non-linear production function, I simulate the model developed in the appendix in order to investigate how informative $\hat{\alpha}^p$ and $\hat{\alpha}^s$ are. The amount of output data used in the analysis is created using the model’s expression for equilibrium $q_{p,d}$ and the empirical hazard rates for a worker leaving a team. Parameter values are chosen so as to approximate the observed mean and standard deviation of output. Estimating $\hat{\alpha}^p$ and $\hat{\alpha}^s$ on the simulated data using additive fixed effects then give correlations of 0.9 between $\hat{\alpha}^p$ and α^p and 0.75 between $\hat{\alpha}^s$ and α^s .

and (c) other processor permanent productivity (across panel A and B).¹⁵ As predicted by the model, processor output is increasing in processor and supplier productivity throughout the range. The reason why supplier productivity has a positive effect on output regardless of how slow the processor is (and vice versa) may for example be that tasks are not clearly separated. In that case a fast supplier can finish more of the work involved in packing a bunch of roses when working with a slow processor. Other processor's productivity appears to have a small but negative effect (unless the supplier is of low productivity) which indicates that upstream workers consider the benefits of supply to both downstream workers when making their supply decisions.

4.2 Productivity in homogeneous and diverse teams: testing the predictions of the model in the pre-conflict, individual pay period

In the context of the plant, the productivity effect of ethnic diversity can be identified by comparing the output of teams of different ethnicity configurations. I begin by focusing on the first year of the sample period, when processors were paid based on own output, and before conflict began.

The histogram in figure 5 displays mean output by team ethnicity configuration in 2007, distinguishing between teams with Kikuyu and Luo suppliers. Confidence intervals are shown but are narrow. The magnitudes in the histogram are in the notes to the figure, along with the standard errors. Note first that there are no significant differences between teams with Kikuyu and Luo suppliers. Most importantly, all-Kikuyu teams are on average as productive as all-Luo teams. Given the nature of work at the plant, this is arguably unsurprising. Focusing instead on output differences that point to discriminatory behavior, it is also the case that the output gap between Kikuyu-Luo-Luo and all-Kikuyu teams is not significantly different from the output gap between Luo-Kikuyu-Kikuyu and all-Luo teams. The same is true for the gap in output between homogeneous and horizontally mixed teams. The evidence in figure 5 thus suggests that Kikuyu and Luo workers are of similar ability and equally discriminatory on average. These results enable a more concise presentation of the evidence to follow. In the remainder of the paper, I do not distinguish between specific ethnic groups and instead focus on the relation between the ethnic backgrounds of workers in a team.

It is clear in figure 5 that team output is highest in homogeneous teams and lowest in vertically mixed teams, with output in horizontally mixed teams falling in between the two. The distribution of team and processor output in teams of different ethnicity configurations is displayed in figure 6. Notably, the density of output for coethnic processors in horizontally mixed teams is shifted to the right of that in homogeneous teams. Conversely, the density of

¹⁵ $\hat{\alpha}^p$ is normalized to have the mean and standard deviation of processor output, and $\hat{\alpha}^s$ the mean and standard deviation of team output. Note also that, because all suppliers in a packing hall obtain roses from the same "pool" of flowers arriving from the greenhouses, mechanically negative across-team "peer effects" should in theory be observed: less flowers are left for other teams if a given team is more productive. But such effects should be small for a sample of the size considered here, and other teams of different configurations should not be differentially affected.

output for non-coethnic processors in horizontally mixed teams is shifted to the left of that in vertically mixed teams. The distributions appear close to normal.

Regression results corresponding to figure 5 are in table 4. The effects are very precisely estimated. Including individual fixed effects in the regressions has little influence on the results, as expected given quasi-random assignment to teams.¹⁶ The output of processors in vertically mixed teams is eight and a half percent lower than that of processors in homogeneous teams, an output gap that is also reflected in the total output of vertically mixed teams. As predicted by the model, upstream workers discriminate against non-coethnics downstream by undersupplying them, it appears. Such discrimination lowers final output.

The results in table 4 also indicate that suppliers discriminate horizontally. It is important to distinguish between the two processors in horizontally mixed teams. The output of the non-coethnic processor is eighteen percent lower than that of processors in homogeneous teams, and nine percent lower than that of processors in vertically mixed teams. The output of the coethnic processor is seven percent higher than that of processors in homogeneous teams. That processor output is lower if the other processor is of the same ethnicity as the supplier points to horizontal favoritism, as predicted by the model. As [Becker \(1957\)](#) emphasized, favored workers benefit from discrimination against non-favored workers. In some situations, such benefits may give favored individuals an incentive to maintain ethnic divisions in society.

Recall that the output loss from horizontal discrimination will depend on the relative productivity of favored and non-favored downstream workers. In the context of the farm, the two ethnic groups are similarly-sized, and we saw above that Kikuyu and Luo workers appear to be of similar ability on average. In such a situation, the output of vertically mixed teams is expected to be lower than that of horizontally mixed teams, which is what we see in table 4. Although vertically mixed are in aggregate four percent less productive than horizontally mixed teams, the lowest output processors are found in horizontally mixed teams. Even if the impact of horizontal discrimination on total output is limited when workers of different ethnic groups are of similar ability, the distribution of output across downstream workers is significantly affected.

Suppose, for purposes of illustration, that in the absence of misallocation of roses *across* the two processors in a team, the output of a coethnic processor in a horizontally mixed team would be equal to that of a processor in a homogeneous team. Similarly, suppose that in such a scenario the output of a non-coethnic processor in a horizontally mixed team would be equal to that of a processor in a vertically mixed team. In that case we can decompose the output gap between homogeneous and horizontally mixed teams: 14 percent would be due to the effect of horizontal misallocation and 86 percent due to vertical misallocation.¹⁷ While the magnitude of the “misallocation multiplier” associated with horizontal

¹⁶The average output associated with all types of teams is slightly higher when individual fixed effects are included, but the estimates of the output gap between teams of different configurations are essentially unaffected.

¹⁷This decomposition is subject to caveats in that it ignores the convexity of effort costs, and it is not clear that the effect of vertical and horizontal misallocation is “additive”.

discrimination will depend on the relative productivity of those being favored and those being discriminated against, generally speaking intermediate goods not being passed downstream will tend to lower final output more than intermediate goods being “invested” in a less productive downstream producer.

The model also predicts that higher ability upstream workers will allocate more of their additional capacity to supplying downstream workers of their own ethnic group. In table 5, processor output is regressed on the proxy for supplier ability estimated above, interacted with team ethnicity configuration dummies. The results show that higher supplier productivity benefits non-coethnic processors less than coethnic processors, suggesting that $\partial q_{HM,C}/\partial \alpha_s > \partial q_{VM}/\partial \alpha_s > \partial q_{HM,NC}/\partial \alpha_s$ and $\partial q_H/\partial \alpha_s > \partial q_{VM}/\partial \alpha_s > \partial q_{HM,NC}/\partial \alpha_s$. The effect of supplier ability is not significantly different for processors of the supplier’s ethnic group in homogeneous and horizontally mixed teams.

In light of the model above, the results we have seen so far suggest that suppliers have discriminatory preferences. The output of a processor depends on her ethnic background in relation to that of the supplier, and on the ethnicity of the other processor in relation to that of the supplier. The reason appears to be that upstream workers undersupply non-coethnics and distort their supply of intermediate flowers to benefit coethnics downstream. By doing so suppliers also lower their own pay. The results thus indicate that upstream workers are willing to pay to discriminate.¹⁸

The contrast between these findings and those of [Bandiera, Barankay, and Rasul \(2005\)](#) is noteworthy. The authors explore how “upstream” supervisors allocate their own effort and rows with different amounts of fruit across favored and non-favored “downstream” workers at a farm in the U.K. The setting in [Bandiera, Barankay, and Rasul \(2005\)](#) is thus comparable to the one studied here in terms of the tasks performed by upstream and downstream workers. The authors find that supervisors discriminate against downstream workers to whom they are not socially connected, but only when supervisors are paid fixed wages - that is, *only when doing so is costless to the supervisor*. It may thus be that ethnic antagonism is of greater importance to workers in Kenya than social (dis)connections are to workers in the U.K.

I now consider the extent to which explanations other than a negative output effect of ethnic diversity in teams may account for the results in table 4. The focus here is on documenting the causal impact of a team’s ethnic composition on output; in section 5 I consider alternative theories that predict negative ethnic diversity effects but do so for reasons unrelated to discriminatory preferences.

The cleanest possible test for ethnic diversity effects in team production would switch the ethnicity of one worker in the team, holding constant everything else about that worker as well as the two other workers in the team. In table 6 I exploit the rotation system at the

¹⁸An alternative interpretation would be that the cost of effort devoted to supplying non-coethnic downstream workers is greater than that of supplying coethnic downstream workers. Such an interpretation provides a less satisfactory account of the observed occurrence of horizontal misallocation of flowers. The ability of theories of “technological” ethnic diversity effects to explain this paper’s findings is discussed in section 5.

plant to provide arguably comparable evidence. The analysis explores what happens when a worker is replaced by another worker of the same productivity tercile but the other ethnicity, controlling for pair fixed effects for the pair of workers that remain in the team before and after the switch.¹⁹ Note that there is no significant change in output when the outgoing and incoming worker are of the same ethnic group: worker switches do not in themselves affect the productivity of a team.

In columns 1 and 3, the output of an *unswitched* processor is regressed on dummies for the change in team ethnicity configuration when a supplier or processor of productivity comparable to the replaced worker joins the team. For clarity, I lay out the effects for a worker in processor position 1 (processor 2 is analogous). The output of a processor 1 who is of the same ethnic group as the supplier increases by five percent when a processor 2 of the other ethnic group replaces a comparably productive processor 2 of the supplier's ethnic group. When a supplier who is not of processor 1's ethnic group replaces a comparably productive supplier of processor 1's ethnic group, processor 1's output falls by nine percent if the two processors are of the same ethnic group. If instead processor 2 is of the incoming supplier's ethnic group, processor 1's output falls by 25 percent. The output of a processor 1 who is not of the supplier's ethnic group increases by nine percent if a processor 2 of processor 1's ethnic group replaces a comparably productive processor 2 of the supplier's ethnic group.

The estimates for team output in columns 3 and 4 are similar, output falling by five percent when a team goes from being homogeneous to horizontally mixed due to a worker switch, by nine percent when a team goes from being homogeneous to vertically mixed, and by four percent when a team goes from being horizontally to vertically mixed.

Comparing teams that share the workers in two positions and the productivity tercile of the worker in the third position thus yields similar estimates to comparing all teams of different ethnicity configurations, providing reassurance that the estimates in table 4 represent the causal effect of ethnic diversity. If the estimates in table 4 were due in part to for example non-random assignment to teams or differences in ability across the two ethnic groups interacting with non-linear complementarities in the production function, then controlling for pair fixed effects and the third worker's productivity tercile should lead to different estimates.

Figure 7 depicts the temporal response of team output to the "event" of a worker substitution leading to a change in a team's ethnicity configuration. Panels A - C plot the dynamic response of the first difference of output (the change in team output from the day before) to a change in a team's ethnicity configuration, and panels D - E the cumulative response over time. The decrease in output when a team "becomes mixed" is apparent. The first differenced response occurs almost entirely on the first day after the switch: the difference in output between homogeneous and diverse teams is relatively constant through teams' duration.

¹⁹The pair fixed effect for processor pair ij is for example a dummy that takes value 1 if workers i and j are processors in a team together.

The tribal categorization used here is meaningful. Recall that this paper distinguishes primarily between workers designated as belonging to the Luo and Kikuyu tribal blocs. Categorization was on the basis of political alliances and relations between specific tribes. 86 percent of the sample belongs to three tribes: the Kikuyu (41 percent), Luo (30 percent) and Luhya (15 percent). I now consider sub-samples of teams in which workers belong to two specific tribes, focusing on the Kikuyu - Luo, Kikuyu - Luhya, and Luo - Luhya sub-samples. The Luo and Luhya tribes are categorized as belonging to the “Luo” ethnic group in this paper.

The estimates in table 7 provide a clear picture. In a sub-sample of teams consisting of workers from two different tribes categorized as belonging to the same tribal bloc, little if any discrimination against non-coethnic processors occurs. The output of vertically mixed teams is for example not significantly different from that of homogeneous teams in the Luo - Luhya sub-sample. But within two different sub-samples of teams consisting of workers of two specific tribes categorized as belonging to different tribal blocs here, discrimination is pervasive and of an extent similar to that seen in the full-sample analysis in table 4. There are only minor differences across the Kikuyu - Luo and the Kikuyu - Luhya sub-samples, analyzed in columns 1 - 2 and 3 - 4 of table 7 respectively.

So far we have seen strong evidence indicating that team-level ethnic diversity lowers productivity in the context of factory production in Kenya. If diversity effects are driven by discriminatory preferences, then we would expect the negative effect of ethnic diversity on private sector output to vary with factors that influence taste for discrimination, such as the political climate and relations between groups. A shift in taste for discrimination should differentially lower the output of mixed teams. In the next sub-section, I analyze differences in output between homogeneous and mixed teams during the period of ethnically-based, political conflict in Kenya in early 2008.

4.3 Ethnic conflict and the impact of diversity in teams on productivity: testing the predictions of the model in the conflict period

The two coalitions in Kenya’s December 27 2007 presidential election were ethnically based. In advance of the election, opinion polls predicted that the coalition led by Luo challenger Raila Odinga would oust the sitting Kikuyu- led coalition represented by incumbent president Mwai Kibaki. But results were delayed and the Kibaki victory announced on December 29 disputed by the opposition and the international community. Widespread violence against Kikuyu and Kikuyu- allied tribes erupted, and counter-attacks soon followed. More than 1,200 people were killed and 500,000 displaced in the months that followed ([Gibson and Long, 2009](#)). On February 28, a peace agreement was reached, though violence continued in many areas, and it was not until after April 3 when the two sides reached an agreement on the composition of a power-sharing government that the political crisis ebbed.

The conflict period significantly disrupted life in parts of Kenya.²⁰ However, plant supervisors reported that logistics and worker absence at the farm was largely unaffected and that

²⁰[Dupas and Robinson \(2011\)](#) document a dramatic fall in income and consumption for the rural poor in Western Kenya during the crisis. Many flower farms also struggled: [Ksoll, Macchiavello, and Morjaria](#)

production continued as usual. Because the workers live on the farm in a gated community it was safest to remain on the farm. If the plant's ability to operate was nevertheless affected, a decrease in productivity, as measured by the econometrician, should be observed in all teams.

The model predicts an increase in the gap between the average output of homogeneous and mixed teams if attitudes towards workers of the other ethnic group worsened when conflict began. I interpret a possible increase in taste for discrimination as a decrease in the weight attached to the well-being of non-coethnics.²¹

In table 8, the difference in output between mixed and homogeneous teams before and after conflict began is compared.²² Data from 2007 and the first six weeks of 2008 (when processors were still paid based on own output) is used.

There was no significant change in the output of homogeneous teams when conflict began. If suppliers have social preferences, the impact of conflict on the productivity of homogeneous teams will reflect a combination of (at least) two factors. First, farm-wide disruption effects may have negatively affected output in all teams. Second, it is possible that conflict led to an increase in workers' weight on the utility of coethnics: the findings of [Eifert, Miguel, and Posner \(2010\)](#) suggest that Africans increasingly identify with coethnics during times of heightened political competition between groups. I cannot rule out general disruption effects or an increase in the utility workers derive from coethnics' output and income. But the combination of supervisors' reports and a conflict coefficient for homogeneous teams that is essentially precisely zero points to little farm-wide disruption effects and little effect on workers' weight on coethnics' utility.

The output gap between homogeneous and vertically mixed teams nearly doubled in early 2008. Output in vertically mixed teams decreased by seven percent when conflict began. The results in table 8 thus indicate that upstream workers undersupply non-coethnic downstream workers to a significantly greater extent during times of ethnic conflict, as predicted by the model if taste for discrimination increased.

Output in horizontally mixed teams decreased by four percent when conflict began, but there was a small but significant *increase* in the output of coethnic processors in horizontally mixed teams. An increase in upstream discrimination against workers of other ethnic groups thus appears to increase the supply of flowers to those downstream workers who belong to the same ethnic group as suppliers, as predicted by the model. The relative benefits of

(2010) report that the official export volumes of affected farms dropped by 38 percent on average, in large part due to worker absence exceeding 50 percent on average on such farms.

²¹[Eifert, Miguel, and Posner \(2010\)](#) show that ethnic identities can vary over time, and [Charness and Rabin \(2002\)](#) and others show that social preferences generally depend on the behavior of others. As long as social preference weights are partly group-based rather than entirely individual-specific, we would then expect the weight on non-coethnic co-workers' output and utility to deteriorate during a period of increased antagonism.

²²Data from both 2007 and 2008 was de-seasonalized as follows. Let m_i be average output in month i of 2007, and $\bar{m} = \frac{1}{12} \sum_i m_i$. Output observations from month i of both 2007 and 2008 were then multiplied by \bar{m}/m_i .

flowers supplied to coethnic processors in horizontally mixed teams go up if conflict lowers the utility upstream workers derive from non-coethnics' output, even if suppliers' weight on coethnics' utility is unaffected.

In light of the model presented above, the results for the conflict period thus suggest that discriminatory attitudes towards co-workers of other ethnic groups worsened in Kenya in early 2008. It appears that the economic costs of ethnic diversity vary with the political environment. A back-of-the-envelope calculation suggests that the increase in supplier discrimination during conflict may have cost the farm as much as US\$560,000 in profit per year, had it not responded.²³

Firms may be forced to take measures to limit distortions that arise from internal, ethnic discrimination, especially in times of conflict. In the next sub-section, I analyze how the gap in output between homogeneous and mixed teams was affected when the plant six weeks into the conflict period changed the pay system for processors and thereby altered the incentives facted by biased upstream workers.

4.4 Firms' response to distortionary favoritism and the impact of diversity in teams on productivity: testing the predictions of the model in the team pay period

On February 11 2008, the farm began paying processors w per rose finalized by the team, rather than $2w$ per rose finalized by the processor herself as before. As in standard incentive models, the framework above predicts that processors will freeride on each others' effort when paid in part based on the output of the other processor. Freeriding should negatively affect output in all teams, but in horizontally mixed teams an offsetting positive effect is expected. Under team pay, suppliers are unable to influence the relative pay of the two processors through relative supply. If the higher output for processors of the supplier's ethnic group observed under individual pay is driven by suppliers' taste for discrimination, a decrease in the output gap between coethnic and non-coethnic processors in horizontally mixed teams is thus expected when team pay is introduced.

To test these predictions, I consider the period after processors' pay system was changed and through the remainder of 2008 as a single team pay period.²⁴ Figure 8 displays team and individual output during the three sample periods: pre-conflict (2007), conflict (the first six weeks of 2008), and the team pay period (February 11 through 2008). The decrease in output in mixed teams during conflict is apparent. Comparing the second and third periods, the figure also clearly indicates that the introduction of team pay had a positive effect on output in horizontally mixed teams.

²³US\$560,000 \approx 286 (average number of teams observed per day) \times 233 (average number of fewer roses produced per team per day during the conflict period, relative to 2007) \times 365 production days per year \times 0.023 \times (estimated average profit per rose grown in Kenya as estimated in [Melese and Helmsing \(2010\)](#)).

²⁴In principle, we could distinguish between a "team pay / conflict" period and a "team pay / post-conflict" period. But it is unclear exactly when conflict effectively ended, and, as discussed below, differences in output between teams of difference ethnicity configurations remained essentially constant after team pay was introduced.

Corresponding regression results are in table 9. The results indicate that team pay leads to some degree of freeriding among processors: output in homogeneous and vertically mixed teams is 1 percent lower under team pay. The modest magnitude of this effect is noteworthy and interesting in itself.²⁵

Output in horizontally mixed teams is four percent higher under team pay, as seen in columns 3-4 and 7-8 in table 9. The difference in output between horizontally mixed and homogeneous teams thus decreased significantly when team pay was introduced. The introduction of team pay essentially canceled out the effect of conflict on output in horizontally mixed teams, returning the difference in output between homogeneous and horizontally mixed teams to pre-conflict levels.

The increase in horizontally mixed teams' output appears to be due to horizontal favoritism being eliminated when biased suppliers' ability to increase the relative income of favored processors through relative supply was removed, as predicted by the model. There is no statistically significant difference in the output of the coethnic processor and the non-coethnic processor in horizontally mixed teams during the last ten and a half months of 2008. An output gap of 32 percent between processors of the supplier's ethnicity and processors who are not of the supplier's ethnicity in horizontally mixed teams was eliminated by the introduction of team pay.

The positive impact on output in horizontally mixed teams, which make up half of all teams, led to an overall increase in output when team pay was introduced. However, output in horizontally mixed teams remains lower than in homogeneous teams under team pay, and output in vertically mixed teams still lower. Under team pay a biased supplier continues to derive greater benefit from flowers supplied the more downstream workers belong to her tribe. The ranking of output of teams of different ethnicity configurations observed under team pay is thus due to incentives for vertical discrimination remaining in place, it appears.

The model presented above, in which the productivity effect of ethnic diversity in teams arises from a taste for discrimination on the part of upstream workers, thus predicts the output response to the introduction of team pay well. Approximately one fourth of the yearly expected profit loss due to the impact of conflict on misallocation of flowers (had the farm not responded) was avoided through the change in suppliers' contractual incentives.²⁶ It is difficult to imagine a standard economic model of joint production that would predict an increase in output when team pay is introduced.

²⁵As is clear from figure 1, processors can easily monitor each others' effort. A triangular organization of production may thus be a situation in which freeriding can be effectively dampened through co-monitoring. Note that I cannot rule out that other differences between the individual and team pay periods of 2008 contribute to the team pay coefficient for homogeneous and vertically mixed teams. Such time-varying factors should not influence the comparison of different types of teams.

²⁶Note that after the conflict period the plant also hired more plant workers, probably to make up for lost capacity due to the decrease in productivity. Though workers are paid piece rates, overhead costs per worker are significant (housing, etc). The actual change in profit when conflict began, and after team pay was introduced, is thus difficult to estimate. Workers that were hired after conflict began are excluded from the analysis in this paper.

In the previous sub-section we saw that the economic costs of ethnic diversity vary with the political environment. The reason appears to be that distortionary discrimination at work increases during times of conflict. The results in this sub-section suggest that, in high-cost environments, firms adopt “second best” policies to limit the distortions caused by ethnic favoritism. Group-based pay leads to freeriding and reduces output in homogeneous teams, but the new pay system introduced by the plant during the conflict period in Kenya in early 2008 was likely designed to remove the ability of biased upstream workers to increase one processor’s pay relative to the other’s through differential allocation of flowers. Distortionary discrimination fell and the net effect was positive. Interestingly, [La Ferrara \(2002\)](#) also finds that ethnically diverse cooperatives in Nairobi are more likely to adopt group-pay. It thus appears that ethnic diversity has an important influence on how firms organize production in the private sector.

In the next section I discuss the ability of non-taste-based ethnic diversity effects to explain the results we have seen so far.

5. Sources of Ethnic Diversity Effects

Taste-based discrimination is only one of many potential reasons why output may be lower in diverse teams than in homogeneous teams. Distinguishing between different sources of diversity effects is important. Unlike if differential allocation of intermediate goods to coethnic and non-coethnic downstream workers is driven by discrimination, higher supply to coethnic downstream workers may for example be *efficient* if individuals are simply more productive when collaborating with others of their own ethnic group.

The sources of non-taste-based, negative ethnic diversity effects discussed in the literature can be classified into three broad categories:²⁷

1. *Informational diversity effects* arise if upstream workers are risk-averse and better able to judge the productivity of coethnics, or have downward-biased beliefs about the productivity of non-coethnics ([Becker, 1957](#)). Such informational effects will lead to higher supply to downstream coethnics. Note that, in the context of the sample plant, a supplier who supplies “too few” roses to a non-coethnic processor may never learn the processor’s true productivity.
2. *Technological diversity effects* arise if individual productivity is higher when working with coethnics, for example due to better communication or peer effects among coethnics ([Lang, 1986](#)). All workers in the sample speak Swahili, but complicated forms of peer effects could explain the results in table 4.²⁸

²⁷[Habyarimana, Posner, and Weinstein \(2007\)](#) uses a similar classification, but their focus is on explaining why public goods provision is lower in diverse societies.

²⁸In particular, non-linear, ethnicity-specific, positive peer effects could explain the results in table 4. Suppose that workers’ effort responds to co-workers’ effort but only that of coethnics. Suppose further that, within ethnic groups, working with highly productive co-workers increases effort but working with less productive workers does not decrease effort. In that case, homogeneous teams will be more productive than mixed teams.

3. "*Cooperational*" diversity effects arise if coethnics are better able to sustain cooperation (Kandel and Lazear, 1992; Habyarimana, Posner, and Weinstein, 2007). Coordinating on a high effort equilibrium is more easily achieved if deviators can be effectively sanctioned. If workers of different ethnic groups segregate socially, it may be easier to punish deviators within an ethnic community.

I consider these possibilities in turn. Note first that informational and technological diversity effects are unlikely to explain this paper's results. Suppose that the higher output observed in homogeneous teams during the pre-conflict, individual pay period was due to inferior technology or information in diverse teams. In that case it is difficult to see why output in mixed teams would fall differentially during conflict, and why the output of the two processors in horizontally mixed teams would be equalized under team pay.

Cooperational effects have proven difficult to distinguish from social preferences (see for example Bandiera, Barankay, and Rasul, 2005), in part because such theories typically have few testable implications. Some forms of cooperational diversity effects could explain the observed decrease in mixed teams' output during conflict. If trust for example facilitates cooperation, an erosion of trust between workers of different ethnic groups during times of ethnic antagonism could lead to a decrease in mixed teams' output. Other forms of cooperational diversity effects could explain the observed increase in the output of non-coethnic processors in horizontally mixed teams under team pay. Coethnic processors that can exert effective social pressure on the upstream worker may for example induce the supplier to supply more to non-coethnic processors in horizontally mixed teams under team pay because processors derive benefits from the output of the other processor under team pay. It is, however, difficult to think of cooperational or other forms of non-taste-based ethnic diversity effects that can simultaneously explain a decrease in mixed teams' output during conflict, equalization of processors' output in horizontally mixed teams when team pay is introduced, and the other results seen in this paper.

Though I cannot rule out that other forms of ethnic diversity effects also play a role, I thus conclude that the leading explanation for the lower output observed in ethnically diverse teams at the plant is taste-based discrimination on the part of suppliers.²⁹

So far we have seen that output in factory production in Kenya is lower when individuals of different ethnic backgrounds work together, and that the reason appears to be that biased upstream workers undersupply downstream workers of other ethnic groups and misallocate intermediate goods across coethnic and non-coethnic downstream workers. We have also seen that distortionary workplace discrimination is greater during times of conflict, and that firms introduce policies in response in order to reduce workers' incentive to discriminate. By studying how discriminatory preferences are shaped, and how firms choose their response to distortionary discrimination, researchers can go beyond identifying a source of ethnic diversity effects in production and begin to address why those effects vary across space and

²⁹More complicated forms of social preferences than the simple differential weight attached to coethnics' and non-coethnics' well-being in the model above may also explain this paper's results.

time and how profit motives in the private sector can reduce the aggregate effect of ethnic diversity. I address these questions in more depth in the next two sections.

6. Understanding the Response of Workplace Behavior to Conflict

6.1 Magnitude of the increase in taste for discrimination

In this section I explore the magnitude, persistence, and heterogeneity of the response of individuals' taste for discrimination to conflict between groups.

By how much did suppliers' weight on the utility of non-coethnics fall when conflict began? A limitation of studying triangular production units is that I am unable to separately identify the structural parameters θ_C and θ_{NC} because suppliers are never observed working purely for their own benefit. But by taking advantage of the plant's worker rotation system I can bound the impact of conflict on θ_{NC} through a reduced form approach along the lines advocated by Chetty (2009). The required assumption is that θ_C was unaffected by conflict, an assumption supported by the fact that average output in homogeneous teams did not change during the conflict period.

Step 1: Ratios. In the Cobb-Douglas model developed in the theoretical appendix, the ability of the supplier does not influence the relative output of the two processors:

$$\frac{q_1}{q_2} = \left(\frac{\alpha_1}{\alpha_2} \right)^{\frac{2\beta}{2-\beta-2\gamma}} \left(\frac{1+2\theta_1}{1+2\theta_2} \right)^{\frac{2\gamma}{2-\beta-2\gamma}}$$

Step 2: Ratio-of-ratios. Recall that two workers in a team stay put when the third worker is switched for another worker returning from leave. Consider a sample of horizontally mixed teams in which a supplier of processor 1's ethnicity is replaced by a supplier of processor 2's ethnicity (say in between dates $d = 0$ and $d = 1$). In the model in the appendix, the relative ability of the two processors does not influence their relative output under one supplier relative to their relative output under another supplier:

$$\frac{q_{1,d=0}/q_{2,d=0}}{q_{1,d=1}/q_{2,d=1}} = \left(\frac{1+2\theta_C}{1+2\theta_{NC}} \right)^{\frac{4\gamma}{2-\beta-2\gamma}}$$

Taking the ratio of the ratio of processors' output before a supplier switch to the same ratio after the switch can be thought of as the multiplicative model analogue of a difference-in-differences analysis in additive models. We are left with a quantity that depends only on the powers of the output function, θ_C and θ_{NC} .

Step 3: Ratio-of-ratio-of-ratios. Finally, if θ_C was unaffected by conflict, suppliers' weight on coethnics' utility should have the same influence on the ratio-of-ratios before and after conflict. Taking the ratio of the pre- and during-conflict quantities, we arrive at an expression

that relates θ'_{NC} , the weight on non-coethnics' utility after conflict began, to the pre-conflict θ_{NC} :

$$\frac{(q_{1,d=0}/q_{2,d=0}) / (q_{1,d=1}/q_{2,d=1})}{(q_{1,d=0'}/q_{2,d=0'}) / (q_{1,d=1'}/q_{2,d=1'})} = \left(\frac{1 + 2\theta_{NC}}{1 + 2\theta'_{NC}} \right)^{\frac{4\gamma}{2-\beta-2\gamma}}$$

In the empirical appendix I implement the ratios approach. I bound $\Delta\theta_{NC} = (\theta_{NC} - \theta'_{NC})$ by considering a wide range of possible values for θ_{NC} , β and γ . I estimate a fall in θ_{NC} of 0.01 – 0.07 or 8 – 127 percent. Averaging across the parameter space considered, θ_{NC} is estimated to fall by approximately 35 percent when conflict begins.

This calculation is subject to caveats, but it illustrates an important point. If the decrease in mixed teams' output when conflict began was primarily due to a worsening of discriminatory attitudes, as the results in the previous sections suggest, then production data points to a relatively large increase in taste for discrimination against co-workers of the rival tribe in Kenya in early 2008.

6.2 Persistence of the effect of conflict on workplace discrimination

The effect of conflict on discriminatory workplace behavior does not decay in the nine months after conflict ended. In the model of taste-based discrimination above, the impact of conflict on output in diverse teams should persist for as long as attitudes towards workers of other ethnic groups are affected. Periods of increased antagonism may entail significant hidden economic costs if “mean reversion” in taste for discrimination is slow (or does not occur). The evolution of output in teams of different ethnicity configurations across the three sample periods was depicted in figure 2. After the introduction of team pay, average output in both homogeneous and mixed teams was steady for the remainder of the sample period, suggesting that the impact of conflict on social preferences was long-lived.

6.3 Heterogeneity in workers' response to conflict

How did the response to conflict of distortionary discrimination at work vary across individuals? Modeling θ_C and θ_{NC} as parameter values shared by all workers is a simplification: in reality some workers will have a higher taste for discrimination than others. Figure 9 plots the distribution, across individual suppliers, of the difference in output between homogeneous and (vertically and horizontally) mixed teams supplied, before and after conflict began. It appears that most suppliers discriminate against non-coethnic processors during the pre-conflict period. Conflict led to an increase in the output gap between homogeneous and mixed teams supplied for most upstream workers, but also to a notable widening of the distribution of the output gap. The figure indicates that some upstream workers respond more to conflict than others, differentially increasing the extent to which they discriminate against non-coethnics downstream.

Some workers in the sample were more exposed to the conflict period of early 2008 than others. Though the workers at the plant and their co-habiting family-members were not themselves directly affected, 22 percent of workers report to have “lost a relative” during

the conflict.³⁰ The decrease in output in mixed teams when conflict began was significantly greater in teams supplied by such workers, as seen in columns 1 and 2 of table 10. These results indicate that personal grievances exacerbate individuals' workplace response to conflict.³¹

Younger individuals may have more malleable social preferences. In columns 3 and 4 of table 10 we see that, although output in homogeneous teams led by old and young suppliers was similar, output in mixed teams with young suppliers was significantly higher during the first year of the sample period. Young suppliers were less discriminatory towards non-coethnic co-workers than old suppliers before conflict began, it appears. This finding is consistent with an expectation expressed by many Kenya commentators before 2008. It was argued that the young coming of age at the time would be the country's first "post-tribal" generation (Buckley, 1997). The results of table 10 also show that the decrease in output in mixed teams when conflict began was significantly greater in teams with young suppliers, however. Output in mixed teams with young suppliers was no higher than in mixed teams with older suppliers during the conflict period. These results suggest that youth start out relatively tolerant, but that the attitudes of the young towards non-coethnics respond more negatively to conflict.

The results discussed in this section paint a consistent picture of how distortionary attitudes towards workers of other ethnic groups respond to ethnic conflict. It appears that conflict may entail significant hidden economic costs because distortionary social preferences are updated in a "Bayesian" fashion when conflict occurs, at least in the Kenyan context. A serious episode of violent, political conflict between the Kikuyu and Luo blocs led to a significant shift in the average weight attached to the well-being of non-coethnics, a shift that did not decay in the nine months after conflict ended. The negative response was greater among those more affected and among those likely to have a less cemented "prior".

In the next section I analyze how the plant responded to lower output in mixed teams in more depth.

7. Understanding Firms' Response to Ethnic Diversity Distortions

7.1 Benefits of ethnicity-based assignment to teams

Segregating workers of different ethnic groups would appear to be the profit-maximizing response to distortionary discrimination, from the viewpoint of the econometrician. The results in tables 4 and 8 suggest that segregation would have increased plant productivity

³⁰Note that the high proportion of workers reporting to have a lost relative implies a broad definition of "relative".

³¹In ongoing work I am analyzing the response of discrimination and productivity to specific conflict events through an event study comparing workers' behavior on days in which events occurred in their "home district" (where their relatives reside) to other days. I am also exploring how the response of discrimination to conflict events depends on the ethnicity configuration of the supplier's team on an event day in relation to the "configuration" of the event itself. Krueger and Mas (2004) find that a labor dispute in Illinois had a greater impact on employees' workplace behavior when replacement workers and returning strikers worked side by side.

by four percent before conflict and by eight percent after conflict began, relative to the status quo of arbitrary assignment to teams. Are these expected benefits of a magnitude that is likely to be salient to supervisors? Consider the output increase expected from optimally assigning workers to teams and positions based on ethnicity, productivity or both. If we view a worker as having three characteristics - the tercile to which she belongs in the distribution of processor productivity, the tercile to which she belongs in the distribution of supplier productivity, and her ethnicity - then an average output will be associated with teams of each of 3 ethnicity configurations, 18 productivity configurations and 63 ethnicity-productivity configurations.³² In theory, supervisors can then solve the linear programming problem of maximizing total output subject to the expected output associated with a given type of team and the “budget set” of workers available (see [Bhattacharya, 2009](#); [Graham, Imbens, and Ridder, 2011](#)).³³

The optimal assignments and associated expected output gains are shown in table 11.³⁴ Throughout the period observed, the output gains expected from assigning workers to teams based on ethnicity were larger than those expected from assigning workers based on productivity - twice as large during the conflict period. In fact segregation achieves about half the output gains of the “complete” solution. The complete solution assigns workers optimally to fully specified teams and thus takes into account interactions between the three workers’ ethnicities and productivities - a complicated “general equilibrium” problem that is likely infeasible for supervisors to solve.³⁵ It thus appears that the expected productivity gain of segregation is sizable relative to the expected effect of changing other comparable factors under supervisors’ control.

7.2 Costs of ethnicity-based assignment to teams

The fact that the plant chose not segregate workers, even after conflict led to a dramatic drop in productivity in mixed teams, indicates that managers expect there to be costs associated with segregation. I consider two specific possibilities. First, it may be that interacting with co-workers of other ethnic groups in itself dampens discriminatory attitudes over time.

³² $63 = [3 * ((3 * (3 + 1))/2)] + 3^3 + [3 * ((3 * (3 + 1))/2)]$. In teams in which the two processors are of the same ethnic group, the processors (i.e., the productivity terciles of the processors) are “interchangeable” so there are $[3 * ((3 * (3 + 1))/2)]$ homogeneous types of teams and $[3 * ((3 * (3 + 1))/2)]$ vertically mixed types of teams. In horizontally mixed teams, the processors’ productivity terciles are not interchangeable because the higher ability processor may or may not be of the supplier’s ethnic group, so there are 3^3 types of horizontally mixed teams.

³³[Dupas and Bhattacharya \(2011\)](#) and [Carrell, Sacerdote, and West \(2011\)](#) compute welfare-maximizing assignments in other contexts using this technique. An added complexity here is the need to assign workers to both positions and teams.

³⁴“Optimal” here means output-maximizing, as inferred from the data. The output-maximizing solution may be undesirable for other reasons discussed below.

³⁵An example of a fully specified team is the following: {(Worker in processor position 1: 1st productivity tercile as processor, 3rd productivity tercile as supplier, Kikuyu), (Worker in processor position 2: 2nd productivity tercile as processor, 3rd productivity tercile as supplier, Luo), (Worker in supplier position: 1st productivity tercile as processor, 1st productivity tercile as supplier, Kikuyu)}.

Second, the expected benefits of segregation computed in table 11 may not give an accurate picture of the “out-of-sample” possibility of plant-wide segregation.

Boisjoly, Duncan, Kremer, Levy, and Eccles (2006) find that white American college students become more friendly towards and supportive of African American students after spending time with a black roommate. It is possible that a similar effect occurs in a Kenyan workplace, although in a situation in which mixed teams are characterized by discriminatory behavior it is also possible that interaction increases tensions and exacerbates ethnic biases. To investigate, I compare the behavior of suppliers with greater versus lower experience working with non-coethnics, in table 12. Focusing on output during the second half of 2007 and the first six weeks of 2008, I contrast teams with suppliers with above-average versus below-average time spent in mixed teams during the first half of 2007. Because most workers at the farm had already spent significant time working with non-coethnics before 2007, columns 3 and 4 restrict the sample to those with below-average tenure. The results show no significant effect of time spent working with non-coethnics on the output gap between mixed and homogeneous teams supplied, neither before nor after conflict began. Workers who have interacted more with individuals of other ethnic groups thus appear no less discriminatory in production.

The results in table 12 do not rule out the possibility that complete segregation between the two ethnic groups over time would have a negative influence on attitudes or behavior towards non-coethnics, however. Carrell, Sacerdote, and West (2011) find that implementing an estimated optimal assignment can have unintended consequences due to unforeseen responses on the part of individuals to out-of-sample assignments. In the context of the sample farm, in a country that has experienced periodical violent clashes between ethnic groups, and where workers of different ethnic groups reside in the same quarters, complete segregation at the plant could for example lead to increased social tensions on the farm.

7.3 Firm’s response to distortions due to ethnic diversity

Nevertheless, it is arguably surprising that a supposedly profit-maximizing firm chose to leave large productivity gains “on the table” by not segregating workers of different ethnicities. Ethical considerations add complexity to the issue of team assignment in Kenya, but we would perhaps expect longer-term costs of segregation to be incurred primarily by society, rather than the firm itself, in which case a case can be made for government intervention to enforce integration within firms.

Becker (1957) pointed out that discriminatory employers should go out of business as their profits suffer. A priori, the same argument should hold for flower farms that allow workplace discrimination to influence productivity. However, the floriculture business is not particularly competitive, as evidenced by high profit margins (Noury, 2011).³⁶ Moreover, as the literature in macroeconomics on across-firm misallocation has highlighted, it is not

³⁶Entry- and exit- barriers are significant: large areas of land and expensive, complicated equipment is needed to produce roses and other cut flowers.

necessarily the most productive firms that survive in poor countries' economies (Banerjee and Moll, 2009; Hsieh and Klenow, 2009).³⁷

Further, plant managers *did* respond to the increase in distortionary discrimination when conflict began, as we have seen. The introduction of team pay for processors was likely motivated by the decrease in productivity in diverse teams in early 2008. It is unsurprising that the dramatic differential decrease in mixed teams' output when conflict began led managers to respond, even though the lower output observed in diverse teams during the first year of the sample period did not. A doubling of the output gap of diverse teams during a short period of time is likely more salient to managers than potential foregone productivity gains from arbitrary assignment to teams.

It appears that managers considered an adjustment to contractual incentives a more desirable response to distortionary discrimination than segregating workers. But note that it is likely not possible to eliminate discrimination through contractual incentives, without entirely breaking the link between workers' output and pay. At the sample plant, vertical discrimination continued to significantly affect output after the introduction of team pay.

8. Conclusion

Evidence suggests that ethnic diversity negatively affects public goods provision and the quality of macroeconomic policies. While the possibility of an additional, direct effect on micro-level productivity has long been recognized, corresponding evidence is largely absent. In this paper, I begin by identifying a sizable, negative productivity effect of ethnic diversity in teams in Kenya. I do so using two years of daily output data for 924 workers, almost equally drawn from two rival tribes, at a flower-packing plant. The packing process takes place in triangular production units, one upstream "supplier" supplying two downstream "processors" who finalize bunches of flowers. I show that an arbitrary position rotation system led to quasi-random variation in teams' ethnicity configuration. As predicted by a model in which different weight is attached to coethnic and non-coethnic downstream workers' utility, suppliers discriminate both "vertically" - undersupplying downstream non-coethnics - and "horizontally" - shifting flowers from non-coethnic to coethnics downstream workers. By doing so, upstream workers lower their own pay and total output.

I show that less distortionary, non-taste-based ethnic diversity effects are unlikely to explain this paper's results. As Becker points out, significant aggregate effects "could easily result from the manner in which individual tastes for discrimination allocate resources within a free-enterprise framework" (Becker, 1957, p. 30). Discrimination should lead to misallocation of resources in most joint production situations in which individuals influence the output and income of others. I take advantage of two natural experiments during the time period observed to begin to explore how the productivity effects of ethnic diversity are likely to vary across time and space. When contentious presidential election results led to political conflict and violent clashes between the two ethnic groups represented in the sample in early 2008, a dramatic, differential decrease in the output of mixed teams followed, as predicted by the

³⁷And even national chain stores in the U.S. with access to large amounts of electronic data and analysis appear to forego profit by not assigning workers to teams optimally (Mas and Moretti, 2009).

model. The reason appears to be that workers' taste for discrimination against non-coethnic co-workers increased. I estimate a decrease in the weight attached to non-coethnics' utility of approximately 35 percent in early 2008, through a reduced form approach. A back-of-the-envelope calculation suggests that the increase in distortionary workplace discrimination may have cost the plant half a million dollars in annual profit, had it not responded.

Six weeks into the conflict period, the plant implemented a new pay system in which downstream workers were paid for their combined output ("team pay"). Under team pay, biased upstream workers are unable to increase the relative pay of favored downstream workers by distorting relative supply. As a result, horizontal misallocation of flowers was eliminated. Total output in teams in which the two processors were of different ethnic groups therefore increased, the introduction of team pay returning the difference in output between such teams and homogeneous teams to pre-conflict levels. Overall output also increased, even though the results indicate that team pay led processors to freeride on each others' effort.

This paper's results indicate that, if taste for discrimination is high enough, firms are forced to adopt "second best" policies to limit the distortions caused by such discrimination. But entirely removing workers' incentives for discrimination is difficult. At the plant, team pay had little effect on the degree of discrimination in teams that were ethnically differentiated vertically rather than horizontally, as also predicted by the model. The obvious "solution" to discrimination - segregating workers - may be undesirable for reasons unrelated to productivity in the short term. The extent and multiplier effects of taste-based misallocation also depend on a number of other factors, such as pay systems, the structure of production, and the "geographical" distribution of ethnic groups in the productive system, however. More speculatively, it is possible that such factors respond endogenously to ethnic diversity. Social segregation is commonly observed in diverse societies but likely becomes harder to achieve as urbanization brings larger groups of workers together. The linkages and specialization required in industrialized production are rarely observed in the most ethnically diverse countries.

My findings also suggest that the economic costs of ethnic diversity vary with the political environment. Relatively brief episodes of ethnic conflict can have a long-lasting impact on economically distortionary attitudes: I find no decay in discrimination in the nine months after conflict ended. Multiple equilibria may thus exist if the occurrence of conflict itself depends on attitudes towards non-coethnics, some diverse societies being characterized by tolerance and little conflict and others by ethnic biases and frequent conflict.

Figures

Figure 1a: Organization of team production

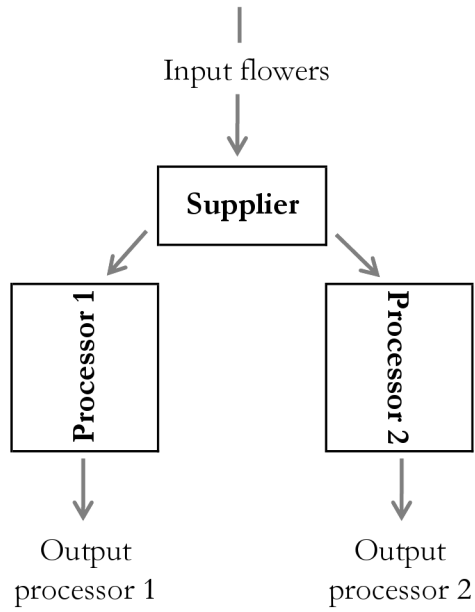


Figure 1b: Team ethnicity configuration categories




Figure 2: Output in homogeneous and mixed teams across time




Figure 3: Model predictions

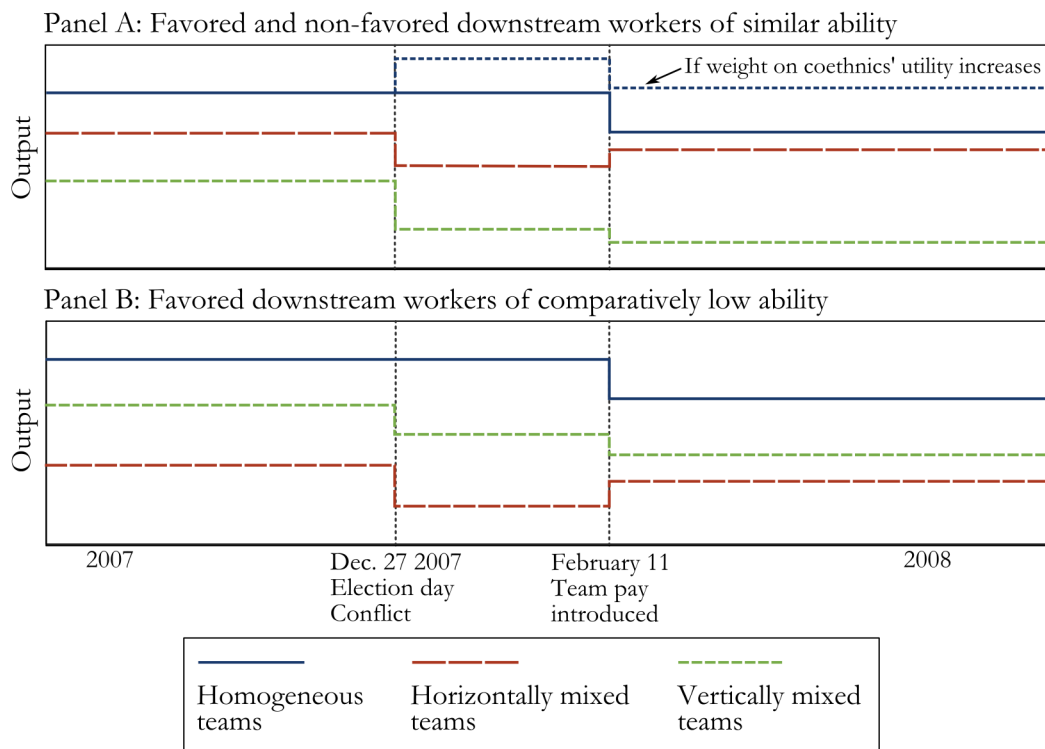

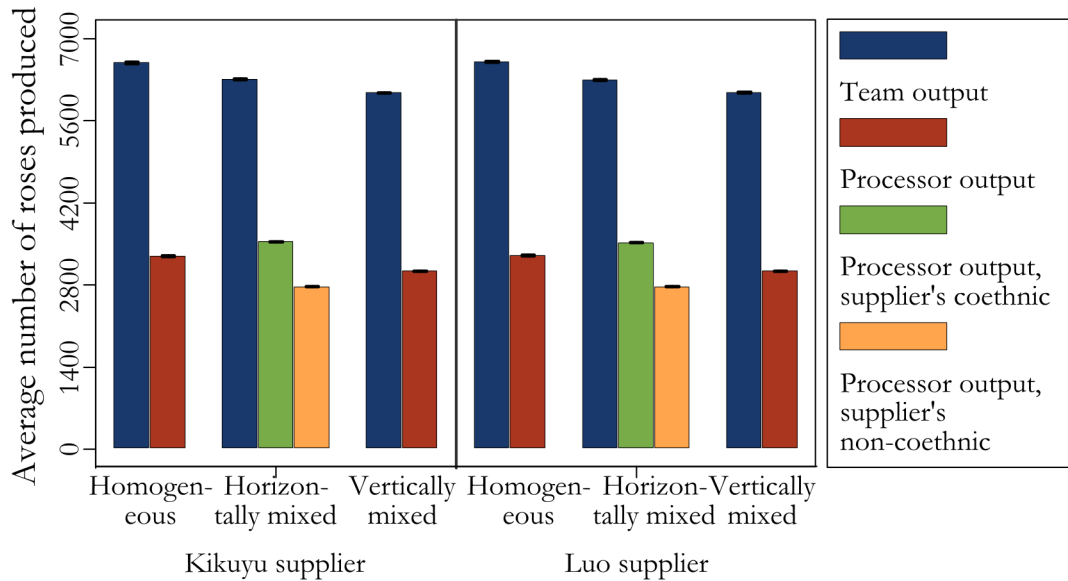


Figure 4: Investigating the shape of the production function



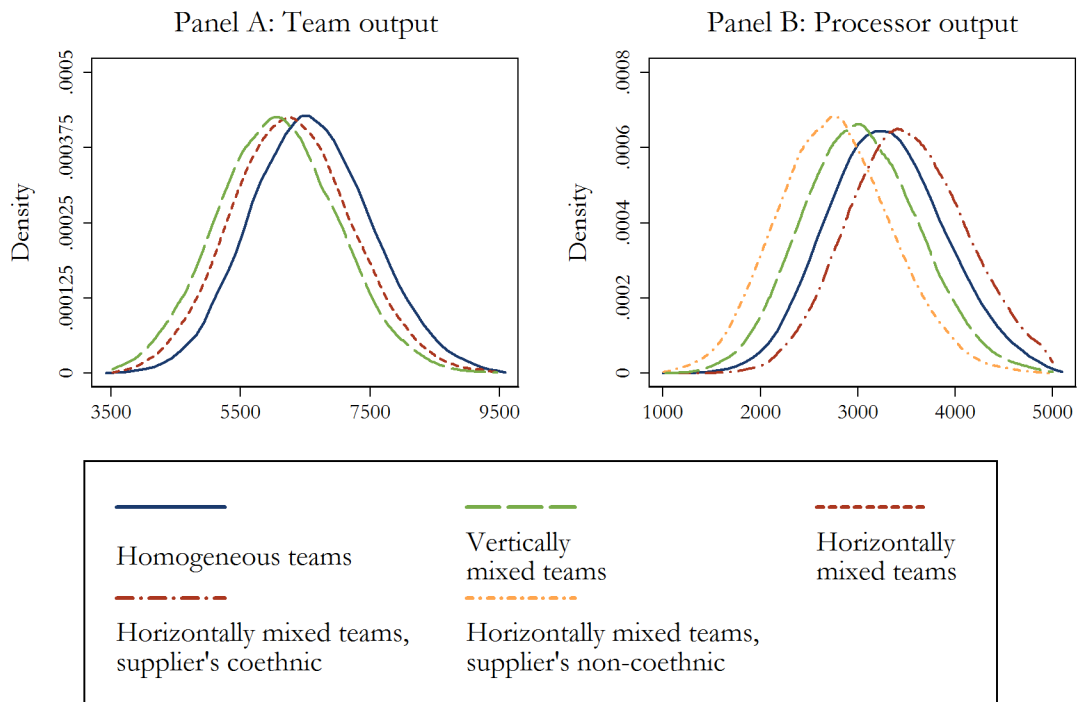
Data from 2007. Outliers (top and bottom percentile) excluded. Local polynomial plots, bandwidth = 350. The processor productivity FE is normalized to have the mean and standard deviation of processor output, and the supplier productivity FE the mean and standard deviation of team output.

Figure 5: Output by team ethnicity configuration



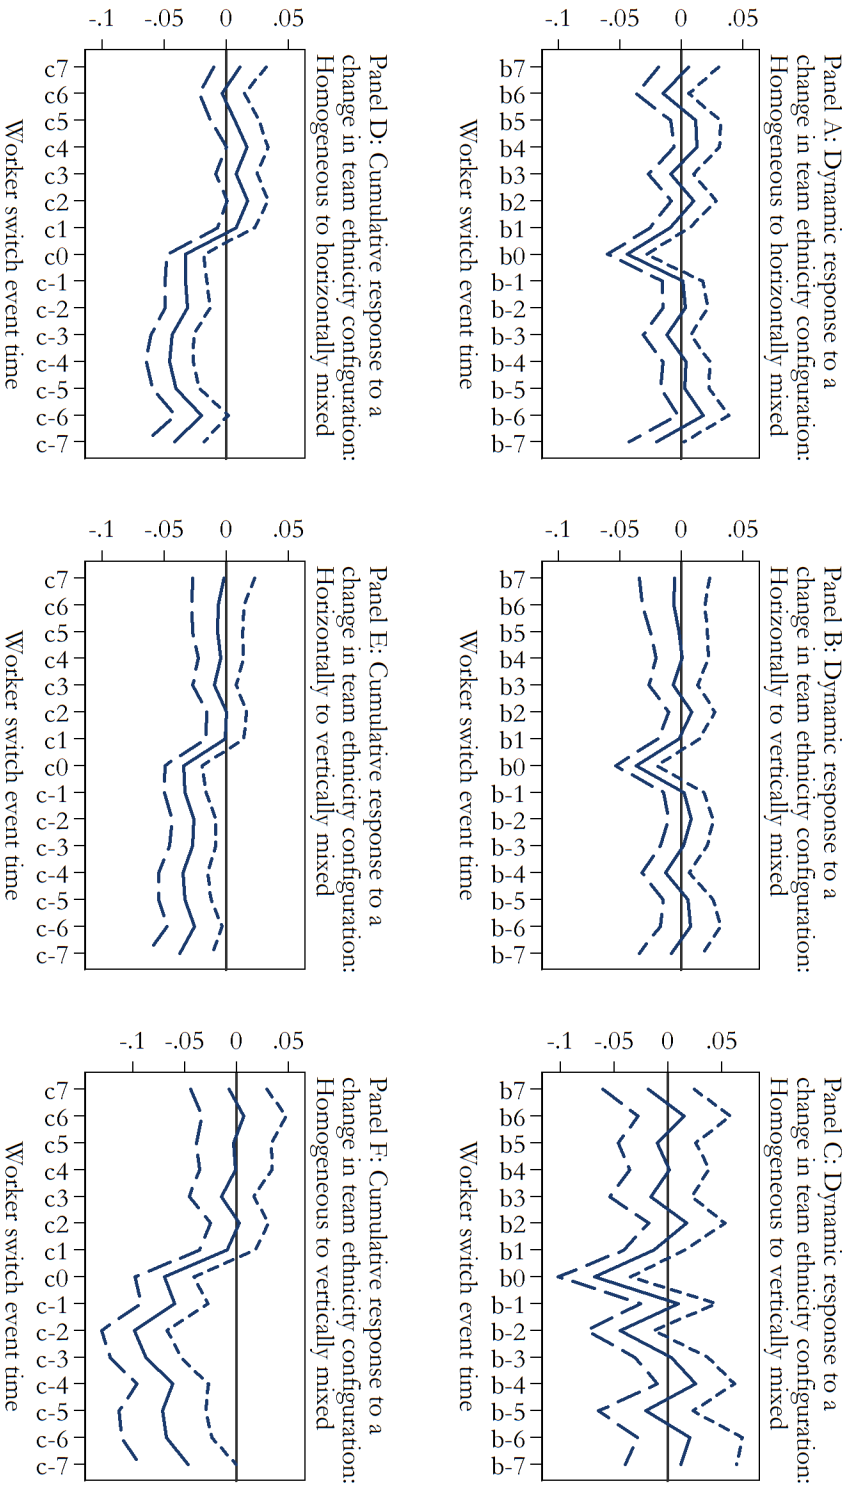
95% confidence intervals are depicted. In teams with Kikuyu suppliers, average output in teams of different ethnicity configurations is as follows (standard errors in parenthesis). Team output in homogeneous teams: 6586 (12). Processor output in homogeneous teams: 3295 (8). Team output in horizontally mixed teams: 6307 (9). Processor output in horizontally mixed teams, supplier's coethnic: 3539 (8). Processor output in horizontally mixed teams, supplier's non-coethnic: 2777 (7). Team output in vertically mixed teams: 6073 (11). Processor output in vertically mixed teams: 3039 (7). In teams with Luo suppliers, average output in teams of different ethnicity configurations is as follows (standard errors in parenthesis). Team output in homogeneous teams: 6606 (12). Processor output in homogeneous teams: 3304 (6). Team output in horizontally mixed teams: 6290 (7). Processor output in horizontally mixed teams, supplier's coethnic: 3521 (7). Processor output in horizontally mixed teams, supplier's non-coethnic: 2762 (7). Team output in vertically mixed teams: 6075 (12). Processor output in vertically mixed teams: 3037 (7)

Figure 6: Distribution of output by team ethnicity configuration




Data from 2007

Figure 7: Team output responses to changes in team ethnicity configuration




Data from 2007. In panels A - C we plot the estimated coefficients from a regression of the first difference (across days) in output on an indicator for a worker switch entailing a change in team ethnicity configuration and its lead and lag terms (the other two workers on the team are unchanged). For example, b1 is the coefficient on the 7th lead term. In panels D - F we plot the cumulative response. The dotted lines represent 95 percent confidence intervals.

Figure 8: Output by team ethnicity configuration
Before and after conflict, and under team pay



95% confidence intervals are depicted but narrow and thus hard to see. 'Conflict' signifies the first 6 weeks of 2008 when ethnically-based violence was taking place but processors were still paid individual piece rates. 'Team pay' signifies the remainder of 2008, after team pay for processors was introduced

Figure 9: Heterogeneity in response to conflict



An observation is the output differential of a given supplier across homogeneous and mixed teams

Tables

Table 1
Sample summary statistics

	Whole sample (N=924)	Kikuyu (N=426)	Luo (N=498)
Ethnicity (% Kikuyu)	0.46 (0.50)		
Gender (% female)	0.59 (0.49)	0.57 (0.49)	0.61 (0.49)
Age (average age)	34.63 (5.21)	34.55 (5.15)	34.82 (5.29)
Experience (average years of tenure)	5.49 (1.48)	5.62 (1.38)	5.38 (1.51)

Standard deviations in parentheses. Individuals of the Kikuyu, Embu, Meru, Kamba, Maasai and Kisii tribes are considered “Kikuyu” and those of the Luo, Luhya and Kalenjin tribes “Luo”.

Table 2 (continued)

		Processor 2								
		0,0,0	0,0,1	0,1,0	0,1,1	1,0,0	1,0,1	1,1,0	1,1,1	TOTAL
P r o c e s s o r 1	0,0,0	0.013 (0.011)	0.010 (0.011)	0.018 (0.017)	0.017 (0.020)	0.012 (0.012)	0.012 (0.010)	0.015 (0.016)	0.014 (0.014)	0.110
	0,0,1	0.009 (0.011)	0.010 (0.011)	0.018 (0.017)	0.023 (0.020)	0.011 (0.012)	0.010 (0.010)	0.014 (0.016)	0.013 (0.013)	0.109
	0,1,0	0.014 (0.014)	0.015 (0.015)	0.025 (0.022)	0.022 (0.026)	0.015 (0.015)	0.013 (0.014)	0.024 (0.021)	0.018 (0.018)	0.144
	0,1,1	0.015 (0.017)	0.021 (0.017)	0.026 (0.026)	0.029 (0.031)	0.017 (0.018)	0.016 (0.016)	0.025 (0.025)	0.021 (0.021)	0.171
	1,0,0	0.011 (0.011)	0.013 (0.011)	0.015 (0.016)	0.019 (0.019)	0.011 (0.011)	0.010 (0.010)	0.015 (0.015)	0.013 (0.013)	0.107
	1,0,1	0.009 (0.009)	0.010 (0.009)	0.015 (0.014)	0.018 (0.017)	0.008 (0.010)	0.008 (0.009)	0.014 (0.013)	0.012 (0.011)	0.093
	1,1,0	0.017 (0.015)	0.011 (0.015)	0.021 (0.023)	0.027 (0.027)	0.020 (0.016)	0.015 (0.014)	0.022 (0.022)	0.018 (0.019)	0.152
	1,1,1	0.012 (0.011)	0.011 (0.011)	0.015 (0.017)	0.023 (0.020)	0.013 (0.012)	0.010 (0.011)	0.016 (0.016)	0.014 (0.014)	0.114
	TOTAL	0.100	0.100	0.152	0.179	0.107	0.094	0.145	0.123	
	p-values:		Whole sample period			Pre-conflict	Conflict	Team pay		
		0.27			0.24	0.24	0.50			

The top number in cell i, j is the observed proportion of position i / position j pairs in which the worker in position i has the 2^3 characteristics listed in row i and the worker in position j the 2^3 characteristics listed in column j . The bottom number is the expected proportion under the null hypothesis of independence. p-values for Pearson's chi-square statistic are shown. Because the worker rotation system leads to complex temporal correlation in team compositions and output, the assumptions required for validity of the chi-square tests would be violated if all data was used. I thus use a periodical "snapshot" of data in this table: team compositions on the first day of every month (team spells do not exceed one month). The chi-square tests are insignificant if data from other dates is used instead. Supplier - Processor 2 is not shown because the two processor positions are "interchangeable" A worker's productivity is her average output in month $t - 2$.

Table 3
Model predictions to be tested

Proposition (see theoretical appendix)	Model predictions when upstream workers have		Predictions of models of non-taste-based diversity effects, when different*
	Discriminatory preferences	Neutral preferences	
2 Output of individuals in teams of different ethnicity configurations	$q_{HM,C} > q_H > q_{VM} > q_{HM,NC}$	$q_{HM,C} = q_H = q_{VM} = q_{HM,NC}$	
3 Output of teams of different ethnicity configurations	$Q_H > Q_{VM}$ and $Q_H > Q_{HM}$	$Q_H = Q_{HM} = Q_{VM}$	
4 Differential effect of upstream capacity across teams of different ethnicity configurations	$\partial q_{HM,C}/\partial \alpha_s > \partial q_H/\partial \alpha_s > \partial q_{VM}/\partial \alpha_s > \partial q_{HM,NC}/\partial \alpha_s$	$\partial q_{HM,C}/\partial \alpha_s = \partial q_H/\partial \alpha_s = \partial q_{VM}/\partial \alpha_s = \partial q_{HM,NC}/\partial \alpha_s$	
5 Effects of change in attitudes towards individuals of other ethnic groups	$\partial q_{HM,C}/\partial \theta_{NC} > 0 = q_H/\partial \theta_{NC} > \partial q_{VM}/\partial \theta_{NC} > \partial q_{HM,NC}/\partial \theta_{NC}$	N/A	No effect
6 Effects of group-based pay for downstream workers	$Q_H^{TP} < Q_H$ $Q_{VM}^{TP} < Q_{VM}$ $Q_H^{TP} > Q_{VM}^{TP}$ $q_{HM,C}^{TP} = q_{HM,NC}^{TP}$ $Q_{HM}^{TP} \geq Q_{HM}$	$Q_H^{TP} < Q_H$ $Q_{VM}^{TP} < Q_{VM}$ $Q_H^{TP} = Q_{VM}^{TP}$ $q_{HM,C}^{TP} = q_{HM,NC}^{TP}$ $Q_{HM}^{TP} < Q_{HM}$	$q_{HM,C}^{TP} > q_{HM,NC}^{TP}$ $Q_{HM}^{TP} < Q_{HM}$

* I refer here to models of technological, informational, or “cooperational” ethnic diversity effects as typically specified in the literature. The predictions of more intricate non-taste-based models may differ.

Table 4
Output by team ethnicity configuration

Dependent variable:	Processor output		Team output		Log (Processor output)		Log (Team output)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Constant	3298.970*** (3.422)	3364.957*** (88.248)	6597.940*** (7.504)	6729.914*** (174.746)	8.084*** (0.001)	8.105*** (0.029)	8.784*** (0.001)	8.806*** (0.028)
Horizontally mixed			-300.192*** (9.193)	-296.986*** (8.351)			-0.048*** (0.001)	-0.047*** (0.001)
Horizontally mixed, processor of supplier's ethnicity	229.705*** (5.057)	231.308*** (4.834)			0.069*** (0.002)	0.069*** (0.001)		
Horizontally mixed, processor not of supplier's ethnicity	-529.896*** (4.902)	-528.293*** (4.744)			-0.182*** (0.002)	-0.182*** (0.002)		
Vertically mixed	-261.902*** (4.929)	-260.200*** (4.767)	-523.803*** (11.027)	-520.400*** (9.910)	-0.086*** (0.002)	-0.085*** (0.002)	-0.085*** (0.002)	-0.084*** (0.002)
Omitted category								
Individual fixed effects?	NO	YES	NO	YES	NO	YES	NO	YES
N	199810	199810	99905	99905	199810	199810	99905	99905

Data from 2007 is used in these OLS regressions. The outcome variables are de-seasonalized, daily output quantities. The standard errors are clustered at the team level in columns 3, 4, 7 and 8, and at the processor \times team level in columns 1, 2, 5, and 6. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5
Supplier ability effect by team ethnicity configuration

Dependent variable:	Processor output	
	(1)	(2)
Supplier permanent productivity	0.05*** (0.00)	0.07*** (0.00)
Supplier permanent productivity × Horizontally mixed, processor of supplier's ethnicity	-0.01 (0.00)	0.00 (0.00)
Supplier permanent productivity × Horizontally mixed, processor not of supplier's ethnicity	-0.02*** (0.00)	-0.04*** (0.00)
Supplier permanent productivity × Vertically mixed	-0.02*** (0.00)	-0.03*** (0.00)
Constant	2996.41*** (21.38)	2619.83*** (24.31)
Horizontally mixed, processor of supplier's ethnicity	263.12*** (31.03)	216.49*** (28.83)
Horizontally mixed, processor not of supplier's ethnicity	-389.63*** (30.80)	-265.38*** (28.89)
Vertically mixed	-154.99*** (31.11)	-83.73*** (28.98)
Omitted category	Homogeneous team / Processor in homogeneous team	
Controls for processor permanent productivity?	NO	YES
	N 197058	197058

Data from 2007 is used in these OLS regressions. The outcome variables are de-seasonalized, daily output quantities. Permanent productivity was estimated as described in section 4. Processor permanent productivity was normalized to have the mean and sd of processor output. Supplier permanent productivity was normalized to have the mean and sd of team output. The sample sizes are slightly reduced because the 5 workers that were observed as a supplier in less than two teams do not have an estimated fixed effect as supplier (and vv for those observed as processors in less than two teams). The standard errors are clustered at the processor × team level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The standard errors should be adjusted for the productivity regressors being estimated, and possible downward bias in the coefficients on the estimated fixed effects corrected for e.g. by using the Fuller estimator, but neither of those adjustments are likely to influence the statistical significance of estimates as precise as those seen here.

Table 6
 Testing for ethnic diversity effects using pair fixed effects and 3rd worker switches
 Sample: pairs of workers observed before and after the 3rd worker is replaced by another
 worker of the same permanent productivity tercile

Change in team ethnicity configuration	Worker switched		Log (Unswitched)		Log (Team output)	
	Processor output	Team output	Processor output	Team output	Processor output	Team output
No change	1.805 (11.208)	-5.094 (9.142)	0.000 (0.004)	-0.000 (0.001)		
Homogeneous to horizontally mixed		-308.462*** (18.210)		-0.050*** (0.003)		
Homogeneous to horizontally mixed, processor of supplier's ethnicity	Other P (23.084)		0.054*** (0.007)			
Homogeneous to vertically mixed	Supplier (33.505)	-525.040*** (36.577)	-0.094*** (0.011)	-0.085*** (0.006)		
Horizontally mixed to vertically mixed	Coethnic P	-209.574*** (19.181)		-0.035*** (0.003)		
Horizontally mixed, processor of supplier's ethnicity to horizontally mixed, processor not of supplier's ethnicity	Supplier	-746.938*** (36.759)	-0.245*** (0.012)			
Horizontally mixed, processor not of supplier's ethnicity to vertically mixed	Coethnic P	245.388*** (24.361)	0.091*** (0.008)			
Constant		3212.938*** (4.551)	8.052*** (0.002)	8.749*** (0.001)		
Omitted category		Pairs of workers "before" the 3rd worker switch (see notes)				
Pair fixed effect for unswitched workers?		YES	YES	YES	YES	
N		109998	51589	109998	51589	

P = Processor, S = Supplier, Coethnic = of supplier's ethnicity. Data from 2007 is used in these OLS regressions. The outcome variables are de-seasonalized output quantities averaged within teams (where a team is defined by a specific worker as supplier and two other specific workers as processors. t thus refers to a team spell below, not a specific day). In (1) and (3) the output of the processor *not being switched* is used. The standard errors are clustered at the pair level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The fixed effects are for the pair of workers held constant. Note that regressors for all types of switches are included. Suppose that pair A goes from being in a team of ethnicity configuration X at $t - 1$ to Y at t , and pair B from being in a Y team at $t - 1$ to an X team at t . Rather than include separate "X to Y" and "Y to X" dummies, the regressions in this table include only "X to Y" and turn the associated dummy on at t for pair A and a $t - 1$ for pair B. For both pairs the coefficient on "X to Y" reflects the difference in output between a period when the pair is in a team of configuration Y and a contiguous period when the pair is in a team of configuration X. For switches involving no change in ethnicity configuration, the dummy is (arbitrarily) turned on at the later of the two periods. Note that the constant reflects a weighted average of "pre" change output in teams of different ethnicity configurations. The non-standard definition of regressors in the table was used so as to ease comparison of the results with those in table 3.

Table 7
Output by tribe-specific team ethnicity configuration

Sample:	Kikuyu - Luo		Kikuyu - Luhya		Luo - Luhya	
	Processor output	Team output	Processor output	Team output	Processor output	Team output
Constant	(1) 3287.80*** (5.49)	(2) 6575.61*** (11.97)	(3) 3294.43*** (6.37)	(4) 6588.85*** (13.88)	(5) 3284.05*** (9.98)	(6) 6589.25*** (17.75)
Horizontally mixed		-304.85*** (15.03)		-255.12*** (20.19)		21.37 (32.84)
Horizontally mixed, processor of supplier's ethnicity	233.84*** (8.43)		251.50*** (12.02)		31.75* (16.67)	
Horizontally mixed, processor not of supplier's ethnicity	-538.69*** (8.13)		-506.62*** (11.27)		17.36 (16.28)	
Vertically mixed	-274.41*** (7.94)	-548.83*** (17.64)	-242.00*** (11.26)	-484.01*** (24.83)	19.60 (16.08)	18.06 (34.49)
Omitted category		Homogeneous team	Processor /	Homogeneous team		
N	72188	36094	34986	17493	16876	8438

Data from 2007 is used in these OLS regressions. The outcome variables are de-seasonalized, daily output quantities. The standard errors are clustered at the team level in columns 2, 4, and 6 and at the processor \times team level in columns 1, 3, and 5. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. In this paper, Luo and Luhya workers are categorized as belonging to the Luo tribal bloc and Kikuyu workers to the Kikuyu bloc.

Table 8
Output by team ethnicity configuration before and after conflict

Dependent variable:	Processor output		Team Output		Log (Processor output)		Log (Team output)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Constant	3298.970*** (3.422)	3580.414*** (92.667)	6597.94*** (7.504)	6973.301*** (220.060)	8.084*** (0.001)	8.164*** (0.029)	8.784*** (0.001)	8.844*** (0.036)
Horizontally mixed			-300.192*** (9.193)	-297.777*** (8.341)			-0.048*** (0.001)	-0.047*** (0.001)
Horizontally mixed, processor of supplier's ethnicity	229.705*** (5.057)	230.912*** (4.822)			0.069*** (0.002)	0.069*** (0.001)		
Horizontally mixed, processor not of supplier's ethnicity	-529.896*** (4.902)	-528.689*** (4.735)			-0.182*** (0.002)	-0.182*** (0.002)		
Vertically mixed	-261.902*** (4.929)	-261.551*** (4.747)	-523.803*** (11.027)	-523.101*** (9.899)	-0.086*** (0.002)	-0.086*** (0.002)	-0.085*** (0.002)	-0.085*** (0.002)
Conflict	-3.469 (9.973)	3.329 (9.252)	-6.937 (22.624)	6.658 (19.536)	-0.001 (0.004)	0.001 (0.003)	-0.001 (0.003)	0.001 (0.003)
Horizontally mixed × Conflict			-251.067*** (27.960)	-272.997*** (24.820)			-0.042*** (0.004)	-0.046*** (0.004)
Horizontally mixed, processor of supplier's ethnicity × Conflict	74.334*** (14.934)	63.369*** (13.749)			0.022*** (0.005)	0.019*** (0.004)		
Horizontally mixed, processor not of supplier's ethnicity × Conflict	-325.319*** (14.360)	-336.366*** (13.963)			-0.133*** (0.005)	-0.137*** (0.005)		
Vertically mixed × Conflict	-218.163*** (14.138)	-219.439*** (13.438)	-436.326*** (31.737)	-438.879*** (28.105)	-0.077*** (0.005)	-0.077*** (0.004)	-0.076*** (0.005)	-0.076*** (0.005)
Omitted category								
Individual fixed effects?	NO	YES	NO	YES	NO	YES	NO	YES
N	224730	224730	112365	112365	224730	224730	112365	112365

Data from 2007 and the first 6 weeks of 2008 is used in these OLS regressions. The outcome variables are de-seasonalized, daily output quantities. The standard errors are clustered at the team level in columns 3, 4, 7 and 8, and at the processor × team level in columns 1, 2, 5, and 6. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9
Output by team ethnicity configuration before and after team pay

Dependent variable:	Processor output		Team output		Log (Processor output)		Log (Team output)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Constant	3295.502*** (9.391)	3325.549*** (80.633)	6591.003*** (21.388)	6651.097*** (160.105)	8.083*** (0.003)	8.085*** (0.029)	8.783*** (0.003)	8.790*** (0.028)
Horizontally mixed			-551.257*** (26.458)	-568.948*** (23.745)			-0.090*** (0.004)	-0.093*** (0.004)
Horizontally mixed, processor of supplier's ethnicity	304.039*** (14.088)	295.194*** (13.058)			0.090*** (0.004)	0.087*** (0.004)		
Horizontally mixed, processor not of supplier's ethnicity	-855.296*** (13.524)	-864.142*** (13.130)			-0.316*** (0.005)	-0.319*** (0.005)		
Vertically mixed	-480.064*** (13.277)	-486.560*** (12.636)	-960.129*** (29.804)	-973.121*** (26.616)	-0.163*** (0.004)	-0.165*** (0.004)	-0.161*** (0.005)	-0.163*** (0.004)
Team pay	-30.283*** (9.989)	-32.151*** (9.346)	-60.566*** (22.804)	-64.302*** (20.059)	-0.010*** (0.003)	-0.010*** (0.003)	-0.010*** (0.004)	-0.011*** (0.003)
Horizontally mixed × Team pay			253.850*** (28.200)	260.498*** (25.287)			0.043*** (0.004)	0.044*** (0.004)
Horizontally mixed, processor of supplier's ethnicity × Team pay	-425.545*** (14.912)	-422.222*** (13.851)			-0.130*** (0.004)	-0.128*** (0.004)		
Horizontally mixed, processor not of supplier's ethnicity × Team pay	679.396*** (14.407)	682.720*** (13.913)			0.258*** (0.005)	0.260*** (0.005)		
Vertically mixed × Team pay	1.351 (14.110)	1.112 (13.466)	2.701 (31.811)	2.224 (28.522)	-0.002 (0.005)	-0.001 (0.005)	-0.001 (0.005)	-0.001 (0.005)
Omitted category								
Individual fixed effects?	NO	YES	NO	YES	NO	YES	NO	YES
N	204148	204148	102074	102074	204148	204148	102074	102074

Data from 2008 is used in these OLS regressions. The outcome variables are de-seasonalized, daily output quantities. The standard errors are clustered at the team level in columns 3, 4, 7 and 8, and at the processor × team level in columns 1, 2, 5, and 6. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 10
Heterogeneity in effect of conflict on discriminatory behavior

Dependent variable:	Supplier lost	Supplier did not	Supplier	Supplier
	relative	lose relative	young	old
Sample:	Team output			
	(1)	(2)	(3)	(4)
Constant	6564.04*** (16.03)	6608.01*** (8.48)	6593.77*** (10.78)	6602.07*** (10.45)
Horizontally mixed	-286.25*** (19.50)	-304.21*** (10.42)	-263.82*** (13.01)	-336.66*** (12.97)
Vertically mixed	-474.25*** (23.57)	-538.01*** (12.47)	-390.33*** (15.05)	-660.79*** (15.21)
Conflict	61.26 (47.55)	-26.00 (25.64)	1.42 (32.67)	-15.18 (31.33)
Horizontally mixed × Conflict	-329.46*** (59.27)	-230.85*** (31.63)	-279.78*** (39.49)	-223.44*** (39.65)
Vertically mixed × Conflict	-584.96*** (77.47)	-402.41*** (35.07)	-567.77*** (43.71)	-301.22*** (45.62)
Omitted category	Homogeneous teams			
N	24860	87505	56432	55933

Data from 2007 and the first six weeks of 2008 is used in these OLS regressions. The outcome variables are de-seasonalized, daily output quantities. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The standard errors are clustered at the team level.

Table 11
Output gains from optimal assignment by ethnicity, productivity or both

Period: No conflict						
Output-maximizing assignment by:	Ethnicity		Productivity as P and S		Ethnicity and productivity as P and S	
Assignment	Homogeneous	100 %	s2p3p3	50.64%	Homogeneous,s2p1p2	17.53%
			s3p1p2	49.03%	Homogeneous,s3p1p2	31.82%
Output gains relative to:			s3p1p3	0.32%	Homogeneous,s3p3p3	50.65%
random assignment	4.4%			3.54%		9.19%
output-minimizing assignment	8.6%			7.13%		16.38%
Period: Conflict						
Output-maximizing assignment by:	Ethnicity		Productivity as P and S		Ethnicity and productivity as P and S	
Assignment	Homogeneous	100 %	s2p2p3	0.65%	Homogeneous,s2p1p2	28.90%
			s2p3p3	50.32%	Homogeneous,s3p1p1	18.18%
			s3p1p1	34.32%	Homogeneous,s3p2p3	4.55%
Output gains relative to:			s3p2p2	14.61%	Homogeneous,s3p3p3	48.38%
random assignment	8.2%			4.33%		15.37%
output-minimizing assignment	17 %			8.47%		30.10%
Period: Team pay						
Optimal assignment by:	Ethnicity		Productivity as P and S		Ethnicity and productivity as P and S	
Assignment	Homogeneous	100 %	s2p1p3	0.32%	Homogeneous,s2p3p3	44.48%
			s2p3p3	50.65%	Homogeneous,s3p1p1	34.42%
			s3p1p1	19.16%	Homogeneous,s3p2p2	14.94%
Output-maximizing assignment by:			s3p2p2	29.87%	Homogeneous,s3p3p3	6.17%
random assignment	6.4%			3.42%		10.65%
output-minimizing assignment	17 %			7.13%		24.85%

sX = supplier productivity of tercile X. pX analogous (only productivity tercile in assigned position is shown). The team type configuration that the average output associated with all types of teams and the “budget set” of workers available suggests will maximize output is displayed. The procedure is described in the empirical appendix.

Table 12
Effect of previous interaction with workers of other ethnic groups
on discriminatory behavior before and during conflict

Sample, tenure:	All		Below median tenure	
Sample, previous interaction	Supplier interaction w/ non-coethnics low	Supplier interaction w/ non-coethnics high	Supplier interaction w/ non-coethnics low	Supplier interaction w/ non-coethnics high
Dependent variable:	Team output			
	(1)	(2)	(3)	(4)
Constant	6612.00*** (14.66)	6579.56*** (14.57)	6575.20*** (19.20)	6595.34*** (21.91)
Horizontally mixed	-307.34*** (18.12)	-275.16*** (17.86)	-312.52*** (24.28)	-294.89*** (27.00)
Vertically mixed	-536.23*** (22.46)	-504.28*** (20.45)	-536.14*** (29.01)	-492.12*** (31.19)
Conflict	-20.81 (33.73)	11.24 (33.20)	0.03 (43.19)	-27.91 (51.55)
Horizontally mixed × Conflict	-258.85*** (41.34)	-263.32*** (41.23)	-263.39*** (53.96)	-270.99*** (65.68)
Vertically mixed × Conflict	-437.18*** (48.82)	-446.72*** (45.58)	-442.10*** (69.34)	-446.37*** (67.83)
Omitted category	Homogeneous teams			
N	29634	32642	15132	13582

Data from 2007 and the first six weeks of 2008 is used in these OLS regressions. The outcome variables are de-seasonalized, daily output quantities. The standard errors are clustered at the team level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix

A1. Theoretical Appendix

In addition to the assumptions in section 3, I make the following assumptions. Let $q_p = f(e_{sp}, \alpha_s, e_p, \alpha_p) = (e_p \alpha_p)^\beta (e_{sp} \alpha_s)^\gamma$. β then measures the slope of processor output in processor ability and effort, and γ the slope in supplier ability and effort. The ability terms are assumed to be positive, and q_p concave in processor and supplier effort. q_p is also assumed to display decreasing returns: $0 < \beta < 1$, $0 < \gamma < 1$, $\beta + \gamma < 1$. The processor's effort carries costs $\frac{1}{2}e_1^2$, and the total effort of the supplier $\frac{1}{2}(e_{s1} + e_{s2})^2$. I assume that $\alpha_p > 1$, $\alpha_s > 1$ and $-\frac{1}{2} < \theta_p < \frac{1}{2}$.³⁸ I also assume that suppliers do not take ethnicity as a signal of ability.

Consider first the processor's problem, focusing here on processor 1 (processor 2's problem is analogous). A processor maximizes her benefit of pay minus her cost of effort:

$$\begin{aligned} \text{Max}_{e_1} \quad & 2w(e_1 \alpha_1)^\beta (e_{s1} \alpha_s)^\gamma - \frac{1}{2}e_1^2 \\ \text{s.t.} \quad & e_1 \geq 0 \end{aligned}$$

which gives

$$e_1 = \left(2w\beta(e_{s1} \alpha_s)^\gamma \alpha_1^\beta \right)^{\frac{1}{2-\beta}}$$

Processor effort is thus increasing in processor and supplier ability and in the supplier's effort. Note that the processor's effort choice depends on the supplier's weight on her utility only through its influence on her supply of intermediate flowers.

A supplier maximizes her benefit of pay minus her cost of effort plus the additional utility (or disutility) she derives from the well-being of each of the processors:

$$\begin{aligned} \text{Max}_{e_{s1}, e_{s2}} \quad & w((e_1 \alpha_1)^\beta (e_{s1} \alpha_s)^\gamma + (e_2 \alpha_2)^\beta (e_{s2} \alpha_s)^\gamma) - \frac{1}{2}(e_{s1} + e_{s2})^2 \\ & + \theta_1 \left(2w(e_1 \alpha_1)^\beta (e_{s1} \alpha_s)^\gamma - \frac{1}{2}e_1^2 \right) + \theta_2 \left(2w(e_2 \alpha_2)^\beta (e_{s2} \alpha_s)^\gamma - \frac{1}{2}e_2^2 \right) \\ \text{s.t.} \quad & e_{s1} \geq 0 \text{ and } e_{s2} \geq 0 \end{aligned}$$

The supplier's first order condition for e_{s1} gives

$$(e_{s1} + e_{s2}) = (1 + 2\theta_1) w(e_1 \alpha_1)^\beta \gamma (e_{s1} \alpha_s)^{\gamma-1}$$

When the supplier's two first order conditions hold simultaneously,

³⁸If this restriction is violated corner solutions arise.

$$e_{s1} = \left(\frac{(1 + 2\theta_1) w(e_1 \alpha_1)^\beta \gamma (\alpha_s)^{\gamma-1}}{1 + \left(\frac{1+2\theta_1}{1+2\theta_2} \right)^{\frac{1}{\gamma-1}} \left(\frac{e_1 \alpha_1}{e_2 \alpha_2} \right)^{\frac{\beta}{\gamma-1}}} \right)^{\frac{1}{2-\gamma}}$$

Because the supplier considers the pay-off (from own pay and processors' utility) of supply to each of the processors, her effort devoted to supplying processor 1 is increasing in that processor's ability and utility weight, but decreasing in the ability and utility weight of the other processor.

The model has the following predictions. Because tedious algebra is involved, the proofs are in the online theoretical appendix.

Proposition 1 (Existence and comparative statics):

- i. There exists a unique equilibrium in which production is given by

$$q_1^* = \frac{k_q \alpha_s^{\frac{2\gamma}{2-\beta-\gamma}} \alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (1 + 2\theta_1)^{\frac{2\gamma}{2-\beta-2\gamma}}}{\left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} (1 + 2\theta_1)^{\frac{2-\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} (1 + 2\theta_2)^{\frac{2-\beta}{2-2\gamma-\beta}} \right)^{\frac{\gamma}{2-\beta-\gamma}}}$$

$$Q^* = \frac{k_q \alpha_s^{\frac{2\gamma}{2-\beta-\gamma}} \left(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (1 + 2\theta_1)^{\frac{2\gamma}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} (1 + 2\theta_2)^{\frac{2\gamma}{2-\beta-2\gamma}} \right)}{\left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} (1 + 2\theta_1)^{\frac{2-\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} (1 + 2\theta_2)^{\frac{2-\beta}{2-2\gamma-\beta}} \right)^{\frac{\gamma}{2-\beta-\gamma}}}$$

where $k_q = (2\beta)^{\frac{\beta}{2-\gamma-\beta}} w^{\frac{\beta+\gamma}{2-\gamma-\beta}} \gamma^{\frac{\gamma}{2-\gamma-\beta}}$ and $Q = q_1 + q_2$ is team output.

- ii. Processor output is increasing in own ability, the ability of the supplier and the weight the supplier attaches to her utility, but decreasing in the ability and weight of the other processor: $\frac{\partial q_1}{\partial \alpha_1} > 0$, $\frac{\partial q_1}{\partial \alpha_s} > 0$, $\frac{\partial q_1}{\partial \alpha_2} < 0$, $\frac{\partial q_1}{\partial \theta_1} > 0$, $\frac{\partial q_1}{\partial \theta_2} < 0$

In principle the θ 's vary continuously. However, to focus on the possibility of supplier discrimination, I consider a simplified case. Let $\theta_i = \theta_C$ if processor i is of the supplier's ethnic group, and $\theta_i = \theta_{NC}$ if not. Processors are then observed in four different positions: in homogeneous teams (H), in vertically mixed teams (VM), and in horizontally mixed teams in which the processor in question may (HM, C) or may not (HM, NC) be of the supplier's ethnic group. From a team perspective there are three types of ethnicity configurations, as illustrated in figure 1b. I then derive the following comparative propositions.

Proposition 2 (Processor output): Processor output is unaffected by the ethnicity of the supplier and the other processor if the supplier has ethnicity-neutral social preferences

($\theta_C = \theta_{NC}$): $q_H = q_{HM,C} = q_{HM,NC} = q_{VM}$. Processor output is higher (a) when working with a coethnic supplier, and (b) when working with another processor who is not of the supplier's ethnicity if the supplier has discriminatory preferences ($\theta_C > \theta_{NC}$): $q_{HM,C} > q_H > q_{VM} > q_{HM,NC}$

Ethnicity-neutral upstream workers' optimal supply to each processor is determined by the abilities of the three workers. Proposition 2 makes clear that biased supplier preferences will lead to "horizontal misallocation" - the relative supply to the two processors deviating from their relative abilities - in horizontally mixed teams, and to "vertical misallocation" - the total quantity of roses supplied deviating from the ethnicity-neutral optimal supply - in both horizontally and vertically mixed teams. Misallocation of roses is predicted to lower team output:

Proposition 3 (Team output): Team output is unaffected by a team's ethnicity configuration if the supplier has ethnicity-neutral social preferences ($\theta_C = \theta_{NC}$): $Q_H = Q_{HM} = Q_{VM}$. Team output in homogeneous teams is higher than in mixed teams if the supplier has discriminatory preferences ($\theta_C > \theta_{NC}$): $Q_H > Q_{VM}$ and $Q_H > Q_{HM}$

Next I consider the framework's predictions for how upstream capacity is allocated across downstream workers:

Proposition 4 (Supplier ability effect): The effect of supplier ability on processor output is unaffected by a team's ethnicity configuration if the supplier has ethnicity-neutral social preferences ($\theta_C = \theta_{NC}$): $\partial q_H / \partial \alpha_s = \partial q_{HM,C} / \partial \alpha_s = \partial q_{HM,NC} / \partial \alpha_s = \partial q_{VM} / \partial \alpha_s$. Higher supplier ability benefits processor output more (a) when working with a coethnic supplier, and (b) when working with another processor who is not of the supplier's ethnic group if the supplier has discriminatory preferences ($\theta_C > \theta_{NC}$): $\partial q_{HM,C} / \partial \alpha_s > \partial q_H / \partial \alpha_s > \partial q_{VM} / \partial \alpha_s > \partial q_{HM,NC} / \partial \alpha_s$

Biased, higher-ability suppliers allocate more of their additional capacity to supplying coethnic processors because they derive greater benefits from coethnics' output.

It is possible that the period of ethnic conflict in Kenya in early 2008 led to a change in attitudes towards co-workers of the other ethnic group, which I model as a change in θ_{NC} :

Proposition 5 (Change in preferences): A decrease in the weight attached to the well-being of non-coethnics leads to an increase in the output of the processor of the supplier's ethnicity in horizontally mixed teams, no change in the output of processors in homogeneous teams, and a decrease in the output of processors who are not of the supplier's ethnicity. The decrease is greater for non-coethnic processors in horizontally mixed teams: $\partial q_{HM,C} / \partial \theta_{NC} < 0 = \partial q_H / \partial \theta_{NC} \leq \partial q_{VM} / \partial \theta_{NC} \leq \partial q_{HM,NC} / \partial \theta_{NC}$

If the gap between the weight attached to coethnics' and non-coethnics' well-being widens, so does the output gap between teams of different ethnicity configurations.

Six weeks into the conflict period the plant began paying processors for their combined output. Under such a pay system a processor's utility from pay is $w(q_1 + q_2)$, rather than $2wq_1$ as under individual pay. Processor 1's problem becomes:

$$\begin{aligned} \text{Max}_{e_1} \quad & w \left((e_1 \alpha_1)^\beta (e_{s1} \alpha_s)^\gamma + (e_2 \alpha_2)^\beta (e_{s2} \alpha_s)^\gamma \right) - \frac{1}{2} e_1^2 \\ \text{s.t.} \quad & e_1 \geq 0 \end{aligned}$$

which gives

$$e_1 = \left(w \beta (e_{s1} \alpha_s)^\gamma \alpha_1^\beta \right)^{\frac{1}{2-\beta}}$$

Under team pay the supplier solves

$$\begin{aligned} \text{Max}_{e_{s1}, e_{s2}} \quad & w \left((e_1 \alpha_1)^\beta (e_{s1} \alpha_s)^\gamma + (e_2 \alpha_2)^\beta (e_{s2} \alpha_s)^\gamma \right) - \frac{1}{2} (e_{s1} + e_{s2})^2 \\ & + w(\theta_1 + \theta_2) \left((e_1 \alpha_1)^\beta (e_{s1} \alpha_s)^\gamma + (e_2 \alpha_2)^\beta (e_{s2} \alpha_s)^\gamma \right) - \theta_1 \frac{1}{2} e_1^2 - \theta_2 \frac{1}{2} e_2^2 \\ \text{s.t.} \quad & e_{s1} \geq 0 \text{ and } e_{s2} \geq 0 \end{aligned}$$

The supplier's first order condition for e_{s1} gives

$$e_{s1} + e_{s2} = w(1 + \theta_1 + \theta_2) (e_1 \alpha_1)^\beta \gamma (e_{s1} \alpha_s)^{\gamma-1} \alpha_s$$

When the supplier's two first order conditions hold simultaneously,

$$e_{s1} = \left(\frac{w(1 + \theta_1 + \theta_2) \gamma (e_1 \alpha_1)^\beta \alpha_s^\gamma}{1 + \left(\frac{e_2 \alpha_2}{e_1 \alpha_1} \right)^{\frac{\beta}{1-\gamma}}} \right)^{\frac{1}{2-\gamma}}$$

Because effort devoted to supplying one processor benefits both processors under team pay, the supplier's effort in supplying processor 1 is increasing in both θ_1 and θ_2 . If the two processors are of the same ability $e_{s1} = e_{s2}$ under team pay.

Solving the model under team pay gives the following predictions:

Proposition 6 (Team pay):

- i. There exists a unique equilibrium under team pay in which production is given by

$$q_1^{TP*} = \frac{k_q^{TP} \alpha_s^{\frac{\gamma}{2-\beta-\gamma}} \alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (1 + \theta_1 + \theta_2)^{\frac{\gamma}{2-\beta-\gamma}}}{\left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} \right)^{\frac{\gamma}{2-\beta-\gamma}}}$$

$$Q^{TP*} = k_q^{TP} \alpha_s^{\frac{\gamma}{2-\beta-\gamma}} \left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} \right)^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}} (1 + \theta_1 + \theta_2)^{\frac{\gamma}{2-\beta-\gamma}}$$

where $k_q^{TP} = \gamma^{\frac{\gamma}{2-\beta-\gamma}} w^{\frac{\beta+\gamma}{2-\beta-\gamma}} \beta^{\frac{\beta+2\gamma}{2-\beta-\gamma}}$.

- ii. Output in homogeneous and vertically mixed teams falls when team pay is introduced: $Q_H^{TP} < Q_H$ and $Q_{VM}^{TP} < Q_{VM}$
- iii. Output in homogeneous teams will continue to exceed that in vertically mixed teams under team pay if suppliers have discriminatory preferences ($\theta_C > \theta_{NC}$): $Q_H^{TP} > Q_{VM}^{TP}$
- iv. The output of the processor of the supplier's ethnicity and the processor who is not of the supplier's ethnicity in horizontally mixed teams is equal under team pay, even if suppliers have ethnic preferences ($\theta_C > \theta_{NC}$): $q_{HM,C}^{TP} = q_{HM,NC}^{TP}$
- v. Output in horizontally mixed teams Q_{HM}^{TP} can decrease or increase when team pay is introduced if suppliers have discriminatory preferences ($\theta_C > \theta_{NC}$): $Q_{HM}^{TP} \gtrless Q_{HM}$

In scenarios in which the two downstream workers are of the same ethnic group - homogeneous and vertically mixed teams - the supplier's problem reduces to the same problem she faced under individual pay. In such teams equilibrium production falls under team pay as processors freeride on each other. $Q_H > Q_{VM}$ is expected to continue to hold because biased suppliers' incentive to discriminate against non-coethnics through total supply remains under team pay.

In addition to the negative freeriding effect, team pay is expected to have an offsetting positive effect in horizontally mixed teams, in which $\theta_1 \neq \theta_2$. Because the two processors in a team are paid the same under team pay, the supplier is unable to increase her own utility by "shifting" roses from less to more favored processors. Eliminating horizontal misallocation will positively affect team output.

A2. Empirical Appendix

A2a. Estimation of utility-weights

I estimate the ratio-of-ratios on a sample of horizontally mixed teams in which a supplier is followed by another supplier of the other ethnic group. Instead of comparing the change in output from one day to the next, I compare average output under the first supplier, $s = 0$, to average output under the second supplier, $s = 1$. The log of the numerator of the left-hand side of the ratio-of-ratios is regressed on the log of the denominator and a constant:

$$\log(q_{1,\overline{s=0}}/q_{2,\overline{s=0}}) = \lambda + \eta \log(q_{1,\overline{s=1}}/q_{2,\overline{s=1}}) + \varepsilon \quad (1.1)$$

The resulting $\widehat{\lambda}$ can be interpreted as an estimate of $\log((1 + 2\theta_1/1 + 2\theta_2)^{\frac{4\gamma}{2-\beta-2\gamma}})$. Arranging the data such that $\log((1 + 2\theta_1/1 + 2\theta_2)^{\frac{4\gamma}{2-\beta-2\gamma}}) = \log((1 + 2\theta_C/1 + 2\theta_{NC})^{\frac{4\gamma}{2-\beta-2\gamma}})$ and estimating (1.1) on pre-conflict data gives $\widehat{\lambda} = 0.36$. $\widehat{\lambda}'$, from estimating (1.1) on data from the conflict period, is 0.52. Both estimates are significantly greater than zero at the 1% level.

Noting that $\widehat{\theta}_C = \frac{1}{2} \left(\left(\exp(\widehat{\lambda}) \right)^{\frac{2-\beta-2\gamma}{4\gamma}} \left(1 + 2\widehat{\theta}_{NC} \right) - 1 \right)$, with $\widehat{\lambda}$ in hand we can evaluate the locus of pairs of utility-weights that can explain the observed change in output when a supplier of one ethnic group replaces a supplier of the other ethnic group. Focusing on the pre-conflict period, suppose for example that $\beta = \gamma = 0.3$. Consider four possibilities: $(\theta_C > 0$ and $\theta_{NC} = 0)$, $(\theta_C = 0, \theta_{NC} < 0)$, $(\theta_C > 0$ and $\theta_{NC} > 0)$, $(\theta_C > 0, \theta_{NC} < 0)$. In the first two cases, $\widehat{\lambda} = 0.36$ implies $\widehat{\theta}_C \approx 0.19$ and $\widehat{\theta}_{NC} \approx -0.14$, respectively. In the latter two cases, further assumptions are required. If both preference parameters are positive, $\theta_C = 0.25$ implies $\widehat{\theta}_{NC} \approx 0.04$ while $\theta_C = 0.4$ implies $\widehat{\theta}_{NC} \approx 0.15$. If individuals attach positive weight to the well-being of coethnics and negative to the well-being of non-coethnics, suppose that $|\theta_C| = |\theta_{NC}|$. In that case $\widehat{\theta}_C \approx 0.08$, $\widehat{\theta}_{NC} \approx -0.08$.

Suppose that θ_C did not change when conflict began, as the results of table 8 suggest. Then,

$$1 = \frac{\widehat{\theta}_C}{\widehat{\theta}'_C} = \frac{\frac{1}{2} \left(\left(\exp(\widehat{\lambda}) \right)^{\frac{2-\beta-2\gamma}{4\gamma}} \left(1 + 2\widehat{\theta}_{NC} \right) - 1 \right)}{\frac{1}{2} \left(\left(\exp(\widehat{\lambda}') \right)^{\frac{2-\beta-2\gamma}{4\gamma}} \left(1 + 2\widehat{\theta}'_{NC} \right) - 1 \right)}$$

which gives

$$\theta'_{NC} = \frac{1}{2} \left(\frac{1 + 2\widehat{\theta}_{NC}}{\left(\exp(\widehat{\lambda}' - \widehat{\lambda}) \right)^{\frac{2-\beta-2\gamma}{4\gamma}}} - 1 \right) = \frac{1}{2} \left(\frac{1 + 2\widehat{\theta}_{NC}}{\left(\exp(0.16) \right)^{\frac{2-\beta-2\gamma}{4\gamma}}} - 1 \right)$$

Assumptions on θ_{NC} , β and γ are need to bound $\Delta\theta_{NC}$; I consider a wide parameter space. Given the plant's chosen triangular organization of production, β and γ are arguably likely to

be of similar magnitude. Consider $\beta \in [\frac{1}{3}, \frac{2}{3}]$, $\gamma \in [\frac{1}{3}, \frac{2}{3}]$ (subject to $\beta + \gamma < 1$). Letting $\theta_{NC} \in [-0.15, 0.15]$, $\beta = \gamma = \frac{1}{3}$ gives $\Delta\theta_{NC} \in [-0.07, -0.04]$ and $\Delta\theta_{NC}/\theta_{NC} \in [-1.27, -0.27]$. $\beta = \frac{1}{3}, \gamma = \frac{2}{3}$ gives $\Delta\theta_{NC} \approx -0.01$ and $\Delta\theta_{NC}/\theta_{NC} \in [-0.23, -0.09]$. $\beta = \frac{2}{3}, \gamma = \frac{1}{3}$ gives $\Delta\theta_{NC} \in [-0.05, -0.03]$ and $\Delta\theta_{NC}/\theta_{NC} \in [-0.71, -0.32]$. Averaging across the parameter space considered gives the bounds and average $\Delta\theta_{NC}/\theta_{NC}$ noted in the paper.

A2b. Optimal assignment procedure

I briefly describe the procedure used to compute the optimal assignments in table 11. See [Bhattacharya \(2009\)](#) for a more detailed description and justification of the procedure. The goal is to maximize the total output of a set of workers with multiple discrete characteristics. Discreteness implies a finite number of worker types, which can be combined into a finite number of team types. Output is maximized by choosing the quantities of each type of team that gives the highest total output, subject to the quantities of each worker type available. A solution to such a system is obtained using integer linear programming.

A worker is fully characterized by a collection of three discrete attributes: tribe, productivity tercile as supplier, and productivity tercile as processor. In turn, the set of possible team types is derived from the set of possible worker types. A team consists of one type of worker as supplier, one type of worker as processor 1, and one type of worker as processor 2.

The two processor positions are considered to be equivalent, and thus the number of processor pairs is calculated as two unordered draws with replacement from the pool of possible workers. There are $\binom{18+2-1}{2} = 171$ ways that these two can be chosen. Combining those with the 18 possibilities for the supplier gives 3078 distinct types of teams, if all possible types were to be considered. Those 3078 team types are mapped into 18 output coefficients when assignment is by productivity, and 63 output coefficients when assignment is by both productivity and tribe, as described in the paper.

An output-maximizing assignment is the solution of an integer linear programming problem with the following objective function:

$$\text{Max}_{t_1, \dots, t_{3078}} Q = \sum_{i=1}^{3078} \bar{Q}_i t_i$$

Each t_i term represents a possible type of team that can be formed from three workers, and \bar{Q}_i is the average output of that type of team.

The maximization of the objective function is constrained by the number of each type of worker that is present at the plant. For each worker type w_j , a constraint equates the number of workers used with the number of workers in the workforce:

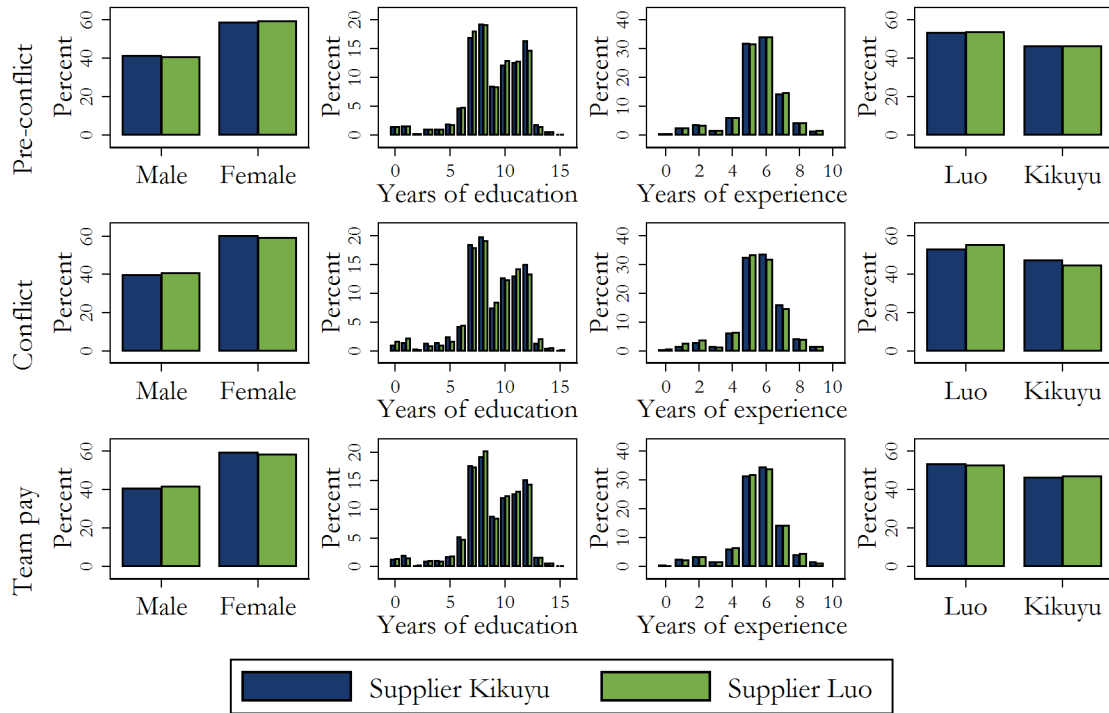
$$\begin{aligned} & \sum \{t_i | \text{there is 1 } w_j \text{ worker in } t_i\} \\ & + 2 \sum \{t_i | \text{there are 2 } w_j \text{ workers in } t_i\} \\ & + 3 \sum \{t_i | \text{there are 3 } w_j \text{ workers in } t_i\} = w_j \end{aligned}$$

The result of building these constraints is an 18×3078 matrix equation for which the columns represent team types and the rows worker types.

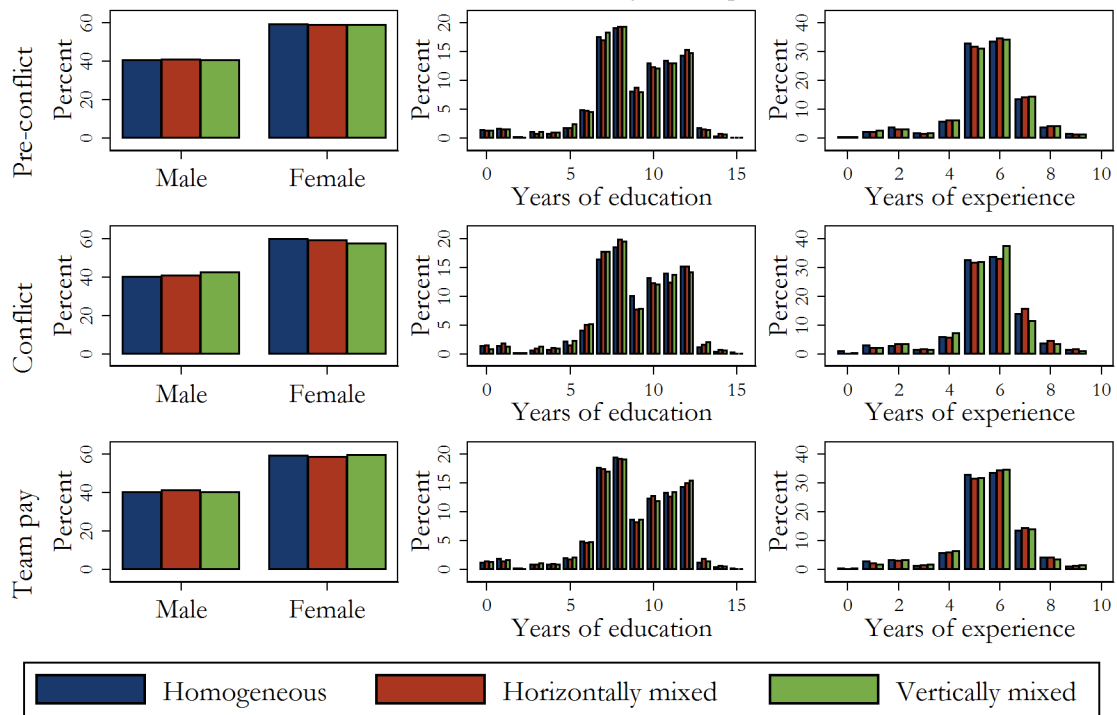
The optimal assignments in table 11 were obtained by solving these problems using the Gurobi solver. The output associated with “random assignment” in the table was computed by drawing 300 random assignments and taking the average output of those.

Appendix Figures

Appendix Figure 1: Distribution of co-worker characteristics for Kikuyu and Luo suppliers



Appendix Figure 2: Distribution of worker characteristics in teams of different ethnicity configurations



A3: Online Appendix

Proof of proposition 1 (Existence and comparative statics)

Existence.

I show solutions for processor 1, processor 2 is analogous. Processor 1's first order condition gives

$$e_1^* = \left(2w\beta(e_{s1}\alpha_s)^\gamma \alpha_1^\beta \right)^{\frac{1}{2-\beta}}$$

The supplier's first order condition for e_{sp} gives

$$(e_{s1} + e_{s2}) = (1 + 2\theta_1) w(e_1\alpha_1)^\beta \gamma (e_{s1}\alpha_s)^{\gamma-1}$$

As $(1 + 2\theta_1) > 0$, and the other terms in the roots are positive, $e_1^* > 0$ and $e_{s1}^* > 0$, which implies $q_1^* > 0$. Solving gives

$$q_1^* = \frac{k_q \alpha_s^{\frac{2\gamma}{2-\beta-\gamma}} \alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (1 + 2\theta_1)^{\frac{2\gamma}{2-\beta-2\gamma}}}{\left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} (1 + 2\theta_1)^{\frac{2-\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} (1 + 2\theta_1)^{\frac{2-\beta}{2-2\gamma-\beta}} \right)^{\frac{\gamma}{2-\beta-\gamma}}}$$

$$Q^* = \frac{k_q \alpha_s^{\frac{2\gamma}{2-\beta-\gamma}} \left(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (1 + 2\theta_1)^{\frac{2\gamma}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} (1 + 2\theta_2)^{\frac{2\gamma}{2-\beta-2\gamma}} \right)}{\left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} (1 + 2\theta_1)^{\frac{2-\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} (1 + 2\theta_2)^{\frac{2-\beta}{2-2\gamma-\beta}} \right)^{\frac{\gamma}{2-\beta-\gamma}}}$$

where $k_q = (2\beta)^{\frac{\beta}{2-\gamma-\beta}} w^{\frac{\beta+\gamma}{2-\gamma-\beta}} \gamma^{\frac{\gamma}{2-\gamma-\beta}}$ and $Q = q_1 + q_2$ is team output.

Call processor p 's utility U_p and supplier s 's utility U_s . As

$$\frac{d^2 U_p}{de_p^2} = 2w\beta(\beta - 1) \alpha_p^\beta (e_{sp}\alpha_s)^\gamma e_p^{\beta-2} - 1$$

$$\frac{d^2 U_s}{de_{sp}^2} = (1 + 2\theta_p) w\gamma(\gamma - 1) \alpha_s^\gamma (e_p\alpha_p)^\beta e_{sp}^{\gamma-2} - 1$$

the second order conditions are globally satisfied as long as output is concave in its arguments, $\beta < 1$ and $\gamma < 1$.

Comparative statics.

To evaluate the comparative statics, replace $(1 + 2\theta_l)$ with $\tilde{\theta}_l$ where $l \in \{C, NC, 1, 2\}$, note that $\tilde{\theta}_l > 0$ as $\theta_l > -\frac{1}{2}$. Define c to be such that $c * (1 + 2\theta_{NC}) = (1 + 2\theta_C) \iff c = \frac{(1+2\theta_C)}{(1+2\theta_{NC})} = \frac{\tilde{\theta}_C}{\tilde{\theta}_{NC}}$. Further define c_p , where $p \in \{1, 2\}$, to be such that $c_p * (1 + 2\theta_{NC}) = (1 + 2\theta_p)$. So $c_p = c$ if p is of the supplier's ethnic group and $c_p = 1$ if not. Define q_T where

$T \in \{H, HM_C, HM_{NC}, VM\}$. So for example $q_T = q_H = q(\theta_1 = \theta_C, \theta_2 = \theta_C)$ if $T = H$. Then

$$\partial q / \partial \alpha_s = \frac{2\gamma}{2 - \beta - 2\gamma} \alpha_s^{\frac{\beta-2}{2-\beta-2\gamma}} * q_T > 0 \text{ because all terms are positive}$$

$$\frac{\partial q_1}{\partial \alpha_1} = H \frac{c_1^{\frac{2\gamma}{2-\beta-2\gamma}} \left(\frac{2-\beta-2\gamma}{2-\beta-\gamma} \alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} c_1^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} c_2^{\frac{2-\beta}{2-\beta-2\gamma}} \right)}{\left(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} c_1^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} c_2^{\frac{2-\beta}{2-\beta-2\gamma}} \right)^{\frac{2-\beta}{2-\beta-\gamma}}}$$

$$\text{where } H = \tilde{\theta}_{NC}^{\frac{2\gamma}{2-\beta-2\gamma}} \left(\frac{2\beta}{2-\beta-2\gamma} \alpha_1^{\frac{2\beta}{2-\beta-2\gamma}-1} \right) \frac{\tilde{\theta}_{NC}^{\frac{2-\beta}{2-\beta-2\gamma}}}{\tilde{\theta}_{NC}^{\frac{2-\beta}{2-\beta-2\gamma}} \frac{2-\beta}{2-\beta-\gamma}} > 0,$$

so $\frac{\partial q_1}{\partial \alpha_1} > 0$ as all terms are positive.

$$\frac{\partial q_1}{\partial \alpha_2} = \tilde{H} \left(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} c_1^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} c_2^{\frac{2-\beta}{2-\beta-2\gamma}} \right)^{\frac{\beta-2}{2-\beta-\gamma}} c_2^{\frac{2-\beta}{2-\beta-2\gamma}} c_1^{\frac{2\gamma}{2-\beta-2\gamma}} < 0$$

where

$$\tilde{H} = \left(k_q \alpha_s^{\frac{2\gamma}{2-\beta-\gamma}} \alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} \alpha_2^{\frac{3\beta-2+2\gamma}{2-\beta-2\gamma}} \left(\frac{-\gamma}{2-\beta-\gamma} \right) \left(\frac{2\beta}{2-\beta-2\gamma} \right) \right) \tilde{\theta}_{NC}^{\frac{2-\beta}{2-\beta-2\gamma} \frac{\beta-2}{2-\beta-\gamma} + \frac{2-\beta+2\gamma}{2-\beta-2\gamma}} < 0,$$

so $\frac{\partial q_1}{\partial \alpha_2} < 0$ as \tilde{H} is negative and all other terms are positive.

To analyze the comparative static of q_1 with respect to θ_1 and θ_2 , it is convenient to let $\theta_{p,old} = \theta_p$ and $\theta_{p,new} = k * \theta_p$. WLOG, let $k > 0$. Consider first an increase in θ_p . For the change in output, Δq_{θ_p} , we then have that:

$$\begin{aligned}
 \Delta q_{\theta_p} &= \frac{K(\theta_{p,new})^{\frac{2\gamma}{2-\beta-2\gamma}}}{\left(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (\theta_{p,new})^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} \theta_o^{\frac{2-\beta}{2-\beta-2\gamma}}\right)^{\frac{\gamma}{2-\beta-\gamma}}} - \\
 &\quad \frac{K(\theta_{p,old})^{\frac{2\gamma}{2-\beta-2\gamma}}}{\left(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (\theta_{p,old})^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} \theta_o^{\frac{2-\beta}{2-\beta-2\gamma}}\right)^{\frac{\gamma}{2-\beta-\gamma}}} \\
 &\quad \text{where } K \text{ is } k_q \alpha_s^{\frac{2\gamma}{2-\beta-\gamma}} \\
 &= K(\theta_p)^{\frac{2\gamma}{2-\beta-2\gamma}} * \left(\frac{\left(\frac{\left(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} \theta_p^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} \theta_o^{\frac{2-\beta}{2-\beta-2\gamma}} \right)^{\frac{\gamma}{2-\beta-\gamma}}}{\left(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (k\theta_p)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} \theta_o^{\frac{2-\beta}{2-\beta-2\gamma}} \right)^{\frac{\gamma}{2-\beta-\gamma}}} \right) k^{\frac{2\gamma}{2-\beta-2\gamma}}}{\left(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} \theta_p^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} \theta_o^{\frac{2-\beta}{2-\beta-2\gamma}} \right)^{\frac{\gamma}{2-\beta-\gamma}}} - 1 \right) \\
 &> 0 \text{ because, } \left(\frac{\left(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} \theta_p^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} \theta_o^{\frac{2-\beta}{2-\beta-2\gamma}} \right)^{\frac{\gamma}{2-\beta-\gamma}}}{\left(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (k\theta_p)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} \theta_o^{\frac{2-\beta}{2-\beta-2\gamma}} \right)^{\frac{\gamma}{2-\beta-\gamma}}} k^{\frac{2\gamma}{2-\beta-2\gamma}} \right) - 1 > 0.
 \end{aligned}$$

$$\begin{aligned}
 \Delta q_{\theta_o} &= \frac{K(\theta_p)^{\frac{2\gamma}{2-\beta-2\gamma}}}{\left(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} \theta_p^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} (\theta_{o,new})^{\frac{2-\beta}{2-\beta-2\gamma}}\right)^{\frac{\gamma}{2-\beta-\gamma}}} - \\
 &\quad \frac{K(\theta_p)^{\frac{2\gamma}{2-\beta-2\gamma}}}{\left(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} \theta_p^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} (\theta_{o,old})^{\frac{2-\beta}{2-\beta-2\gamma}}\right)^{\frac{\gamma}{2-\beta-\gamma}}} \\
 &< \frac{K(\theta_p)^{\frac{2\gamma}{2-\beta-2\gamma}}}{\left(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} \theta_p^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} \theta_o^{\frac{2-\beta}{2-\beta-2\gamma}}\right)^{\frac{\gamma}{2-\beta-\gamma}}} - \\
 &\quad \frac{K(\theta_p)^{\frac{2\gamma}{2-\beta-2\gamma}}}{\left(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} \theta_p^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} \theta_o^{\frac{2-\beta}{2-\beta-2\gamma}}\right)^{\frac{\gamma}{2-\beta-\gamma}}} = 0
 \end{aligned}$$

□

Proof of proposition 2 (Individual output)

With $\tilde{\theta}_C = \tilde{\theta}$ as a baseline, express $\tilde{\theta}_{NC} = c\tilde{\theta}$, where $c = \frac{\tilde{\theta}_{NC}}{\tilde{\theta}_C}$. Replace $(1 + 2\theta_l)$ with $\tilde{\theta}_l$ where $l \in \{C, NC, 1, 2\}$. Let $\tilde{\theta} = \tilde{\theta}_{NC}$ and define c to be such that $\tilde{\theta}_C = c\tilde{\theta}_{NC} = c\tilde{\theta}$. Define q_T where $T \in \{H, HM_C, HM_{NC}, VM\}$. So for example $q_T = q_H = q(\theta_1 = \theta_C, \theta_2 = \theta_C)$ if

$T = H$. Let

$$\begin{aligned} A &= K\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (\tilde{\theta})^{\frac{2\gamma}{2-\beta-2\gamma}}, \text{ where } K = k_q\alpha_s^{\frac{2\gamma}{2-\beta-\gamma}} \\ B &= (\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} \tilde{\theta}^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} \tilde{\theta}^{\frac{2-\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}} \\ B' &= (\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} \tilde{\theta}^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} \tilde{\theta}^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}} \\ B'' &= (\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} \tilde{\theta}^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} \left(\frac{\tilde{\theta}}{c}\right)^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}} \end{aligned}$$

Noting that $\frac{1}{2} > \theta_l > -\frac{1}{2}$ implies $2 > \tilde{\theta}_l > 0$, $B'' = (\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} \tilde{\theta}^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} \left(\frac{\tilde{\theta}}{c}\right)^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}} > 0$ i.f.f. $\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} \tilde{\theta}^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} \left(\frac{\tilde{\theta}}{c}\right)^{\frac{2-\beta}{2-\beta-2\gamma}} > 0$, which holds.

Then

$$\begin{aligned} q_{VM} &= \frac{K\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (\tilde{\theta})^{\frac{2\gamma}{2-\beta-2\gamma}}}{(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} \tilde{\theta}^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} \tilde{\theta}^{\frac{2-\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} = \frac{A}{B'} \\ q_{HM,NC} &= \frac{K\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (\tilde{\theta})^{\frac{2\gamma}{2-\beta-2\gamma}}}{(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} \tilde{\theta}^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} \tilde{\theta}^{\frac{2-\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} = \frac{A}{B} \\ q_H &= \frac{K\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (c\tilde{\theta})^{\frac{2\gamma}{2-\beta-2\gamma}}}{(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (c\tilde{\theta})^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} (c\tilde{\theta})^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} = \frac{A}{B'} c^{2-\beta-2\gamma - (\frac{2-\beta}{2-\beta-2\gamma})(\frac{\gamma}{2-\beta-\gamma})} \\ q_{HM,C} &= \frac{K\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (c\tilde{\theta})^{\frac{2\gamma}{2-\beta-2\gamma}}}{(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (c\tilde{\theta})^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} \tilde{\theta}^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} = \frac{A}{B''} c^{2-\beta-2\gamma - (\frac{2-\beta}{2-\beta-2\gamma})(\frac{\gamma}{2-\beta-\gamma})} \end{aligned}$$

As $B > B' > B'' > 0$ and $c > 1$ we have that $q_{HM,C} > q_H > q_{VM} > q_{HM,NC}$.

□

Proof of proposition 3 (Team output)

Q_H vs Q_{VM} :

$$\begin{aligned} Q_H &= k_q\alpha_s^{\frac{2\gamma}{2-\beta-\gamma}} \left(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} \right)^{\frac{2-2\gamma-\beta}{2-\gamma-\beta}} (1 + 2\theta_C)^{\frac{\gamma}{2-\beta-\gamma}} \\ &> k_q\alpha_s^{\frac{2\gamma}{2-\beta-\gamma}} \left(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} \right)^{\frac{2-2\gamma-\beta}{2-\gamma-\beta}} (1 + 2\theta_{NC})^{\frac{\gamma}{2-\beta-\gamma}} = Q_{VM} \\ &\iff \\ \theta_C &> \theta_{NC} \end{aligned}$$

Q_H vs Q_{HM} :

Let $\tilde{\theta}_C = 1 + 2\theta_C$, $\tilde{\theta}_{NC} = 1 + 2\theta_{NC}$ and $c = \frac{\tilde{\theta}_C}{\tilde{\theta}_{NC}}$ (note that $c > 1$). To ease the notation, let $\tilde{\theta}_{NC} = \tilde{\theta}$.

$$\frac{Q_H}{Q_{HM}} = \frac{\left(\frac{K(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}}(c\tilde{\theta})^{\frac{2\gamma}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}}(c\tilde{\theta})^{\frac{2\gamma}{2-\beta-2\gamma}})}{(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}}(c\tilde{\theta})^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}}(c\tilde{\theta})^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} \right)}{\left(\frac{K(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}}\tilde{\theta}^{\frac{2\gamma}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}}(c\tilde{\theta})^{\frac{2\gamma}{2-\beta-2\gamma}})}{(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}}\tilde{\theta}^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}}(c\tilde{\theta})^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} \right)}, \text{ where } K = k_q \alpha_s^{\frac{2\gamma}{2-\beta-\gamma}}$$

Rearranging terms, canceling out the K 's, $\tilde{\theta}$'s and factoring out the common terms $c^{\frac{2\gamma}{2-\beta-2\gamma}}$ and $\frac{1}{(c^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}}$,

$$\begin{aligned} \frac{Q_H}{Q_{HM}} &= c^{\frac{2\gamma}{2-\beta-2\gamma} - \frac{2-\beta}{2-\beta-2\gamma} \frac{\gamma}{2-\beta-\gamma}} \frac{(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}})^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}}}{(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2\gamma}{2-\beta-2\gamma}})^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}}} (\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}} \\ &= c^{\frac{\gamma}{2-\beta-\gamma} (\frac{2-\beta-2\gamma}{2-\beta-\gamma} + \frac{\gamma}{2-\beta-\gamma})} \frac{(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}})^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}}}{(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{\gamma}{2-\beta-\gamma}})^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}}} (\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}} \\ &= c^{(\frac{\gamma}{2-\beta-\gamma})^2} \frac{(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{\gamma}{2-\beta-\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{\gamma}{2-\beta-\gamma}})^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}}}{(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{\gamma}{2-\beta-\gamma}})^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}}} (\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}} \\ &> c^{(\frac{\gamma}{2-\beta-\gamma})^2} \frac{(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{\gamma}{2-\beta-\gamma}})^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}}}{(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{\gamma}{2-\beta-\gamma}})^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}}} (\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}} \\ &\text{as } \alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{\gamma}{2-\beta-\gamma}} > \alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} \\ &= c^{(\frac{\gamma}{2-\beta-\gamma})^2} \frac{(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}}{(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{\gamma}{2-\beta-\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} > c^{(\frac{\gamma}{2-\beta-\gamma})^2} \\ &\text{as } \frac{\gamma}{2-\beta-\gamma} < \frac{2-\beta}{2-\beta-2\gamma} \text{ implies} \\ &\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}} > \alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{\gamma}{2-\beta-\gamma}} \\ &\text{So } \frac{Q_H}{Q_{HM}} > c^{(\frac{\gamma}{2-\beta-\gamma})^2} > 1 \text{ and thus } Q_H > Q_{HM} \end{aligned}$$

Q_{HM} vs Q_{VM} :

The ranking of Q_{HM} and Q_{VM} is ambiguous in general. Consider the case where $\alpha_1 = \alpha_2$ and let α denote the common value for α_1 and α_2 . Then

$$\begin{aligned}
 \frac{Q_{HM}}{Q_{VM}} &= \frac{\left(\frac{K(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (\tilde{\theta}_{NC})^{\frac{2\gamma}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} (\tilde{\theta}_C)^{\frac{2\gamma}{2-\beta-2\gamma}})}{(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (\tilde{\theta}_{NC})^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} (\tilde{\theta}_C)^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} \right)}{\left(\frac{K(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (\tilde{\theta}_{NC})^{\frac{2\gamma}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} (\tilde{\theta}_{NC})^{\frac{2\gamma}{2-\beta-2\gamma}})}{(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (\tilde{\theta}_{NC})^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} (\tilde{\theta}_{NC})^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} \right)} \\
 &= \frac{\left(\frac{(\tilde{\theta}_{NC})^{\frac{2\gamma}{2-\beta-2\gamma}} + (\tilde{\theta}_C)^{\frac{2\gamma}{2-\beta-2\gamma}}}{((\tilde{\theta}_{NC})^{\frac{2-\beta}{2-\beta-2\gamma}} + (\tilde{\theta}_C)^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} \right)}{\left(\frac{(\tilde{\theta}_{NC})^{\frac{2\gamma}{2-\beta-2\gamma}} + (\tilde{\theta}_{NC})^{\frac{2\gamma}{2-\beta-2\gamma}}}{((\tilde{\theta}_{NC})^{\frac{2-\beta}{2-\beta-2\gamma}} + (\tilde{\theta}_{NC})^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} \right)} \\
 &= \frac{(\tilde{\theta}_{NC})^{\frac{2\gamma}{2-\beta-2\gamma}} + (\tilde{\theta}_C)^{\frac{2\gamma}{2-\beta-2\gamma}}}{(\tilde{\theta}_{NC})^{\frac{2\gamma}{2-\beta-2\gamma}} + (\tilde{\theta}_{NC})^{\frac{2\gamma}{2-\beta-2\gamma}}} \left(\frac{(\tilde{\theta}_{NC})^{-1}}{(\tilde{\theta}_{NC})^{-1}} * \frac{(\tilde{\theta}_{NC})^{\frac{2-\beta}{2-\beta-2\gamma}} + (\tilde{\theta}_{NC})^{\frac{2-\beta}{2-\beta-2\gamma}}}{(\tilde{\theta}_{NC})^{\frac{2-\beta}{2-\beta-2\gamma}} + (\tilde{\theta}_C)^{\frac{2-\beta}{2-\beta-2\gamma}}} \right)^{\frac{\gamma}{2-\beta-\gamma}} \\
 &= \frac{(\tilde{\theta}_{NC})^{\frac{2\gamma}{2-\beta-2\gamma}} + (\tilde{\theta}_C)^{\frac{2\gamma}{2-\beta-2\gamma}}}{((\tilde{\theta}_{NC})^{\frac{2\gamma}{2-\beta-2\gamma}} + (\tilde{\theta}_{NC})^{-1}(\tilde{\theta}_C)^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} \left(\frac{1}{(\tilde{\theta}_{NC})^{\frac{2\gamma}{2-\beta-2\gamma}} + (\tilde{\theta}_{NC})^{\frac{2\gamma}{2-\beta-2\gamma}}} \right)^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}} \\
 &> \frac{(\tilde{\theta}_{NC})^{\frac{2\gamma}{2-\beta-2\gamma}} + (\tilde{\theta}_C)^{\frac{2\gamma}{2-\beta-2\gamma}}}{((\tilde{\theta}_{NC})^{\frac{2\gamma}{2-\beta-2\gamma}} + (\tilde{\theta}_C)^{-1}(\tilde{\theta}_C)^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} \left(\frac{1}{(\tilde{\theta}_{NC})^{\frac{2\gamma}{2-\beta-2\gamma}} + (\tilde{\theta}_{NC})^{\frac{2\gamma}{2-\beta-2\gamma}}} \right)^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}} \\
 &= \left(\frac{(\tilde{\theta}_{NC})^{\frac{2\gamma}{2-\beta-2\gamma}} + (\tilde{\theta}_C)^{\frac{2\gamma}{2-\beta-2\gamma}}}{(\tilde{\theta}_{NC})^{\frac{2\gamma}{2-\beta-2\gamma}} + (\tilde{\theta}_{NC})^{\frac{2\gamma}{2-\beta-2\gamma}}} \right)^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}} > 1
 \end{aligned}$$

so $Q_{HM} > Q_{VM}$ if $\alpha_1 = \alpha_2$.

□

Proof of proposition 4 (Supplier ability effect)

Define q_T where $T \in \{H, HM_C, HM_{NC}, VM\}$. So for example $q_T = q_H = q(\theta_1 = \theta_C, \theta_2 = \theta_C)$ if $T = H$. We have that $\partial q_T / \partial \alpha_s = \frac{2\gamma}{2-\beta-2\gamma} \alpha_s^{\frac{\beta-2}{2-\beta-2\gamma}} * q_T$. As shown in proposition 2, $q_{HM,C} > q_H > q_{VM} > q_{HM,NC}$. Thus, $\partial q_{HM,C} / \partial \alpha_s > \partial q_H / \partial \alpha_s > \partial q_{VM} / \partial \alpha_s > \partial q_{HM,NC} / \partial \alpha_s$.

□

Proof of proposition 5 (Change in preferences):

WLOG consider an improvement in attitudes towards non-coethnics. Denote by $\tilde{\theta}'_i$ the new value of $\tilde{\theta}_i$ (where $\tilde{\theta}_i = 1 + 2\theta_i$). With $\tilde{\theta}_C = \tilde{\theta}$ as a baseline, express $\tilde{\theta}_{NC} = c\tilde{\theta}$, where $c = \frac{\tilde{\theta}_{NC}}{\tilde{\theta}_C}$. Further let $\tilde{\theta}'_{NC} = k\tilde{\theta}_{NC}$, where $k = \frac{\tilde{\theta}'_{NC}}{\tilde{\theta}_{NC}}$. Note that $c < 1$, $k > 1$ and $ck < 1$. (Unlike previous propositions, here we express θ_{NC} in terms of θ_C). Let p designate the processor in question and o the other processor.

We have that

$$\begin{aligned}
 q_p &= \frac{K(\tilde{\theta}_p)^{\frac{2\gamma}{2-\beta-2\gamma}}}{\left(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (\tilde{\theta}_p)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} (\tilde{\theta}_o)^{\frac{2-\beta}{2-\beta-2\gamma}}\right)^{\frac{\gamma}{2-\beta-\gamma}}}, \text{ where } K = k_q \alpha_s^{\frac{2\gamma}{2-\beta-\gamma}} \\
 \Delta q_\theta &= \frac{K(\theta'_p)^{\frac{2\gamma}{2-\beta-2\gamma}}}{\left(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (\theta'_p)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} (\theta'_o)^{\frac{2-\beta}{2-\beta-2\gamma}}\right)^{\frac{\gamma}{2-\beta-\gamma}}} - \frac{K(\theta_p)^{\frac{2\gamma}{2-\beta-2\gamma}}}{\left(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} \theta_p^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} \theta_o^{\frac{2-\beta}{2-\beta-2\gamma}}\right)^{\frac{\gamma}{2-\beta-\gamma}}} \\
 &= \frac{K(\tilde{\theta})^{\frac{2\gamma}{2-\beta-2\gamma}}}{\tilde{\theta}^{\frac{\gamma}{2-\beta-\gamma}}} * \left[\frac{(c'_p)^{\frac{2\gamma}{2-\beta-2\gamma}}}{\left(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (c'_p)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} (c'_o)^{\frac{2-\beta}{2-\beta-2\gamma}}\right)^{\frac{\gamma}{2-\beta-\gamma}}} \right. \\
 &\quad \left. - \frac{c_p^{\frac{2\gamma}{2-\beta-2\gamma}}}{\left(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} c_p^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} c_o^{\frac{2-\beta}{2-\beta-2\gamma}}\right)^{\frac{\gamma}{2-\beta-\gamma}}} \right]
 \end{aligned}$$

where $c_i = c$ and $c'_i = kc$ if processor i is not of the supplier's ethnic group and $c_i = c'_i = 1$ if processor i is of the supplier's ethnic group. So

$$\begin{aligned}
 &\frac{\tilde{\theta}^{\frac{\gamma}{2-\beta-\gamma}}}{K(\tilde{\theta})^{\frac{2\gamma}{2-\beta-2\gamma}}} \Delta q_\theta = \\
 &\frac{(c'_p)^{\frac{2\gamma}{2-\beta-2\gamma}}}{\left(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (c'_p)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} (c'_o)^{\frac{2-\beta}{2-\beta-2\gamma}}\right)^{\frac{\gamma}{2-\beta-\gamma}}} - \frac{c_p^{\frac{2\gamma}{2-\beta-2\gamma}}}{\left(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} c_p^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} c_o^{\frac{2-\beta}{2-\beta-2\gamma}}\right)^{\frac{\gamma}{2-\beta-\gamma}}} \\
 &= g(\alpha_p, \alpha_o, T, k).
 \end{aligned}$$

The values for c_i and c'_i are given by the T :

$$\begin{aligned}
 g_H &= \frac{1}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} - \frac{1}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} = 0 \\
 g_{VM} &= \frac{(kc)^{\frac{2\gamma}{2-\beta-2\gamma}}}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} - \frac{c^{\frac{2\gamma}{2-\beta-2\gamma}}}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} \\
 &= c^{\frac{\gamma}{2-\beta-2\gamma}} \left(k^{\frac{\gamma}{2-\beta-2\gamma}} - 1 \right) \left(\frac{1}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} \right) > 0 \\
 g_{HM,C} &= \frac{1}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} - \frac{1}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} < 0 \\
 g_{HM,NC} &= \frac{(kc)^{\frac{2\gamma}{2-\beta-2\gamma}}}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} - \frac{c^{\frac{2\gamma}{2-\beta-2\gamma}}}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} \\
 &= c^{\frac{2\gamma}{2-\beta-2\gamma}} \left(\frac{k^{\frac{2\gamma}{2-\beta-2\gamma}}}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} - \frac{1}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} \right) \\
 &> 0
 \end{aligned}$$

Furthermore,

$$\begin{aligned}
 g_{HM,NC} &= c^{\frac{2\gamma}{2-\beta-2\gamma}} \left(\frac{k^{\frac{2\gamma}{2-\beta-2\gamma}}}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} - \frac{1}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} \right) \\
 &< c^{\frac{2\gamma}{2-\beta-2\gamma}} \left(\frac{k^{\frac{2\gamma}{2-\beta-2\gamma}}}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} - \frac{1}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} \right) \\
 &= g_{VM}
 \end{aligned}$$

Because, keeping in mind that $1 > kc > c > 0$, $g_{HM,NC} > 0$, and $g_{VM} > 0$

$$\begin{aligned}
 \frac{g_{HM,NC}}{c^{\frac{2\gamma}{2-\beta-2\gamma}}} &= \frac{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} \\
 &\quad \frac{k^{\frac{2\gamma}{2-\beta-2\gamma}}}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} - \\
 &\quad \frac{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} \\
 &\quad \frac{1}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} \\
 &> \frac{k^{\frac{2\gamma}{2-\beta-2\gamma}}}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} - \\
 &\quad \frac{1}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} \\
 &= \frac{g_{VM}}{c^{\frac{2\gamma}{2-\beta-2\gamma}}}
 \end{aligned}$$

if

$$\begin{aligned}
 &\frac{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} > \frac{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} \\
 \Leftrightarrow &\frac{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} > \frac{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} \\
 \Leftrightarrow &\frac{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} - d)^{\frac{\gamma}{2-\beta-\gamma}}}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} - D)^{\frac{\gamma}{2-\beta-\gamma}}} > \frac{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}}
 \end{aligned}$$

where $d = (1 - kc) \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}}$ and $D = (1 - c) \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}}$.

The above inequality holds as $\frac{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}} > 1$, both d and D are positive and the numerator and denominator of $\frac{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} (kc)^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}}{(\alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}} + \alpha_o^{\frac{2\beta}{2-\beta-2\gamma}} c^{\frac{2-\beta}{2-\beta-2\gamma}})^{\frac{\gamma}{2-\beta-\gamma}}}$

is positive. So, $g_{HM,NC} > g_{VM}$.

Thus $g_{HM,NC} > g_{VM} > g_H > g_{HM,C}$ which implies $\Delta q_{\theta HM,NC} > \Delta q_{\theta VM} > \Delta q_{\theta H} > \Delta q_{\theta HM,C}$.

□

Proof of proposition 6 (Team pay):

i.

Processor 1's first order condition gives

$$e_1 = (w\beta(e_{s1}\alpha_s)^\gamma \alpha_1^\beta)^{\frac{1}{2-\beta}}$$

The supplier's first order condition for e_{s1} gives

$$e_{s1} + e_{s2} = w(1 + \theta_1 + \theta_2)(e_1\alpha_1)^\beta \gamma (e_{s1}\alpha_s)^{\gamma-1} \alpha_s$$

Solving gives

$$q_1^{TP} = \frac{k_q^{TP} \alpha_s^{\frac{\gamma}{2-\beta-\gamma}} \alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (1 + \theta_1 + \theta_2)^{\frac{\gamma}{2-\beta-\gamma}}}{\left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} \right)^{\frac{\gamma}{2-\beta-\gamma}}}$$

$$Q^{TP} = k_q^{TP} \alpha_s^{\frac{\gamma}{2-\beta-\gamma}} \left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} \right)^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}} (1 + \theta_1 + \theta_2)^{\frac{\gamma}{2-\beta-\gamma}}$$

where $k_q^{TP} = \gamma^{\frac{\gamma}{2-\beta-\gamma}} w^{\frac{\beta+\gamma}{2-\beta-\gamma}} \beta^{\frac{\beta+2\gamma}{2-\beta-\gamma}}$. As $(1 + \theta_1 + \theta_2) > 0$ and the other terms in the root are positive, $e_1^* > 0$ and $e_{s1}^* > 0$ which implies $q_1^{TP*} > 0$.

As,

$$\frac{d^2 U_p}{de_p^2} = w\beta(\beta - 1) \alpha_p (e_{sp} \alpha_s)^\gamma (e_p \alpha_p)^{\beta-2} - 1$$

$$\frac{d^2 U_s}{de_{sp}^2} = (w(1 + \theta_1 + \theta_2)) (e_1 \alpha_1)^\beta \alpha_s \gamma (\gamma - 1) (e_{s1} \alpha_s)^{\gamma-2} - 1$$

the second order conditions are globally satisfied as long as output is concave in its arguments, $\beta < 1$ and $\gamma < 1$.

ii.

To show that output in homogeneous and vertically mixed teams falls when team pay is introduced, it suffices to show that $Q^{TP} < Q$ when $\theta_1 = \theta_2$. Let $\theta_1 = \theta_2$, and denote their common value by θ . Then:

$$\begin{aligned}
 Q^* &= \frac{k_q \alpha_s^{\frac{2\gamma}{2-\beta-\gamma}} \left(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (1+2\theta)^{\frac{2\gamma}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} (1+2\theta)^{\frac{2\gamma}{2-\beta-2\gamma}} \right)}{\left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} (1+2\theta)^{\frac{2-\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} (1+2\theta)^{\frac{2-\beta}{2-2\gamma-\beta}} \right)^{\frac{\gamma}{2-\beta-\gamma}}} \\
 &= k_q \alpha_s^{\frac{2\gamma}{2-\beta-\gamma}} (1+2\theta)^{\frac{2\gamma}{2-\beta-2\gamma} - \frac{2-\beta}{2-2\gamma-\beta} \frac{\gamma}{2-\beta-\gamma}} \left(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} \right)^{1 - \frac{\gamma}{2-\beta-\gamma}} \\
 &= (2\beta)^{\frac{\beta}{2-\gamma-\beta}} w^{\frac{\beta+\gamma}{2-\gamma-\beta}} \gamma^{\frac{\gamma}{2-\gamma-\beta}} \alpha_s^{\frac{2\gamma}{2-\beta-\gamma}} (1+2\theta)^{\frac{\gamma}{2-\beta-\gamma}} \left(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} \right)^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}}
 \end{aligned}$$

and

$$\begin{aligned}
 Q^{TP*} &= k_q^{TP} \alpha_s^{\frac{\gamma}{2-\beta-\gamma}} \left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} \right)^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}} (1+\theta+\theta)^{\frac{\gamma}{2-\beta-\gamma}} \\
 &= k_q^{TP} \alpha_s^{\frac{\gamma}{2-\beta-\gamma}} (1+2\theta)^{\frac{\gamma}{2-\beta-\gamma}} \left(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} \right)^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}} \\
 &= \beta^{\frac{\beta+2\gamma}{2-\beta-\gamma}} w^{\frac{\beta+\gamma}{2-\gamma-\beta}} \gamma^{\frac{\gamma}{2-\gamma-\beta}} \alpha_s^{\frac{\gamma}{2-\beta-\gamma}} (1+2\theta)^{\frac{\gamma}{2-\beta-\gamma}} \left(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} \right)^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}}
 \end{aligned}$$

So we have,

$$\begin{aligned}
 \frac{Q^*}{Q^{TP*}} &= \frac{(2\beta)^{\frac{\beta}{2-\gamma-\beta}} w^{\frac{\beta+\gamma}{2-\gamma-\beta}} \gamma^{\frac{\gamma}{2-\gamma-\beta}} \alpha_s^{\frac{2\gamma}{2-\beta-\gamma}} (1+2\theta)^{\frac{\gamma}{2-\beta-\gamma}} \left(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} \right)^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}}}{\beta^{\frac{\beta+2\gamma}{2-\beta-\gamma}} w^{\frac{\beta+\gamma}{2-\gamma-\beta}} \gamma^{\frac{\gamma}{2-\gamma-\beta}} \alpha_s^{\frac{\gamma}{2-\beta-\gamma}} (1+2\theta)^{\frac{\gamma}{2-\beta-\gamma}} \left(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} \right)^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}}} \\
 &= \frac{(2\beta)^{\frac{\beta}{2-\gamma-\beta}} \alpha_s^{\frac{2\gamma}{2-\beta-\gamma}}}{\beta^{\frac{\beta+2\gamma}{2-\beta-\gamma}} \alpha_s^{\frac{\gamma}{2-\beta-\gamma}}} = \frac{2^{\frac{\beta}{2-\gamma-\beta}} \alpha_s^{\frac{\gamma}{2-\beta-\gamma}}}{\beta^{\frac{2\gamma}{2-\beta-\gamma}}} \\
 &= \alpha_s^{\frac{\gamma}{2-\beta-\gamma}} \left(\frac{2\beta}{\beta^{2\gamma}} \right)^{\frac{1}{2-\gamma-\beta}} > \alpha_s^{\frac{\gamma}{2-\beta-\gamma}}
 \end{aligned}$$

As it was assumed that $\alpha_s > 1$ we have $Q^* > Q^{TP*}$ in homogeneous and vertically mixed teams.

iii.

$$\begin{aligned}
 Q_H^{TP} &= k_q^{TP} \alpha_s^{\frac{2-\beta+4\gamma-\gamma\beta}{2-\beta}} \left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} \right)^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}} (1+2\theta_C)^{\frac{\gamma}{2-\beta-\gamma}} \\
 &> k_q^{TP} \alpha_s^{\frac{2-\beta+4\gamma-\gamma\beta}{2-\beta}} \left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} \right)^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}} (1+2\theta_{NC})^{\frac{\gamma}{2-\beta-\gamma}} = Q_{VM}^{TP} \\
 &\iff \theta_C > \theta_{NC}
 \end{aligned}$$

iv.

To show that $q_{HM,C}^{TP} = q_{HM,NC}^{TP}$, it suffices to show that in horizontally mixed teams, we have that individual output is invariant to that individual's ethnicity. To that effect we observe that the utility-weights enter in the expression for individual output,

$$q_1^{TP} = \frac{k_q^{TP} \alpha_s^{\frac{2\gamma}{2-\beta-\gamma}} \alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (1+\theta_1+\theta_2)^{\frac{\gamma}{2-\beta-\gamma}}}{\left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} \right)^{\frac{\gamma}{2-\beta-\gamma}}}$$

through the sum $(1+\theta_1+\theta_2)$, which is invariant to the whether processor 1 is the coethnic individual in the team.

v.

Let $\theta_{NC} = \theta$ and $\theta_C = c(\theta_{NC}) = c\theta$. We have that

$$\begin{aligned}
 &(Q_{HM}^{TP} - Q_{HM}) \\
 &= k_q^{TP} \alpha_s^{\frac{\gamma}{2-\beta-\gamma}} \left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} \right)^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}} (1+\theta_C+\theta_{NC})^{\frac{\gamma}{2-\beta-\gamma}} \\
 &\quad \frac{k_q \alpha_s^{\frac{2\gamma}{2-\beta-\gamma}} \left(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (1+2\theta_C)^{\frac{2\gamma}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} (1+2\theta_{NC})^{\frac{2\gamma}{2-\beta-2\gamma}} \right)}{\left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} (1+2\theta_C)^{\frac{2-\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} (1+2\theta_{NC})^{\frac{2-\beta}{2-2\gamma-\beta}} \right)^{\frac{\gamma}{2-\beta-\gamma}}} \\
 &= k_q \alpha_s^{\frac{\gamma}{2-\beta-\gamma}} \left[\frac{\left(\frac{2\beta}{\beta^{2\gamma}} \right)^{\frac{1}{2-\gamma-\beta}} \left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} (1+c\theta+\theta)^{\frac{2\gamma}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} (1+c\theta+\theta)^{\frac{2\gamma}{2-2\gamma-\beta}} \right)}{\left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} (1+c\theta+\theta)^{\frac{2-\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} (1+c\theta+\theta)^{\frac{2-\beta}{2-2\gamma-\beta}} \right)^{\frac{\gamma}{2-\beta-\gamma}}} \right. \\
 &\quad \left. \frac{\alpha_s^{\frac{\gamma}{2-\beta-\gamma}} \left(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (1+2c\theta)^{\frac{2\gamma}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} (1+2\theta)^{\frac{2\gamma}{2-\beta-2\gamma}} \right)}{\left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} (1+2c\theta)^{\frac{2-\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} (1+2\theta)^{\frac{2-\beta}{2-2\gamma-\beta}} \right)^{\frac{\gamma}{2-\beta-\gamma}}} \right]
 \end{aligned}$$

The last expression will be positive if

$$\begin{aligned}
 & \frac{\left(\frac{2^\beta}{\beta^{2\gamma}}\right)^{\frac{1}{2-\gamma-\beta}} \left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} (1+c\theta+\theta)^{\frac{2\gamma}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} (1+c\theta+\theta)^{\frac{2\gamma}{2-2\gamma-\beta}}\right)}{\alpha_s^{\frac{\gamma}{2-\beta-\gamma}} \left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} (1+c\theta+\theta)^{\frac{2-\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} (1+c\theta+\theta)^{\frac{2-\beta}{2-2\gamma-\beta}}\right)^{\frac{\gamma}{2-\beta-\gamma}}} \\
 & > \frac{\left(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (1+2c\theta)^{\frac{2\gamma}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} (1+2\theta)^{\frac{2\gamma}{2-\beta-2\gamma}}\right)}{\left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} (1+2c\theta)^{\frac{2-\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} (1+2\theta)^{\frac{2-\beta}{2-2\gamma-\beta}}\right)^{\frac{\gamma}{2-\beta-\gamma}}}
 \end{aligned}$$

and negative if not. There are parameter values such that the left-hand-side is greater and other parameter values such that the right hand side is greater. For example, for all values of $\alpha_1, \alpha_2, \theta, \beta, \gamma, c$: $\left(\frac{2^\beta}{\beta^{2\gamma}}\right)^{\frac{1}{2-\gamma-\beta}} / \alpha_s^{\frac{\gamma}{2-\beta-\gamma}}$ as a function of $\alpha_s \in \mathfrak{R}^+$ is surjective on positive real numbers and the remaining two fractions are also positive. So the left-hand-side is greater for some values of α_s and the right-hand side greater for other values of α_s .

□

Chapter 2

Parents' Employment and Child Outcomes: Experimental Evidence from Ethiopia

Abstract

Parent's employment is frequently championed as the best way to improve the lives of children in poor countries and mother's employment argued to especially benefit daughters. Exploiting a field experiment in rural Ethiopia that randomized long-term job offers, the first of its kind, this paper presents direct causal evidence on the impact of a parent's employment on children's lives, focusing primarily on time use. Daughters take over the majority of house-work left undone when mothers get employed and end up spending 24 percent less time in school, but daughters' time use is unaffected by father's employment. An increase in sons' school-time of about ten percent when a mother *or* a father gets employed appears to be due to higher household income. The auxiliary predictions of a collective framework in which female time use substitution is key to household employment and schooling decisions find empirical support. The higher the proportion of daughters, the less negative the impact of mother's employment on a given daughter's school-time and the more likely that a mother seeks employment. The more weight parents attach to daughters' welfare, the less negative the impact of mother's employment on daughters school-time and the less likely that a mother seeks employment.

"...when women are...empowered to participate in...the workplace...children and families benefit. Both boys are and girls are more likely to have access to adequate nutrition...and education", 2010 UNICEF "State of the World's Children Report", p. ii

"Women's lack of economic empowerment...has a host of other negative impacts, including less favorable education and health outcomes for children", 2006 "Gender Equality as Smart Economics: A World Bank Group Gender Action Plan", p. 2

1. Introduction

The quotes above illustrate a prevailing view among policymakers which sees the creation of job opportunities for parents - especially mothers - as a quintessential tool for improving the lives of children in poor countries. The view appears to be based in large part on extrapolation of findings from studies of the effects in the household of increases in unearned income. Relying on such extrapolation may be adequate if the dominant household models - in which children typically appear only as an expenditure category for the decision-making parents¹ - provide an accurate picture of a poor country household. If instead there is substitution between parents' and childrens' time use, then employment may be a fundamentally different "treatment" than pure income transfers due to its implications for the employed parent's time use. In that case the lack of causal evidence on the consequences for children of parent's employment is a problematic gap in the literature on poor countries.

Taking advantage of a field experiment that randomized long-term job offers this paper presents direct evidence on the impact of a parent's employment on children's lives. Five Ethiopian flower farms agreed to allocate fall 2008 job offers through a lottery system. The experiment was "natural" in the sense that parents sought employment in the exact same way they would have done in the absence of the research team. Because households thus themselves determined if the mother or the father applied, I analyze the two sub-samples separately.

The farms were willing to randomize job offers because open positions attracted large numbers of mostly inexperienced applicants and screening was difficult. Before the lottery took place, enumerators surveyed acceptable applicants. Winners and losers were re-surveyed five to seven months after employment commenced. The randomization was effectively stratified on gender.

The main results are as follows. As daughters take over house-work left undone when a mother gets employed, their school-time falls by 24 percent per week. Daughters' time use is unaffected by father's employment. An increase in sons' school time of about ten percent when a mother *or* a father gets employed appears to be due to higher household income; sons' house-work time is unaffected by parents' employment.

After documenting the impact of parents' employment on childrens' time use, I present a simple collective framework in which each parent attaches weight to daughters' well-being

¹A modeling choice that may be sensible for the rich country settings for which the models were developed.

and daughters derive utility from going to school, but only females can do house-work (Chiappori, 1988, 1992). The framework highlights the variables upon which heterogeneity in the response to mother's employment is likely to depend if the primary underlying force is time use substitution between mothers and daughters.

Testing the framework's predictions, I find that, (1) the higher the proportion of daughters - a variable that is shown to be exogenous in the sample studied - the less negative the impact of mother's employment on a given daughter's school-time, (2) the greater the weight attached to daughters' well-being, the less negative the impact of mother's employment on a daughter's school-time, and (3) the greater the initial bargaining power of the mother, the greater the reduction in daughters' school-time when mothers get employed. Daughters themselves appear to have little influence over the change in their time use when mothers get employed.

Interestingly, selection into mother's versus father's employment appears to depend on the same covariates, providing further evidence of the importance of female house-work substitution.

These results have important implications for the design of employment programs and for how selection into parent's employment and its effects in the household should be modeled. If full-time school enrollment is not universal, explicitly accounting for children's time use is important. In situations where the house-work necessary to run a household is time-consuming, the substitutability between parents' and children's effort introduces a potential trade-off between parents' and children's preferences. If house-work is effectively gender-specific, then the conventional wisdom - that economically empowering mothers is of greater benefit to daughters than empowering fathers - is not necessarily the full story when it comes to parent's employment, even if mothers weigh daughters' well-being more than fathers do. The reason is that mothers may face a trade-off between own and daughters' time use that fathers do not. If female participation in the market economy over time influences the norms governing the division of labor in the household, then the longer-term effects of mother's employment may differ from those observed here, but such norms are likely slow to change.

This paper builds on and extends the overlapping literatures on adult employment, child labor and schooling, and intra-household decision-making in poor countries. Causal evidence on the effects in the household of long-term parental employment in poor countries is to my knowledge largely absent, credible exogenous variation in employment rarely being available. Indirect inference - for example on the basis of findings from studies of unearned income - has been attempted, but there are good reasons to study parent's employment directly. Beyond the implied time use reconfiguration, employment may for example affect the two parents' relative bargaining power differently than government transfers or income from other sources do. This paper presents the first experimental evidence on the effects in the household of a parent's long-term employment.

Children's time use is one of the primary determinants of human capital accumulation and child well-being. The degree of substitutability (or complementarity) between parents' and children's time use is therefore important. Several existing studies find correlations between a mother's employment status and children's time use in poor countries (for example

Francavilla, Giannelli, and Grilli, 2010; Connelly, DeGraff, and Levison, 1996). Doran (2008) convincingly shows that adults in Mexico work more when children work less due to an exogenous increase in time spent in school. But his focus is on paid child labor; though understudied in the literature due to a lack of data child house-work is much more common than paid work in most of the developing world, and the effect of parents' time use on children's time use is typically of greater relevance for policy than the converse. Gender-specificity of house-work in combination with the typically greater time requirements of "female" responsibilities may be a particularly important though often overlooked form of son favoritism, especially because child labor and schooling are negatively related (Boozer and Suri, 2001; Ravallion and Wodon, 1999). I take advantage of an exogenous increase in mother's and father's work hours to provide causal evidence on time use substitution between mothers, fathers, daughters and sons.

Existing evidence on the relationship between child labor and household income and wealth is mixed (Basu, Das, and Dutta, 2010; Edmonds and Pavcnik, 2005; Boozer and Suri, 2001). Bar and Basu (2009) argue that an inverted-U relationship can arise because of missing labor markets for children²: the results in this paper suggest that missing labor markets for adults can also lead to a range in which child labor may appear to be increasing in parents' income. As formal employment opportunities arise for mothers, daughters may be forced to take over house-work.

The preferences of children and parents are not perfectly aligned, even if parents are partially altruistic. An important question is how much influence children have over their own lives: the review in Edmonds (2008) argues that our almost complete lack of knowledge about parent-child agency and who makes child time use decisions is the most pressing issue in the literature on child labor. This paper's results indicate that the reconfiguration of a daughter's time that occurs when a mother gets employed in rural Ethiopia is decided by parents, primarily mothers, while daughters themselves have little influence over the change in their time use.

The paper is organized as follows. In section 2, I present the setting and the experiment. The reduced form time use estimates are in section 3. In section 4, I present a simple theoretical framework of household work and schooling decisions that illustrates the forces that underlie the results in section 3, and derive auxiliary predictions. The predictions are tested in sections 5. Section 6 provides further evidence on how time use decisions are made and section 7 analyzes selection into employment. Section 8 concludes.

2. The Setting

2.1 Floriculture in Ethiopia and research design

With a GDP per capita of approximately US\$700 at ppp, Ethiopia is one of the poorest countries in Africa. The adult literacy rate is 36 percent and 45 percent of illiterate young people - a category that encompasses the majority of this project's sample - are underemployed (International Labor Organization, 2005; Unicef, 2011). 84 percent of the population

²That is, poor parents may prefer their children to work but opportunities for child labor arise only as households' wealth rises.

is found in rural areas of the country, where most households engage in smallholder agriculture (Economist Intelligence Unit, 2008).

Growth in the commercial floriculture sector in Ethiopia has been explosive in recent years, fueled in part by government incentives and in part by the abundant availability of cheap land and labor in rural areas. In 2008, 81 flower farms employed around 50,000 unskilled workers. Most flower farm workers work in greenhouses, growing and harvesting flowers, or in “packhouses”, packaging flowers and preparing them for shipping. Over 70 percent of flower farm workers are women (Gebreyesus and Iizuka, 2010).

Hiring on Ethiopian flower farms typically takes place in October and November, before the main growing and harvesting season. The supervisors on five flower farms agreed to randomize job offers during fall 2008 because of an unusual situation in the labor market for flower farm workers at the time.³ Because comparable jobs were seldom available in the areas suitable for flower growing, applicants almost always outnumbered the positions to be filled by large margins. Ethiopian flower farms - still getting to grips with cost components significantly larger than labor and with little ability to predict the productivity of the mostly uneducated, illiterate and inexperienced applicants - did not prioritize optimization of the unskilled workforce (Gebreyesus and Iizuka, 2010).⁴ Because supervisors were already allocating job offers relatively arbitrarily when approached by the research team, explicit randomization was a modest procedural change.

When Ethiopian flower farms hire, word is typically spread in nearby villages. Job-seekers arrive at the farm on announced “hiring days”. At the participating farms, supervisors first excluded any unacceptable applicants. A team of enumerators then carried out the baseline survey with the remaining applicants. Finally, the names of the number of female and male workers to be hired were drawn randomly from a hat. The full sample thus consists of 527 households in which at least one spouse applied to a flower farm job and was deemed acceptable for hiring. There are 346 women in the sample and 188 men: in almost all cases one of two spouses applied.⁵

We attempted to re-interview everyone in the treatment and control groups 5 - 7 months after employment commenced. Because few farms were hiring workers in the season that followed the randomization, only 6 re-interviewed individuals in the control group had managed to obtain employment. Careful tracking procedures led to a re-interview rate of 88 percent and no statistically significant differential attrition.⁶

³Two additional farms were originally part of the sample but were dropped before the baseline survey was completed and the randomization took place.

⁴Labor costs make up a relatively small portion of flower farms' total costs (Gebreyesus and Iizuka, 2010). The reason why wages for flower farm workers did not fall is an important unanswered question for future research.

⁵Most of the farms had informal rules against hiring couples, but in seven households both spouses applied. I exclude this subsample from the econometric analysis.

⁶Differential attrition was an ex ante concern because control group individuals had to be tracked down in their home locations, whereas hired individuals could be re-interviewed on the farms. But mobility is low in the relevant regions, and extensive efforts were made to keep contact with the sample individuals. When asked about the treated sample workers not present at follow-up, supervisors reported that most were on

2.2 The sample

Almost all the job-seekers are parents: the focus here is on the effects of a parent's employment for children in the household. The field experiment was "natural" in the sense that job-seekers applied to a position in the exact same way that they would have done in the absence of the research team.⁷ Because households themselves selected into the "mother applied" and "father applied" groups, I analyze the two sub-samples separately. In section 7, I analyze selection into the two groups. Tasks performed by men and women on flower farms can differ somewhat so that male and female applicants did not compete for the same jobs. The randomization was thus effectively stratified on gender.

The research design allowed us to survey only one member of the household, the applicant. A limitation of the data is thus that questions about, for example, the time use of other household members were answered by the applicant. Both unemployed and employed parents appear very knowledgeable about their families' time use, however. I discuss potentially resulting biases below.

Summary statistics are presented in table 2. As expected, there are no statistically significant differences between the treatment and control groups in household characteristics, nor in the outcomes analyzed below, in either of the two sub-samples. Literacy rates are low, especially for women. Income and wealth indicators, such as the material that the applicant's floor is made of, indicate the severe poverty of the sample. It appears that individuals are unlikely to apply to a flower farm job if they have infants or very young children.

In addition to socioeconomic background questions, the survey included a basic expenditure module and questions about the time use of the household members. Because, as is typical for Ethiopia, some families in the sample are very large, average expenditure and time use values for younger sons and younger daughters were recorded, in addition to individual values for the oldest son and the oldest daughter.

The primary focus in this paper is on intra-household time use substitution. Table 3 lays out the time use of an average mother, father, daughter and son, in each of the two sub-samples, at baseline. Fathers spend more hours than mothers on paid work, although both fathers and mothers appear to be underemployed. Notably, mothers spend almost thirty more hours on house-work per week than fathers do. Agricultural work is included in the definition of house-work: the figures in table 3 essentially suggest that men in rural Ethiopia have a fair amount of leisure time, at least at the time of our baseline survey.⁸

Daughters are on average 9.21 years old and sons on average 9.70 years old at baseline, but there is substantial variation in the age of children in the sample. Sons spend on average 12 hours per week in school. Daughters spend 12 hours per week in the sub-sample in which

leave. Some had been fired for disciplinary issues, a few had left because of health problems, and one had moved to Addis Ababa.

⁷The applicants and most members of the supervisor and enumerator teams were not aware of the explicit randomization, but were likely aware of the generally arbitrary character of hiring procedures.

⁸The primary responsibility for cultivating household plots lies with men, but the required effort is concentrated in the harvesting periods. This project was carried out in between the two main harvests in Ethiopia.

mothers applied to a flower farm job, and seven hours in the sub-sample in which fathers applied. Travel times and the fixed cost of attending school on a given day are significant: many school-children in the sample did not attend school every day of the week.⁹ Parents face both financial and opportunity costs of sending their children to school. On average daughters do seven hours of house-work per week and sons three hours, but some children spend significantly more time on house-work.¹⁰

3. Parent's Employment and Child Time Use

3.1 What changes when a parent gets employed

Let Y_{it} be an outcome of interest, X_{it} controls (included for precision), φ_j a farm fixed effect, T_{it} household or individual i 's treatment status at time t so that $T_{it} = 1$ signifies that the applying parent was chosen for employment, and ε_{it} an iid error term. Given random assignment, consistent estimates of a treatment effect of interest, $\beta = \mathbf{E}[Y_i|T_i = 1] - \mathbf{E}[Y_i|T_i = 0]$, can then be obtained through simple OLS regressions of the form

$$Y_{ijt} = \alpha + \varphi_j + \beta T_{it} + \theta X_{it} + \varepsilon_{ijt}$$

Before analyzing the impact on childrens' time use, it is important to lay out how a household's income and parents' time use are affected when a mother or father gets employed. Note first that the applicants consider the flower farm positions to be long-term. While workers are aware that the duration of flower farm jobs can be uncertain due to seasonal and demand variation, 84 percent of the treatment group at follow-up expected to be working at the flower farm a year later.

The applicants randomly chosen for employment were paid 302 Ethiopian Birr (approximately US\$30 in 2008/9) per month on average, which represents a 154 percent increase in the average earned income of mothers that applied and a 41 percent increase in the earned income of fathers that applied. Total household income increased by 28 percent in households in which applying mothers were chosen for employment and by 25 percent in households in which applying fathers were chosen for employment.

Flower farm employment entails 47 work-hours per week on average, typically six days of full-time work. Employment is estimated to decrease the time that mothers devote to house-work by 14 hours or 37 percent per week. Father's time use is essentially unaffected by mother's employment.¹¹ There is no significant impact of father's employment on a his hours of house-work, nor on the time use of the mother.

3.2 Effect of a parent's employment on children's time use

⁹Children in rural Ethiopia spend substantial time traveling to and from school. In the Ethiopian Rural Household Survey, children are reported to spend 45 minutes on average traveling to and from school. Transport time was not recorded in the survey used here but appears to have been higher for many children in the sample.

¹⁰The house-work required of boys may be somewhat greater in social strata or rural areas with more livestock.

¹¹There is in fact a marginally significant increase in the paid work done by fathers of 0.4 hours per week when mothers get employed.

Childrens' time use is a primary determinant of human capital accumulation and also has direct welfare consequences of interest. We have seen that parental employment involves a significant increase in household income, a shift in parents' *relative* incomes towards the newly employed spouse, and, in the case of mother's employment, a large decrease in the time the employed parent is able to devote to house-work. If childrens' time use responds to household income, the relative economic position of the mother and father, or parents' time use, sons' and daughters' activities may thus be affected when a parent gets employed.

I now analyze the impact of parent's employment on children's time use. In table 4, I aggregate the time that all daughters and all sons, respectively, spend in a given activity. Children aged 5 or older are included in the schooling regressions. Daughters spend six fewer hours in school per week when mothers get employed, a highly significant decrease of 24 percent. The reason is a need to spend more time helping out at home: daughters' spend nine, or 48 percent, more hours on house-work when mothers get employed.

Sons' school-time, on the other hand, increases by two-and-a-half hours, or nine percent, when mothers get employed. This increase appears to be unrelated to time use substitution: sons' house-work time is unaffected by mothers' employment.

The only significant effect of father's employment on children's time use is an increase in sons' school-time of three-and-a-half hours, or 11 percent.

Recalling that children's time use was reported by the applicant, a potential alternative interpretation of the results in table 4 is that employment leads to systematic changes in the accuracy of a treated parent's beliefs about children's time use rather than, or in addition to, actual changes in children's time use. While I cannot entirely rule out such a possibility, comparison of the answers given by mothers and fathers in the seven households in which both parents applied suggests that any resulting biases are likely to be minor. First, at baseline mothers and fathers have remarkably similar beliefs about childrens' time use. Second, at follow-up there were only negligible changes in the difference between the answers given by a spouse who was randomly chosen for employment and the answers given by his or her spouse who also applied but was not chosen employment.¹²

Table 5 investigates the impact of mother's employment on daughters' time use in more depth. The dependent variables in columns 1 - 6 are indicators that take value 1 if the household-member(s) primarily responsible for a given house-work task is/are daughter(s). Mother's employment significantly increases the probability that daughters are responsible for fetching water, grinding grains / cooking, cleaning / washing / ironing, food shopping, and caring for (other) children. In panel B, I separately analyze how mother's employment affects the extensive and intensive margin of the oldest and younger daughters' school-time. The oldest daughter is most affected: the probability that she is enrolled in school decreases by 11 percentage points, or 12 percent, and her hours of schooling, conditional on being enrolled, fall by almost four hours, or 24 percent, when the mother starts working. The

¹²For example, at baseline the average ratios of the mother's and the father's beliefs about daughters' and sons' school hours were 0.92 and 0.91 in the households where both parents applied. Those ratios changed little in the follow-up survey, regardless of whether none of the applying spouses were hired, only the mother was hired, or only the father was hired.

estimated effect on the probability that younger daughters are enrolled in school is also negative, but not significant. There is a marginally significant 15 percent decrease in the school-hours of enrolled younger daughters when mothers get employed.

The evidence in tables 4 and 5 provides a clear picture of intra-household time use substitution in rural Ethiopia. House-work substitution between male and female members of the household is limited. Part of the reason why sons are not required to take over duties from the father when he gets employed appears to be that fathers spend little time on house-work in the first place. Even in households in which fathers spent substantial time on house-work before applying to a flower farm, sons do not take over those responsibilities when the father starts working, however. Moreover, sons' school-time increases both when a mother gets employed and when a father gets employed. The increase in schooling thus appears unrelated to the shift in parents' relative bargaining power that likely follows employment. The evidence we have seen indicates that sons' school-time increases when a parent gets employed due to the resulting increase in household income. Primary school (grades 1 - 8) is supposed to be free in Ethiopia, but parents face costs associated with uniforms, material and clothing for school - costs that can be significant for households as poor as those in this paper's sample.

The picture for daughters is a different one. The results in tables 4 and 5 indicate that daughters take over several of the mother's house-work tasks when she gets employed and end up making up for about two thirds of the decrease in the mother's house-work time. To do so, daughters are forced to attend school less. In contrast, daughters' schooling is unaffected by father's employment. These results leave little doubt that intra-household time use substitution is key to schooling outcomes in rural Ethiopia. The evidence in table 5 that some daughters drop out of school entirely when mothers get employed is particularly worrisome because school attendance is path-dependent: it may be difficult for a daughter to return to school later on once she has dropped out. The reason why daughters, unlike sons, do not attend school more when fathers get employed is not immediately clear. It may be that parents have "lexicographical" preferences over childrens' schooling and prioritize sons before daughters. It is also possible that fathers weigh sons' schooling relative to daughters' schooling more than mothers and that employment increases fathers' influence over child schooling decisions.

In the next section I present a simple theoretical framework of the household that captures the time use effects we have seen so far and derive auxiliary predictions that further illustrate the trade-offs that arise when the house-work of different familymembers is substitutable.

4. Theoretical Framework

I take a simple version of Chiappori's collective model (1988, 1992) as a starting point and illustrate the additional trade-offs that arise in a situation with gender-specific intrahousehold labor substitution. I model a household with no sons and introduce a house-work constraint - a fixed amount of house-work H which must be carried out by the females of the household. I take bargaining power as fixed in the current period but assume that relative incomes may influence bargaining power over time. The time use and expenditures of the parents and

daughters are then determined so as to maximize a weighted average of the parents' utility functions, where the weights are given by the two parents' respective bargaining power.

Define the utility of a wife/mother, husband/father, and each of n daughters as:

$$\begin{aligned} U_w &= u_w^c(c_p) + u_w^l(l_w) + u_w^h(h_w) + n\alpha_w U_d(\cdot), \\ U_h &= u_h^c(c_p) + u_h^l(l_h) + n\alpha_h U_d(\cdot), \\ U_d &= u_d^c(c_d) + u_d^f(f_d, s_d, h_d) + u_d^h(h_d). \end{aligned}$$

u_i^t represents the utility of family member i (p is for "parent") with respect to variable t , where c, h, l, f represent non-food consumption, house-work, leisure, and food consumption. α_i represents the weight parent i attaches to a daughter's utility; I assume $0 < \alpha_i < 1/n$ and $\alpha_h < \alpha_w$. Also assume $c_d = g(y_w, y_h)$. Schooling is assumed to be free, and the arguments of u_d^f ("food utility" or hunger) are assumed to enter separately (i.e., the cross partial derivatives of u_d^f are zero).

θ_i is the bargaining power of parent i : $\theta_i(m_w, m_h)$, where m_i is the time spent working (in hours per month) of parent i , and $\theta_i(m_w, m_h)$ is increasing in m_i and $\theta_w + \theta_h = 1$. Also define $\theta_d = \theta_w\alpha_w + \theta_h\alpha_h$ as the (indirect) bargaining power of an individual daughter. To begin with, I take m_w, m_h and thus $\theta_i(m_w, m_h)$ as fixed. The household (HH) then maximizes:

$$\begin{aligned} U_{HH} &= \theta_w U_w + \theta_h U_h \\ &= \theta_w [u_w^c + u_w^l + u_w^h] + \theta_h [u_h^c + u_h^l] + n(\theta_w\alpha_w + \theta_h\alpha_h) [u_d^c + u_d^f + u_d^h] \end{aligned}$$

subject to the following constraints:

- (i) Budget constraint: $y = y_w + y_h = w_w m_w + w_h m_h = 2c_p + nc_d + nf_d$
- (ii) House-work constraint: $H = h_w + nh_d$
- (iii) Time constraints: $1 = m_w + l_w + h_w, \quad 1 = m_h + l_h, \quad 1 = h_d + s_d$

where the time available to each household member is normalized to 1.

Focusing on the effect of mother's employment on daughters' time use, the model has the following predictions:

Prediction 1 (Mother's Employment): Daughters' school time is decreasing in the mother's formal work hours.

The intuition is of course that employed mothers have less time available for house-work so that daughters must pick up the slack, leaving less time for school.

Prediction 2 (Mother's bargaining power): Daughters' school time is decreasing in the mother's bargaining power.

Mothers that prefer both formal work and leisure to house-work spend less time on house-work as their bargaining power rises. Daughters pick up the slack and end up spending less time in school.

Prediction 3 (Mother's weight on daughters' well-being): Daughters' school time is increasing in the mother's weight on their well-being.

Mothers that care more about their daughters do more house-work so as to leave more time for daughters to attend school.

Prediction 4 (Number of daughters): Individual daughters' school time is increasing in the total number of daughters in the household.

The more daughters who can participate in house-work activities there are, the more time is left for each daughter to attend school. This prediction would not be unambiguous if costs of schooling were explicitly accounted for in the model; in that case the effect of a given daughter attending school on the money left for other daughters' schooling would need to be accounted for. But note that the model suggests that the presence of adult females - who do not attend school - other than the mother should also increase each daughter's school time. The intuition of prediction 4 can thus be tested in two alternative ways.

5. Heterogeneity in the Impact of Parent's Employment on Child Outcomes

The magnitude of the effect of mother's employment on daughters' schooling is far from uniform across households. In table 6, I test predictions 2 - 4 by interacting proxies for the household characteristics that the framework above suggests should induce heterogeneity in the response of daughter's schooling to mother's employment with the treatment. While the tests are thus indirect - focusing on the mitigating (or exacerbating) effect of a given covariate on the response to employment rather than the covariate itself - they provide informative evidence on the intuition behind each prediction.

As a proxy for mother's bargaining power I use her share of baseline earned income, a "distribution factor" commonly used in empirical research that has been shown to influence individuals' control over household decisions (for example, [Bonke and Browning, 2009](#); [Dauphin, El Lahga, Fortin, and Lacroix, 2008](#); [Browning and Chiappori, 1998](#)). In the first two columns we see that the impact of mother's employment on daughters' house-work is two hours bigger, and the impact on daughters' school-time three hours more negative, for every one standard deviation increase in the mother's initial bargaining power. Prediction 2 thus finds empirical support: it appears that, relative to fathers, mothers prefer daughters to take over more house-work when mothers get employed and are left with less available time to spend on house-work duties.

In columns 3 and 4 I proxy for mother's weight on daughters' well-being with daughters' initial expenditure share. Daughters' expenditure share will reflect a combination of mother's and father's weight on daughters' well-being (in addition to other factors). The interaction

between the treatment and daughters' weight is negative but not significant in the house-work regression. The decrease in schooling is significantly smaller, by 3 hours, for every one standard deviation increase in daughters' expenditure share, as predicted by prediction 3. It thus appears that the negative effect of mother's employment on daughters' schooling is dampened in households that value daughters more.

If the negative effect of mother's employment on daughters' schooling is due to gender-specific time use substitution, as I have argued, then the effect for a given daughter should be smaller the more females there are in the household. The reason is of course that the house-work previously done by the mother can be spread across females. In columns 5 and 6 we see that, while the impact of mother's employment on daughter's time use remains significant even when other adult women are present, the impact is greatly reduced in such households.¹³ In columns 9 and 10 we see that the effect on the oldest daughter's house-work time is two-and-half hours reduced, and the effect on the oldest daughter's school-time one-and-half hours reduced, for every additional daughter that is present in the household. It is thus clear that other female household-members share the extra house-work burden that mother's employment implies, as the framework predicts.

In sum, the evidence in table 6 indicates that the extent of time use substitution between mothers and daughters depends on parents' preferences and on the number of females available for house-work. The framework above can thus account for important determinants of heterogeneity in how children's time use responds to a parent's employment. In the next section I explore how child time use decisions are made.

6. Household Decision-making and the Impact of Mother's Employment on Daughters' Time Use

If time use substitution between parents and children is significant and parents are imperfectly altruistic, then parent-child agency issues are likely important for children's well-being (see [Edmonds, 2008](#)). The framework above follows the literature in modeling parents as the relevant decisionmakers in the household. The evidence in table 6 that mother's bargaining power influences the amount of house-work that daughters take over when mothers get employed suggests that mothers have influence over daughters' time. But should daughters themselves and fathers also be seen as decisionmakers participating in decisions about daughters' time use?

In table 7, I interact proxies for the mother's, father's and the oldest daughter's preferences - answers to survey questions about each of the three family-members' attitude towards girls' schooling - with the treatment.¹⁴ In households in which the mother considers girls' schooling more important, the negative effect of mother's employment on daughters' schooling is significantly smaller. The father's attitude towards girls schooling appears to have less

¹³The presence of other adult men in the household has no effect on the impact of mother's employment on daughters' time use.

¹⁴The survey questions asked the applicant "On a scale from 0 to 10, how important would [you/your spouse] say that it is that girls attend school?" and "On a scale from 0 to 10, how much would you say that your oldest daughter likes/would like going to school?" (this question was not asked for younger daughters).

influence on time use substitution between mothers and daughters, and a daughter's own preferences have no significant effect on the amount of house-work she is expected to take over when her mother gets employed. It thus appears that daughters in rural Ethiopia have little control over their own time use in times of need.

7. Parents' Employment Decisions and Intra-household Time Use Substitution

Analysis of how households select into mother's versus father's employment is important in its own right but also represents a powerful auxiliary test of the main message of the framework above. If, as this paper has argued, a key determinant of rural Ethiopians' time use is gender-specific, intra-household labor substitution, then household characteristics that influence the impact of mother's and father's employment on other family-members - such as the gender composition of a couple's children - should also influence selection into mother's versus father's employment.

The sample analyzed consists of households in which either the mother or the father applied to a flower farm. To pool the two sub-samples and explore selection into the two groups, we must thus assume that, in for example a household in which the mother applied, the father would have applied had the mother not done so. While this assumption is ultimately untestable, it is arguably reasonable. As noted, there were only seven households in which both spouses applied - for most households the relevant choice options appear to have been for one or none of the two spouses to apply. There are few households in the sample in which the spouse of the applicant was already formally employed. In table 8 I investigate the comparability of the two sub-samples. Excluding the right-hand-side variables that the framework predicts *should* influence selection into the two groups (which I analyze below), the only significant difference is that husbands are one year older in households in which the mother applied. I thus control for husband's age in the analysis below.

As we saw in table 6, perhaps the most important variable governing heterogeneity in the impact of mother's employment on daughters' time-use is the gender composition of the couple's children because the presence of more daughters means that house-work can be shared between more hands. As such, we would expect the number of daughters to have an important influence on selection into mother's versus father's employment. But testing for a causal relationship is possible only if the number of daughters is exogenous conditional on the total number of children. If parents follow differential stopping rules - that is, if the probability of having another child depends on the gender composition of existing children - then the number of daughters is not exogenous even conditional on family size, as pointed out by Clark (2000) and discussed in detail in Washington (2008). It turns out that parents in the sample do not follow such stopping rules: neither a variable equal to the total number of children, nor dummies for having a given number of children, predict the proportion of daughters, as seen in table 9. The explanation may be that desired family sizes in rural Ethiopia are so large that almost all couples have one or more sons through "natural" fertility behavior. Parents with son preference typically want "at least X number of sons" (Clark, 2000), where X is a positive but relatively low number.

We can thus test if the gender composition of a couple's children has a causal effect on the probability that a mother (rather than a father) seeks employment. I do so in table 10, including interactions with the proxies for mother's weight on daughters' well-being and mother's bargaining power to mirror the heterogeneity regressions in table 6. The selection analysis results are supportive of the idea that female time use substitution is key to household employment and schooling decisions in Ethiopia. For example, one additional daughter increases the probability that the mother applies by 8 percentage points, or 12 percent, controlling for the total number of children, in households with low weight on daughter well-being and low mother's bargaining power.

The results also indicate that the higher the weight on daughters' well-being, the lower the influence of the number of daughters on the couple's employment decision. The reason appears to be that highly valued daughters are expected to take over less household work when mothers get employed.

Finally, mother's bargaining power at baseline has a marginally significant positive effect on the influence of the number of daughters on the probability that the mother applies. The presence of daughters has a direct influence on the mother's well-being under mother's employment relative to father's employment because a mother can likely decrease her time spent on house-work when employed more when more daughters are present. A father's well-being under mother's versus father's employment may, in contrast, be less dependant on the gender composition of the couple's children because "male" house-work is less time-consuming. It appears that greater bargaining power for the mother therefore increases the weight given to the gender composition of the couple's children when the employment decision is made.

The findings in table 10 thus suggest that parents take into account substitutability between a mother's and daughters' time use when making adult employment decisions. If daughters taking over house-work duties when mothers get employed is difficult to avoid, it may be that the best way to take daughters' well-being into account is at the employment decision stage.

8. Conclusion

In this paper, I argue that parents' time use is likely to influence how much time children in poor countries spend in school when house-work is time-consuming, due to gender-specific intra-household labor substitution. Exploiting a field experiment in rural Ethiopia that randomized long-term job offers on five flower farms, I show that daughters spend significantly less time in school when mothers work because daughters take over house-work duties from employed mothers, but not from employed fathers. Sons spend significantly more time in school when either parent works: greater labor income enables parents to cover the costs of schooling for prioritized children, and the amount of house-work left undone by employed fathers is limited. As predicted by a simple collective framework of household time use with gender-specific house-work duties, the decrease in (individual) daughters' school-time when mothers get employed is greater the fewer daughters there are, the less parents value daughters' well-being, and the greater the bargaining power of the mother. While mothers

may “care” more about daughters’ well-being than fathers do, mothers nevertheless prefer daughters to do more house-work than fathers do because “female” house-work duties must otherwise be carried out by mothers. Finally, the factors that influence the relative impact of mother’s and father’s employment on household utility unsurprisingly also influence selection into mother’s versus father’s employment. For example, the mother is more likely to seek employment the higher the proportion of children that are female.

Overall this paper’s results highlight the centrality of intra-household time use substitution as a determinant of children’s human capital accumulation in poor countries. Modeling children as an expenditure category without explicitly accounting for time use substitution is likely to provide a misleading picture of reality in situations where house-work is time-consuming. It is thus not surprising that the commonly held view that mother’s employment will especially benefit daughters - an argument that appears to be based on extrapolating evidence from studies of unearned income - may be inaccurate in some contexts. It also appears that the gender-specificity of intra-household labor substitution in combination with the typically greater time requirements of “female” house-work duties may represent a form of son bias that becomes increasingly important as more adults enter the formal workforce.

Tables

Table 1
Stated time use substitution

After you started working at the farm, did you continue to do the same amount of house-work or did you do more, or less?	Female worker:		Male worker:		If Less: Who started doing the house-work that you were doing before?	Female worker:		Male worker:	
Less	78 %		19 %		Mother	6 %		0 %	
Same	20 %		78 %		Sister	5 %		0 %	
More	2 %		3 %		Daughter	66 %		0 %	
					Son	0 %		12 %	
					Cousin	3 %		35 %	
					No one	20 %		53 %	

Table 2: Summary statistics and randomization balance

	Panel A: Mother applied (339)			Panel B: Father applied (181)		
	Treatment (174)	Control (165)	Difference	Treatment (103)	Control (78)	Difference
Parent	0.95	0.92	0.03 (0.31)	0.96	0.99	-0.03 (0.02)
Sons	1.47	1.68	-0.21 (0.17)	2.69	2.5	0.19 (0.28)
Daughters	2.09	1.92	0.17 (0.16)	1.80	1.62	0.18 (0.25)
Has cement/wood floor	0.21	0.26	-0.05 (0.05)	0.23	0.24	-0.01 (0.06)
Household income	982.66	1041.18	-59.53 (62.45)	911.12	988.49	-77.37 (77.88)
Literate	0.19	0.15	0.04 (0.04)	0.27	0.31	-0.04 (0.07)
Spouse literate	0.28	0.31	-0.02 (0.05)	0.20	0.16	0.04 (0.06)
Age	26.67	26.16	0.50 (0.71)	29.17	29.54	-0.37 (0.91)
Spouse age	31.04	29.97	1.06 (0.72)	25.94	25.58	0.36 (0.66)
Paid work hours per month	39.44	36.50	2.93 (5.41)	67.50	68.74	-1.24 (5.30)
Spouse paid work hours per month	65.05	64.05	1.00 (3.79)	46.95	43.14	3.81 (7.13)

Standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01

Table 3
Baseline time use

	Panel A: Mother applied	Panel B: Father applied
Mother		
House-work hours per week	38 (12)	33 (11)
Paid work hours per week	9 (11)	10 (11)
Father		
House-work hours per week	5 (2)	5 (2)
Paid work hours per week	15 (8)	16 (8)
Daughter		
House-work hours per week	6 (6)	8 (7)
School hours per week	12 (6)	7 (7)
Son		
House-work hours per week	3 (2)	3 (2)
School hours per week	12 (6)	12 (6)

Standard deviations in parentheses. Child house-work hours in households with daughters/sons, child school hours in households with school-age daughters/sons

Table 4
Effect of parent's employment on child time use

Sample:	Mother applied				Father applied			
	Daughters' House-work hours per week	School hours per week	House-work hours per week	Sons' School hours per week	Daughters' House-work hours per week	School hours per week	House-work hours per week	Sons' School hours per week
Treatment	9.26*** (1.35)	-6.35*** (1.44)	0.42 (0.37)	2.42*** (0.81)	-1.84 (1.18)	0.28 (0.69)	0.34 (0.40)	3.45*** (0.99)
N	290	290	243	243	126	126	172	172
Dep. var. mean	19.23	26.31	5.73	26.55	8.38	23.75	7.29	31.25
Covariates?	YES	YES	YES	YES	YES	YES	YES	YES

Robust standard errors in parentheses. * p < 0.10, ** p < 0.05, *** p < 0.01. OLS regressions. The dependent variables are the total hours that all daughters/sons spent in the relevant activity during the previous seven days. Only households with school-age children are used in (2), (4), (6) and (8). The covariates are applicant's age and literacy, an indicator for the household having a cement/wood floor, and farm fixed effects

Table 5
Unpacking the effect of mother's employment on daughter's schooling

Sample:	Mother applied						Panel B: Effect on extensive and intensive margin of schooling					
	Panel A: Effect on daughter's house-work duties						School hours of already enrolled daughter		Younger daughters enrolled in school		School hours of already enrolled younger daughters	
Dependent variable:	fetches water	fetches firewood	does cooking	Daughter does cleaning	does shopping	does child-care	Oldest daughter enrolled in school	of already enrolled oldest daughter	Younger daughters enrolled in school	of already enrolled younger daughters		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)		
Treatment	0.09*** (0.03)	0.00 (0.03)	0.18*** (0.04)	0.14*** (0.03)	0.05*** (0.02)	0.12*** (0.03)	-0.11*** (0.04)	-3.63*** (0.71)	-0.07 (0.05)	-1.56* (0.89)		
N	290	290	290	290	290	290	260	190	143	98		
Dep. var. mean	0.37	0.48	0	0.17	0	0.20	0.91	15.23	0.94	10.42		
Covariates?	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES		

Robust standard errors in parentheses. * p < 0.10, ** p < 0.05, *** p < 0.01. Linear probability models in columns (1) - (7) and (9), OLS regressions in (8) and (10). Households with positive hours of schooling reported for oldest/younger daughters at baseline are used in (8) and (10). Households with an oldest daughter of school-age are used in (7) and (8). Average values were recorded younger children: households with younger daughters that are on average of school-age are used in (9) and (10). In (1) - (6) the dependent variables are dummies that equal one if the person primarily responsible for the relevant task is a daughter. The covariates are applicant's age and literacy, an indicator for the household having a cement / wood floor, and farm fixed effects.

Table 6
Heterogeneity in the effect of mother's employment on child time use
Mother applied

Dependent variable:	Daughters'										Oldest daughter's	
	House-work hours per week	School hours per week	House-work hours per week	School hours per week	House-work hours per week	School hours per week	House-work hours per week	School hours per week	House-work hours per week	School hours per week	House-work hours per week	School hours per week
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)		
Treatment	6.66*** (1.81)	-3.31 (2.02)	11.73*** (1.92)	-10.94*** (2.23)	9.92*** (1.52)	-6.51*** (1.57)	10.52*** (2.58)	-8.86*** (3.19)	11.50*** (1.23)	-8.35*** (1.54)		
Treatment* Mother's bargaining power at baseline	8.08*** (3.78)	-9.53*** (4.46)					7.47* (3.88)	-8.38*** (4.23)				
Treatment* Mother's weight on daughter well-being at baseline			-70.74 (43.52)	124.17*** (49.15)			-85.15*** (42.47)	131.87*** (51.28)				
Treatment* Other female present at baseline					-4.11*** (1.55)	5.09*** (1.74)	-4.40*** (1.86)	6.33*** (2.26)				
Treatment* Number of daughters at baseline									-2.59*** (0.41)	1.58*** (0.60)		
N	290	260	290	260	290	260	290	260	290	260	290	260
Dep. var. mean	19.23	26.31	19.23	26.31	19.23	26.31	19.23	26.31	19.23	26.31	9.38	12.03
Covariates?	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES

Robust standard errors in parentheses. * p < 0.10, ** p < 0.05, *** p < 0.01. OLS regressions. The covariates are applicant's age and literacy, an indicator for the household having a cement / wood floor, and farm fixed effects. Other regressors: mother's bargaining power at baseline as proxied by her income share (in (1), (2), (7) and (8)), mother's weight on daughter well-being at baseline as proxied by daughters' expenditure share (in (3), (4), (7) and (8)), an indicator for at least one other adult female residing in the household at baseline (in (5) - (8)), and the number of daughters at baseline (in (9) and (10))

Table 7
Control over daughter's time use and effect of mother's employment

Sample:	Mother applied	
	Daughters' school hours per week	Oldest daughter's school hours per week
Dependent variable:	(1)	(2)
Treatment	-22.54*** (5.07)	-8.74*** (3.01)
Treatment* Father's attitude to girls' schooling at baseline	1.10* (0.58)	0.22 (0.26)
Treatment* Mother's attitude to girls' schooling at baseline	1.50*** (0.44)	0.70*** (0.24)
Treatment* Oldest daughter's preference for attending school at baseline	0.22 (0.54)	-0.14 (0.26)
N	260	260
Dep. var. mean	26.31	12.03
Covariates?	YES	YES

Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. OLS regressions. The covariates are applicant's age and literacy, an indicator for the household having a cement /wood floor, and farm fixed effects. Other regressors: mother's and father's attitude to girls' schooling and oldest daughter's preference for attending school at baseline. Mother's attitude to girls' schooling was measured through responses to the question "On a scale from from 0 to 10, how important would you say that it is that girls attend school?", father's attitude to girls' schooling through mothers' responses to the question "On a scale from from 0 to 10, how important would you say that your spouse thinks it is that girls attend school?", and the oldest daughter's preference for attending school through mothers' responses to the question "On a scale from 0 to 10, how much would you say that your oldest daughter likes /would like going to school?"

Table 8
 Investigating the baseline comparability of households
 in which mothers vs fathers applied

	Wife applied (339)	Husband applied (181)	Difference
Has cement/wood floor	0.24	0.24	-0.00 (0.04)
Household income	1010.63	944.46	66.17 (51.17)
Wife literate	0.17	0.18	-0.02 (0.04)
Husband literate	0.30	0.29	0.01 (0.04)
Wife age	26.42	25.79	0.63 (0.54)
Husband age	30.52	29.33	1.20** (0.59)
Wife paid work hours per month	38.01	45.31	-7.30 (4.50)
Husband paid work hours per month	64.56	68.04	-3.48 (3.22)

Standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01

Table 9
Exogeneity of child gender make-up

	Proportion of Daughters	
Children	-0.00	
	(0.01)	
Children=2	-0.07	
	(0.06)	
Children=3	-0.01	
	(0.06)	
Children=4	-0.00	
	(0.06)	
Children=5	-0.01	
	(0.06)	
Children=6	-0.03	
	(0.07)	
Children=7	-0.05	
	(0.07)	
Children=8	-0.06	
	(0.09)	
Children=9	-0.02	
	(0.13)	
Children=10	-0.03	
	(0.15)	
Children=11	-0.35	
	(0.32)	
Children=13	-0.11	
	(0.16)	
Children=14	-0.21	
	(0.23)	
Children=15	-0.06	
	(0.32)	
N	493	493

Standard errors in parentheses. * p< 0.10, ** p<0.05, *** p<0.01

Table 10
 Determinants of mothers seeking employment

Dependent variable:	Mother applied
	(1)
Daughters	0.08*** (0.03)
Daughters* Mother's bargaining power at baseline	0.08* (0.04)
Daughters* Mother's weight on daughter well-being at baseline	-0.61** (0.27)
N	414
Dep. var. mean	0.65
Covariates?	YES

Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Linear probability models. The dependent variable equals one if the mother applied and zero if the father applied. The covariates are applicant's age and literacy, an indicator for the household having a cement /wood floor, and farm fixed effects. Other regressors: total number of children, mother's weight on daughter well-being at baseline as proxied by daughters' expenditure share and mother's bargaining power at baseline as proxied by her income share

Theoretical Appendix

Using the constraints, household utility can be written

$$\begin{aligned}
 U_{HH} = & \theta_w \left[u_w^c \left(\frac{w_w m_w + w_h m_h - ng(w_w m_w, w_h m_h) - n f_d}{2} \right) + u_w^l (l_w) + u_w^h (1 - m_w - l_w) \right] \\
 & + (1 - \theta_w) (1 - \theta_w) \left[u_h^c \left(\frac{w_w m_w + w_h m_h - ng(w_w m_w, w_h m_h) - n f_d}{2} \right) + u_h^l (1 - m_h) \right] \\
 & + n (\theta_w \alpha_w + \theta_h \alpha_h) \left[u_d^c (c_d) + u_d^f \left(f_d, 1 - \frac{H + m_w + l_w - 1}{n}, \frac{H + m_w + l_w - 1}{n} \right) \right] \\
 & \quad \quad \quad + u_d^h \left(\frac{H + m_w + l_w - 1}{n} \right)
 \end{aligned}$$

In the period after mother's and father's work hours are determined, w_i, n, α_i, H, m_i and θ_i are taken as exogenous. Taking first order conditions with respect to the remaining choice variables are f_d and l_w gives:

$$\begin{aligned}
 f_1 = FOC_{f_d} &= -\frac{n}{2} \theta_w \frac{\partial u_w^c}{\partial f_d} - \frac{n}{2} (1 - \theta_w) \frac{\partial u_h^c}{\partial f_d} + n (\theta_w \alpha_w + \theta_h \alpha_h) \frac{\partial u_d^f}{\partial f_d} \\
 &= -\frac{n}{2} \left(\theta_w \frac{\partial u_w^c}{\partial f_d} + \theta_h \frac{\partial u_h^c}{\partial f_d} \right) + n (\theta_w \alpha_w + \theta_h \alpha_h) \frac{\partial u_d^f}{\partial f_d} = 0 \\
 f_2 = FOC_{l_w} &= \theta_w \left[\frac{\partial u_w^l}{\partial l_w} - \frac{\partial u_w^h}{\partial l_w} \right] + n (\theta_w \alpha_w + \theta_h \alpha_h) \left[\frac{1}{n} \frac{\partial u_d^f}{\partial l_w} + \frac{1}{n} \frac{\partial u_d^h}{\partial l_w} \right] \\
 &= \theta_w \left(\frac{\partial u_w^l}{\partial l_w} - \frac{\partial u_w^h}{\partial l_w} \right) + (\theta_w \alpha_w + \theta_h \alpha_h) \left(\frac{\partial u_d^f}{\partial l_w} + \frac{\partial u_d^h}{\partial l_w} \right) = 0
 \end{aligned}$$

In the following I assume that the matrix of second-order partials of U_{HH} :

$$\left(\frac{\partial^2 U_{HH}}{\partial f_i \partial f_j} \right) = \begin{pmatrix} \frac{\partial f_1}{\partial f_d} & \frac{\partial f_1}{\partial l_w} \\ \frac{\partial f_2}{\partial f_d} & \frac{\partial f_2}{\partial l_w} \end{pmatrix}.$$

is negative semidefinite and non-singular so that

$$\begin{aligned}
 \frac{\partial f_1}{\partial f_d} &\leq 0 \\
 \frac{\partial f_1}{\partial f_d} \frac{\partial f_2}{\partial l_w} - \frac{\partial f_1}{\partial l_w} \frac{\partial f_2}{\partial f_d} &> 0.
 \end{aligned}$$

Proof of prediction 1 (Mother's Employment): Daughters' school time is decreasing in the mother's formal work hours.

We have that:

$$\begin{aligned} s_d &= 1 - h_d \\ \implies \frac{\partial s_d}{\partial m_w} &= -\frac{\partial h_d}{\partial m_w} \\ &= -\frac{\partial}{\partial m_w} \left(\frac{H + 1 - m_w - l_w}{n} \right) = \frac{1}{n} \frac{\partial l_w}{\partial m_w} \end{aligned}$$

As long as $\frac{\partial l_w}{\partial m_w} < 0$ and $\frac{1}{n} \rightarrow 0$ as $n \rightarrow \infty$, the prediction follows. □

Proof of prediction 2 (Mother's bargaining power): Daughters' school time is decreasing in the mother's bargaining power.

We have that:

$$\begin{aligned} \frac{\partial s_d}{\partial \theta_w} &= \frac{\partial}{\partial \theta_w} \left(1 - \frac{(H - h_w)}{n} \right) \\ &= \frac{1}{n} \frac{\partial (h_w)}{\partial \theta_w} < 0 \text{ if } \frac{\partial (h_w)}{\partial \theta_w} < 0 \end{aligned}$$

As long as $\frac{\partial (h_w)}{\partial \theta_w} < 0$, the prediction follows. Note that the the magnitude of the effect decreases with the number of daughters. □

Prediction 3 (Mother's weight on daughters' well-being): Daughters' school time is increasing in the mother's weight on their well-being.

We have that:

$$\begin{aligned} \frac{\partial s_d}{\partial \alpha_i} &= \frac{\partial}{\partial \alpha_i} \left(\frac{(n + 1) - H - m_w - l_w}{n} \right) \\ &= -\frac{1}{n} \left(\frac{\partial l_w}{\partial \alpha_i} \right). \end{aligned}$$

By the implicit function theorem, we have

$$\frac{\partial l_w}{\partial \alpha_i} = \frac{\overbrace{\frac{\partial f_2}{\partial f_d} \frac{\partial f_1}{\partial \alpha_i} - \frac{\partial f_1}{\partial f_d} \frac{\partial f_2}{\partial \alpha_i}}^?}{\underbrace{\frac{\partial f_1}{\partial f_d} \frac{\partial f_2}{\partial l_w} - \frac{\partial f_1}{\partial l_w} \frac{\partial f_2}{\partial f_d}}_{>0}}.$$

We have that $\frac{\partial f_2}{\partial f_d} = \frac{\partial f_1}{\partial l_w} = 0$, and, by assumption, $\frac{\partial f_1}{\partial f_d} < 0$. Further,

$$\begin{aligned} \frac{\partial f_2}{\partial \alpha_i} &= \theta_i \left(n \frac{\partial u_d^f}{\partial l_w} + \frac{\partial u_d^h}{\partial h_d} \right) \\ &= \theta_i \left(n \left[\frac{\partial u_d^f}{\partial f_d} \frac{\partial f_d}{\partial l_w} + \frac{\partial u_d^f}{\partial s_d} \frac{\partial s_d}{\partial l_w} + \frac{\partial u_d^f}{\partial h_d} \frac{\partial h_d}{\partial l_w} \right] + \frac{\partial u_d^h}{\partial h_d} \right) \\ &= \theta_i \left(n \left[0 - \frac{1}{n} \frac{\partial u_d^f}{\partial s_d} + \frac{1}{n} \frac{\partial u_d^f}{\partial h_d} \right] + \frac{\partial u_d^h}{\partial h_d} \right) \\ &< 0, \end{aligned}$$

so that

$$\frac{\partial l_w}{\partial \alpha_i} = \frac{\overbrace{\frac{\partial f_2}{\partial f_d} \frac{\partial f_1}{\partial \alpha_i} - \frac{\partial f_1}{\partial f_d} \frac{\partial f_2}{\partial \alpha_i}}^0}{\underbrace{\frac{\partial f_1}{\partial f_d} \frac{\partial f_2}{\partial l_w} - \frac{\partial f_1}{\partial l_w} \frac{\partial f_2}{\partial f_d}}_{>0}} < 0,$$

and thus

$$\frac{\partial s_d}{\partial \alpha_i} = -\frac{1}{n} \left(\frac{\partial l_w}{\partial \alpha_i} \right) > 0.$$

□

Prediction 4 (Number of daughters): Individual daughters' school time is increasing in the total number of daughters in the household.

I assume that $\frac{\partial(h_w)}{\partial\theta_w} < 0$: a wife with more bargaining power will commit more of her time to working. We have that

$$\begin{aligned}\frac{\partial s_d}{\partial n} &= \frac{\partial}{\partial n} \left(\frac{(n+1) - H - m_w - l_w}{n} \right) \\ &= \frac{n(1 - \frac{\partial l_w}{\partial n}) - (n+1 - H - m_w - l_w)}{n^2} \\ &= \frac{H + m_w + l_w - 1 - n \frac{\partial l_w}{\partial n}}{n^2} \\ &= \frac{H - h_w - n \frac{\partial l_w}{\partial n}}{n^2} \\ &= \frac{h_d - \frac{\partial l_w}{\partial n}}{n}.\end{aligned}$$

By the implicit function theorem, we have

$$\frac{\partial l_w}{\partial n} = \frac{\overbrace{\frac{\partial f_2}{\partial f_d} \frac{\partial f_1}{\partial n}}^0 - \overbrace{\frac{\partial f_1}{\partial f_d} \frac{\partial f_2}{\partial n}}^?}{\underbrace{\frac{\partial f_d}{\partial f_d} \frac{\partial l_w}{\partial l_w} - \frac{\partial f_1}{\partial l_w} \frac{\partial f_2}{\partial f_d}}_{>0}}.$$

Further,

$$\begin{aligned}\frac{\partial f_2}{\partial n} &= (\theta_w \alpha_w + (1 - \theta_w) \alpha_h) \frac{\partial u_d^f}{\partial l_w} \\ &= (\theta_w \alpha_w + (1 - \theta_w) \alpha_h) \left[\frac{\partial u_d^f}{\partial f_d} \frac{\partial f_d}{\partial l_w} + \frac{\partial u_d^f}{\partial s_d} \frac{\partial s_d}{\partial l_w} + \frac{\partial u_d^f}{\partial h_d} \frac{\partial h_d}{\partial l_w} \right] \\ &= \underbrace{(\theta_w \alpha_w + (1 - \theta_w) \alpha_h)}_{>0} \left[\frac{\partial u_d^f}{\partial f_d} \underbrace{\frac{\partial f_d}{\partial l_w}}_0 + \underbrace{\frac{\partial u_d^f}{\partial s_d} \frac{-1}{n}}_{>0} + \underbrace{\frac{\partial u_d^f}{\partial h_d} \frac{1}{n}}_{<0} \right].\end{aligned}$$

Again using inequality (??), this becomes

$$\frac{\partial f_2}{\partial n} = \underbrace{(\theta_w \alpha_w + (1 - \theta_w) \alpha_h)}_{>0} \left[\frac{\partial u_d^f}{\partial f_d} \underbrace{\frac{\partial f_d}{\partial l_w}}_0 + \underbrace{\frac{\partial u_d^f}{\partial s_d} \frac{-1}{n}}_{<0} + \frac{\partial u_d^f}{\partial h_d} \frac{1}{n} \right] < 0,$$

so that

$$\frac{\partial l_w}{\partial n} = \frac{\overbrace{\frac{\partial f_2}{\partial f_d} \frac{\partial f_1}{\partial n}}^0 - \overbrace{\frac{\partial f_1}{\partial f_d} \frac{\partial f_2}{\partial n}}^{>0}}{\underbrace{\frac{\partial f_1}{\partial f_d} \frac{\partial f_2}{\partial l_w} - \frac{\partial f_1}{\partial l_w} \frac{\partial f_2}{\partial f_d}}_{>0}} < 0,$$

which finally tells us that

$$\frac{\partial s_d}{\partial n} = \frac{h_d - \frac{\partial l_w}{\partial n}}{n} > 0,$$

giving the prediction.

□

Chapter 3

Backlash: Female Employment and Domestic Violence

with Espen Villanger

Abstract

We explore the relationship between female employment and domestic violence. Conventional economic models predict a decrease in violence when women get employed; a prediction that is central to existing anti-violence policy. Through a field experiment in Ethiopia that randomized job offers, we document a significant 13 percent increase in physical violence when women get employed, and a 34 percent increase in emotional abuse. In further analysis we find limited support for models in which violence is used as a tool to gain control over household resources. Instead it appears that it is emotionally costly to men when household roles deviate from those prescribed by gender norms, and that violence is seen as a way to restore a traditional order.

1. Introduction

Domestic violence represents a serious violation of women's rights and imposes substantial costs on society. In parts of Ethiopia, 71 percent of ever-partnered women have been physically assaulted by a male partner (Garcia-Moreno et al., 2006). In the U.S., domestic violence assault is more common than all other forms of violence combined (Tauchen and Witte, 1995: 1). But despite its prevalence throughout much of the world, the nature of physical abuse of women remains poorly understood. Little is therefore known about how to address the issue.

In this paper, we analyze the effect of female employment on domestic violence through a field experiment in rural Ethiopia that randomized job offers, the first of its kind. Conventional economic models of domestic violence are "optimistic" in the sense of predicting a decrease in abuse when women get employed; we find the opposite. We then begin to distinguish between "pessimistic" models. We find limited support for models in which violence is used as a tool to gain control over household resources, and more support for models that allow men to see violence as a way to restore their dominance in the household.

Five Ethiopian flower farms agreed to randomize fall 2008 long-term job-offers. The sample consists of 329 households in which an adult woman applied to a flower farm job and was deemed acceptable for hiring by the farm. The treatment and control groups were re-surveyed 5 - 7 months after employment commenced.

Our research design has important advantages. Because we directly vary job offers, we can attribute changes in violence to the causal effect of employment. There is to our knowledge no existing experimental evidence from poor countries on the effects of permanent female employment, by many thought to be the most effective way to reduce physical abuse. Policy and arguments are therefore made on the basis of assumptions on which clear-cut causal evidence is largely missing: the World Health Organization argues, for example, that "women's access to...employment should...be strongly supported as part of overall anti-violence efforts" (WHO, 2005: 23). In the absence of sufficient evidence, there is little consensus on which model of domestic violence best describes reality.

In the main result of the paper, we find a 13 percent increase in the probability that a woman is experiencing physical domestic violence, when she gets employed. We also find a 34 percent increase in emotional abuse, and a 32 percent increase in the number of violent incidents per month. As discussed below, the effects are unlikely to represent a change in reporting behavior.

Our results are hard to reconcile with conventional models, most of which are optimistic in the sense that employment and other forms of economic empowerment of women is predicted to decrease abuse. We thus explore the ability of more recent, pessimistic violence models to explain our findings. Authors of instrumental violence models argue that a husband may turn more violent when his wife's income goes up in order to counteract a rise in her bargaining power, or to increase the husband's slice of a bigger income pie. But there is no indication that violent husbands in our sample have greater control over household

resources, neither before nor after female employment. Alternatively, physical abuse may be seen as a way to restore a (perceived) traditional order in the household; either used by husbands to influence wives' behavior, or generating direct, expressive utility for husbands. We argue that a natural adjustment to existing expressive violence models would allow the marginal utility that a husband derives from violence to increase when he is "disempowered" by his wife's employment. Consistent with this, the increase in the incidence of violence is greater in households in which the newly employed wife was likely to end up further ahead of her husband in income because her baseline income was comparatively high relative to her husband's.

This paper's findings have significant implications for theory and policy. We document that the form of female empowerment most forcefully advocated in the effort to reduce abuse of women – employment – increases rather than decreases domestic violence in the context of rural Ethiopia, and that the reason appears to be that men act upon the emotional costs implied by deviations from traditional household roles.

We do not attempt to survey the literature on domestic violence here, but briefly summarize some of the most relevant papers. There are (at least) two cross-cutting dichotomies of domestic violence models: optimistic versus pessimistic models, and instrumental models in which violence is used to gain control over household resources versus models in which violence is not used to gain control over resources. Examples of conventional optimistic models include Chwe (1990) and Aizer (2010). In Chwe (1990), a male principal can use financial disincentives to discourage low effort (for example in home production) from a high income female agent but must instead use costly violence disincentives to motivate a low income female. In Aizer (2010), improvements in a woman's expected utility outside of marriage, for example due to employment, is expected to reduce the level of violence she is willing to "offer" a husband who derives utility from violence. Aizer finds that decreases in the male-female wage gap in the U.S. reduce violence against women.

There are several potential reasons why Aizer's findings differ from ours. One possibility is that, in more male-dominated cultures such as that of many developing countries, the marginal utility men derive from violence may increase as women's standing improves. Though not all the findings of previous studies can necessarily be interpreted causally, our results add to increasing evidence that nominal empowerment of women in poor countries can increase domestic violence. Eswaran and Malhotra (2009) find that women in India who work outside of the home are subjected to more violence. Gonzalez-Bernes (2004) concludes that female labor force participation in Zambia, Rwanda and Tanzania is not associated with lower levels of violence. The evidence for middle income countries is mixed at best.¹

Instrumental models typically argue that men use violence as a tool to gain control over household resources, rather than as an end in itself. Examples of pessimistic instrumental violence models include Bloch and Rao (2002) and Bobonis et al. (2010). Alternatively, men may derive "expressive" utility directly from violence, in which case physical abuse can

¹Angelucci (2008) finds that, while small conditional cash transfers to women in Mexico can decrease violence, larger transfers increase the aggressive behavior of husbands. Bobonis and Castro (2010) find no effect of cash transfers on domestic violence in the long run.

be triggered by events that have purely symbolic meaning (see Card and Dahl, 2010). This paper’s findings are most supportive of the expressive “male backlash” theories emphasized by sociologists (e.g., Macmillan and Gartner, 1999). We argue that backlash models reinterpreted in an economic framework do not necessarily “ignore the individual rationality constraints faced by women” (Aizer, 2010: 4), but rather take seriously an additional motive on the part of men – that of restoring a self-image of dominance in the household to which they may feel entitled, for example due to cultural norms. A similar theory, in an instrumental framework, would be that men use violence to attempt to address unwanted female behavior associated with employment.

The paper is organized as followed. Section 2 describes the rural Ethiopian context and the experiment. In section 3 the main treatment effects are presented and analyzed in light of existing domestic violence models. Section 4 concludes.

2. The Setting

Ethiopia has some of the highest poverty, illiteracy and underemployment rates in Africa, especially for women.² Domestic violence is unusually prevalent; for example, 54 percent of women in a provincial site surveyed by the WHO report to have been victimized by a partner during the last year (Garcia-Moreno et al., 2006). At least until recently, a role for domestic violence was accepted in Ethiopian culture – even by many women. In a nationally representative survey conducted in 2005, 81 percent of Ethiopian women found it justified for a husband to beat his wife if the wife had violated norms (ECSA and ORC Macro 2006, p. 244).

In recent years it has become more common for Ethiopian women to hold formal jobs. In rural areas an important contributing factor has been the explosive rise of the floriculture sector, which mostly employs women. In 2008, 81 flower farms in Ethiopia employed around 50,000 workers (Gebreyesus and Iizuka, 2010). Hiring on Ethiopian flower farms typically takes place in October and November, before the main growing and harvesting season.

The supervisors on five flower farms agreed to randomize job offers during the fall 2008 hiring season because of an unusual situation in the labor market for flower farm workers.³ At the time, applicants almost always outnumbered the positions to be filled by large margins. Ethiopian flower farms - still getting to grips with cost components significantly larger than labor, and with little ability to predict the productivity of the mostly uneducated, illiterate and inexperienced applicants - did not prioritize optimization of the unskilled workforce (Gebreyesus and Iizuka, 2010).⁴ Because supervisors were already allocating job offers relatively arbitrarily when approached by the researchers, explicit randomization was a modest procedural change.

²This section follows Hjort (2011). See that paper for more details on the field experiment.

³Two additional farms were originally part of the sample but were dropped before the baseline survey was completed and the randomization took place.

⁴Labor costs make up a relatively small portion of flower farms’ total costs (Gebreyesus and Iizuka, 2010). The reason why wages for flower farm workers did not fall is an important unanswered question for future research.

The five farms are located in rural areas two and a half to five hours from Addis Ababa and employ local workers who live in small towns nearby the farms. On hiring days, supervisors first excluded any unacceptable applicants. A team of enumerators then carried out the baseline survey with the remaining applicants. Finally, the names of the number of workers to be hired were drawn randomly from a hat.⁵ The sample thus consists of 339 households in which a woman applied to a flower farm job and was deemed acceptable for hiring; we focus on the 329 households in which the applicant was married or living with a steady partner. We attempted to re-interview everyone in the treatment and control groups 5 - 7 months after employment commenced.⁶ Careful tracking procedures led to a re-interview rate of 88 percent and no statistically significant differential attrition. Summary statistics are displayed in table 1. There are no statistically significant differences between the characteristics of the treatment and control groups. Literacy rates are low. Almost all the applicants are parents. Income and wealth indicators, such as the material that the applicant's floor is made of, indicate the severe poverty of the sample.

Flower farm employment typically entails six days of full-time work a week, totaling on average 202 hours per month. The alternative for the women in our sample was typically domestic work, and perhaps a few hours of informal paid work per week. The applicants randomly chosen for employment spent 102 more hours per month working (i.e., doing paid or house-work). The income of treated women increased by 154 percent on average, which translates into a 28 percent increase in total household income.

3. The Impact of Female Employment on Domestic Violence

3.1 Results

In this section we present estimates of the impact of female employment on domestic violence in Ethiopia, and interpret the findings in light of existing theories. Table 2 displays the baseline incidence of different forms of emotional abuse, such as insults or threats, as well as physical violence, such as pushes, slaps, punches or sexual assault. On average, women experienced 0.96 violent incidents per month before seeking employment, or 1.57 incidents among the 63 percent of women who were physically abused. These figures are comparable to those found in other parts of Ethiopia (Garcia-Moreno et al., 2006).

Due to the experimental design, we can identify the impact of female employment on violence through simple OLS regressions of the form

$$y_{ijt} = \alpha + \beta E_{it} + \gamma X_{it} + \eta_j + \varepsilon_{ijt}$$

where y_{ijt} is a violence outcome, E_{it} a dummy for treatment, η_j farm fixed effects and X_{it} other controls (which have little influence on the estimated effects).

⁵The applicants and most members of the supervisor and enumerator teams were not aware of the explicit randomization, but were likely aware of the generally arbitrary character of hiring procedures.

⁶Only 4 re-interviewed women in the control group had managed to obtain employment, probably because few farms hire workers in the season that followed the time of our randomization and there were few other employment options in the areas where the sample farms are located.

The estimated treatment effect are in table 3. The probability of experiencing physical violence increases by 8 percentage points or 13 percent when a woman gets employed in rural Ethiopia. There is also a 19 percentage point or 34 percent increase in emotional abuse. Finally, the intensive margin of violence is affected: the number of violent incidents experienced per month goes up by 0.31 or 32 percent following employment.

An alternative interpretation of these results is that employment affects women's willingness to report violence to an enumerator rather than, or in addition to, violence itself. While we cannot rule out a reporting effect, greater willingness to report violence after employment is unlikely to represent the primary explanation of our findings. Specific, detailed survey questions were used. As noted above, the majority of both men and women in Ethiopia find domestic violence justifiable in some situations, and 63 percent of women in our sample were comfortable reporting abuse at baseline.⁷

The prediction that physical abuse will decrease when women are "empowered" by employment is central to the most-cited domestic violence models. The estimates in table 3 represent strong evidence against such models, in the context of rural Ethiopia. In the next two sub-sections we categorize pessimistic models on the basis of the hypothesized male motivation for abuse, and explore the ability of different categories of pessimistic models to explain our findings.

3.2 Domestic Violence, Financial Control and the Effect of Female Employment

Domestic violence is often modeled as a tool used by husbands to extract rents from wives, typically control over household resources. In such "instrumental" models, a husband may turn more violent when his wife gets employed if employment leads to an increase in the wife's bargaining power that the husband wants to counteract, and/or because more money is "on the table" when the wife is employed.

Instrumental violence should occur only if contracts to avoid it cannot be written: in the benchmark bargaining scenario a Pareto-improving agreement to avoid costly abuse should be reached.⁸ If violence is used to achieve financial control and a market failure leads to the existence of violence in equilibrium, the extent of male control over resources and the extent of domestic violence should arguably be correlated across households. Taking advantage of a survey question that asked women whether or not they had money that they alone could decide how to spend, we present the following evidence:

1. At baseline, 37 percent of women in the sample said that they had money that they alone could decide how to spend. There was no statistically significant difference in

⁷Note also that women interviewed by male enumerators at baseline were as likely to report physical abuse as those interviewed by female enumerators. This fact highlights the openness surrounding domestic violence in our sample.

⁸Asymmetric information about a spouse's "gains from marriage" has been proposed as one possibility (Bloch and Rao, 2002; Bobonis et al., 2010); another possibility is commitment problems.

initial financial control across women experiencing domestic violence and women not being abused.⁹

2. At follow-up, 47 percent of women randomly chosen for employment said that they had control over some money, while 38 percent of women in the comparison group said that they controlled some money. There was no statistically significant difference in financial control across women experiencing domestic violence and women not being abused at follow-up.
3. Husbands do not appear to “counteract” an increase in wives’ bargaining power due to employment by turning more violent. We classify a woman as economically independent if she said that “If I wanted to divorce my spouse I would be able to support my family on my own”.¹⁰ Because, by definition, only economically dependent women can become independent through employment, the increase in violence should be lower for already independent women if violence is used to counteract an increase in a woman’s bargaining power due to employment enhancing her economic independence. As seen in columns one and three of table 4, the interaction between a dummy for baseline economic independence and the treatment is not significant.¹¹

Because instrumental models typically allow men to use violence either to gain “extra” control over resources or to compensate for an initial perceived “deficit” of control, such models are rarely explicit about implied, empirically testable correlations between control and violence. But it is difficult to reconcile the combination of evidence in bullets 1, 2 and 3 with a model in which violence is used by men to gain control over household resources.

3.3 Domestic Violence, Gender Norms and the Effect of Female Employment

This paper’s primary result is that domestic violence increases significantly when women get employed in rural Ethiopia. Knowledge of the mechanism that underlies the increase in abuse is important for theory and policy. It appears that there are two categories of models that may be able to explain our results: expressive models in which a husband’s marginal utility from violence is increasing in the economic standing of his wife, and instrumental models in which violence is used to achieve male goals other than control over household resources. We consider these two possibilities in turn.

Aizer (2010) is an example of an influential class of expressive domestic violence models in which men derive utility directly from violence. Women with better options outside of marriage should be willing to accept less violence at a given “price”: employment is predicted

⁹We take a woman reporting that she has money that she alone can decide how to spend to mean that she has true control over the money.

¹⁰The alternative was “I need to be married in order to be able to provide for a family”.

¹¹The reader might prefer to see employment as an improvement in the outside option of *all* women, rather than only those who report increased independence in a survey. In that case, however, male control over household resources should arguably decrease in all households in which husbands did not turn more violent in response to female employment, if violence is used to maintain male control. This is not what we observe in the data.

to shift a woman's violence "supply curve" up and thus decrease violence. Consider, however, that a husband's violence "demand curve" may also shift up when his wife gets employed, if the husband's marginal utility from violence is increasing in the wife's relative or absolute economic standing. The net outcome may be that the couple's contract curve – the set of feasible bargaining solutions – shifts up in (violence, consumption) space, and that violence itself therefore increases.

Why would the marginal utility that men in Ethiopia derive from violence go up when women get employed? Suppose that there are emotional costs to men of perceived violations of traditional gender roles. In that case "violence may be a means of reinstating [a husband's] authority over his wife" (Macmillan and Gartner, 1999: 949). If improvements in women's economic standing carry emotional costs to men, events that symbolize the perceived challenge to traditional gender roles can likely lead to violence. In columns two and four of table 4 we interact the treatment indicator with the wife's ex ante income as a share of the combined income of the husband and wife. The results show that the impact of employment on violence is bigger in households in which the newly employed wife is likely to end up further ahead of her husband in income because her share of baseline income was high relative to that of other women in the sample.¹² The increase in the probability of violence when a wife gets employed is seven percentage points higher for every one standard deviation increase in the wife's share of baseline income, almost as much as the average effect.¹³ There is also a small but marginally significant *increase* in male labor supply when women get employed in rural Ethiopia. Though alternative explanations are possible, these results are consistent with a plausible story in which improvements in the relative economic standing of women carry emotional costs to men; costs that some men choose to act upon through violence.

A similar possibility is that violence serves an instrumental purpose, but is used not to gain control over household resources but instead to influence the behavior of wives (see Tauchen et al. 1991; Eswaran and Maholtra 2009). Husbands may see some dimensions of female behavior associated with employment as undesirable and potentially "correctable" through violence. The arguably most plausible "real" cost to husbands of female employment is that employed wives devote less time to house-work. In our sample, most of the house-work of women randomly chosen for employment is taken over by daughters (see Hjort, 2011), however. This suggests that costs to husbands of a reallocation of women's time may, if anything, be due the overturning of traditional responsibilities in the household, rather than house-work being left undone.

In sum, the evidence presented here suggests that emotional costs associated with violations of traditional gender roles belong in theories of domestic violence in gender-unequal societies. If so, identity models, in which disutility is associated with a self-image that de-

¹²An ideal analysis would have compared the response of violence in households close to but on opposite sides of the threshold at which the wife was expected to surpass her husband in income due to her employment. But there are very few sample households around that threshold: almost all women in our sample surpass their husband in income when employed.

¹³These results are not driven by the husband's absolute level of income. Note that the effect on the number of violent incidents per month is not significant.

viates from the individual's view of his or her "appropriate" role in the household, are a natural starting point (Akerlof and Kranton, 2010). In the appendix we present an example of a framework in which a husband's incentive to engage in violence depends on his wife's economic standing relative to his own - as does, in turn, the wife's response to violence. The framework allows a male "backlash" when women get employed and predicts how domestic violence responds to female employment in Ethiopia well.

4. Conclusion

This paper has analyzed the impact of female employment on domestic violence through a field experiment in which women's long-term job offers on Ethiopian flower farms were randomized. We estimate a significant 13 percent increase in physical violence when women get employed, as well as large increases in emotional abuse and the intensity of physical violence. These results put into question the relevance of conventional economic models of domestic violence in male-dominated developing countries. Like much existing anti-violence policy, conventional models are "optimistic" in the sense of considering labor force participation a promising route to empowering women and reducing the prevalence of domestic violence.

Most "pessimistic" models argue that physical abuse can increase when employment enhances wives' incomes and bargaining power because husbands use violence as a tool to get access to and control over household resources. But we find no significant correlation between levels of violence and control over household resources, nor changes in violence and control when women get employed, and the reason does not appear to be that violence is used to counteract female bargaining power.

Rather than a male quest for control over household resources, it appears that the models that best explain our results would allow men to care about roles in the household deviating from the roles prescribed by traditional norms, and violence being seen as a way to restore a preferred order. We find that the increase in the probability of violence following female employment is greater in households in which the newly employed woman is likely to end up further ahead of her husband in income. The costs to a husband of lost economic dominance are presumably primarily emotional, suggesting that the benefits of turning to violence in response may also be emotional. It may be that men derive "expressive" utility from violence and, while a woman's "violence supply curve" likely shifts up when her outside option improves, her husband's "violence demand curve" also shifts up because his marginal utility from violence depends on his wife's relative (or absolute) economic standing. A similar "instrumental violence" interpretation would be that men abuse their wives not to achieve financial control but rather, for example, to influence their wives' behavior in the household.

We conclude that: (1) conventional optimistic economic models of domestic violence are unlikely to accurately describe the situation in most households in male-dominated developing countries such as Ethiopia; and (2) not all men will passively accept challenges to their economic dominance, and successful models of domestic violence will likely need to account for the male reaction to female economic progress.

Finally, it is worth emphasizing that the increase in domestic violence we observe when women get employed does not mean that women are not empowered by employment. For

example, it may be that some women previously acquiesced in the face of demands from their husbands but choose not to when emboldened by employment. It is also possible that gender norms themselves slowly respond to female employment, in which case the longer term impact on violence could differ from the deleterious effect observed here. Rather than suggesting that female employment should not be encouraged, the evidence presented here indicates that economic theory, domestic violence policy and female employment programs should take the costs to men of violations of traditional gender roles seriously – insofar as such violations prove costly for women.

Tables

Table 1: Summary Statistics and Randomization Balance

	Treatment (174)	Control (165)	Difference
Parent	0.96	0.96	0.00 (0.02)
Sons	1.49	1.75	-0.25 (0.17)
Daughters	2.12	1.99	0.14 (0.16)
Has cement/wood floor	0.21	0.26	-0.05 (0.05)
Household income	996.46	1073.81	-77.35 (62.49)
Literate	0.19	0.14	0.04 (0.04)
Spouse literate	0.29	0.30	-0.01 (0.05)
Age	26.74	26.26	0.48 (0.72)
Spouse age	31.04	29.97	1.06 (0.72)
Applicant paid work hours per month	39.43	37.08	2.35 (5.54)
Spouse paid work hours per month	64.46	64.13	0.33 (3.84)
Experienced emotional abuse	0.54	0.58	-0.04 (0.05)
Experienced physical violence	0.64	0.62	0.01 (0.05)
Violent incidents per month	0.95	0.96	-0.01 (0.11)

Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2: Baseline prevalence of domestic violence

Emotional abuse	
Does your husband ever...	
threaten to hurt you or someone close to you?	0.44 (0.50)
insult you or make you feel bad about yourself?	0.48 (0.50)
Physical violence	
Does your husband ever...	
push you, shake you or throw something at you?	0.42 (0.49)
slap you?	0.31 (0.46)
punch you with his fist or with something that could hurt you?	0.15 (0.36)
physically force you to have sexual intercourse with him even when you do not want to?	0.12 (0.32)
Violent incidents per month	
How many times per month does your spouse usually act violently towards you?	0.96 (1.00)
Standard deviations in parentheses	

Table 3: Impact of female employment on domestic violence

	Experienced emotional abuse	Experienced physical violence	Violent incidents per month
	(1)	(2)	(3)
Treatment	0.19*** (0.03)	0.08*** (0.02)	0.31*** (0.07)
N	329	329	329
Dep. var. mean	0.56	0.63	0.96
Covariates?	YES	YES	YES

Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The covariates are applicant's age and literacy, and an indicator for the household having a cement / wood floor

Table 4: Exploring existing models' ability to explain the impact of female employment on domestic violence

	Experienced physical violence		Violent incidents per month	
	(1)	(2)	(3)	(4)
Treatment	0.09*** (0.03)	0.01 (0.03)	0.35*** (0.07)	0.25*** (0.09)
Treatment \times Wife economically independent at baseline	-0.05 (0.06)		-0.19 (0.19)	
Treatment \times Wife's baseline income as share of household income		0.21*** (0.08)		0.19 (0.19)
N	329	329	329	329
Dep. var. mean	0.63	0.63	0.96	0.96
Covariates?	YES	YES	YES	YES

Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The covariates are applicant's age and literacy, and an indicator for the household having a cement / wood floor

Theoretical Appendix

If gender roles belong in theories of domestic violence in gender-unequal societies, then identity models are a natural starting point. To illustrate, we present a simple model that captures the idea that the marginal utility that a husband derives from violence may depend on the relative economic standing of his wife. There is evidence that the roles of moderate and extreme violence in the family differ; we distinguish between the two here (see e.g., Johnson, 2009). To capture the idea that violence can take place even when women are “empowered”, we allow women to resist or hit back when they are assaulted.

Let a household consist of a husband (h) and a wife (w), and let $s = h, w$ denote the spouse in question, and $-s$ the other spouse. The husband’s income, y_h , is exogenously given. The wife’s income, y_w , is low if she does informal or no work, y_w^0 , and high if she is employed, y_w^e , where $y_w^e > y_w^0$. Let $Y = y_w + y_h$. Each spouse derives utility from a self-image, i . Let $i = a$ if the spouse sees his/her position in the household as fully in accordance with some reference point, represented for example by the role the individual sees traditional norms as prescribing for someone of his or her gender. So $\partial U/\partial i = 0$ for $i \geq a$ and $\partial U/\partial i > 0$ for $i < a$.

Suppose that a spouse’s self-image is determined by (1) the individual’s perception of his/her “inherent” role in the relationship, modeled as an individual specific constant, $j > 0$; (2) the individual’s share of total household income, y_s/Y , where a higher share yields an improved self-image; and (3) violence. An individual can act violently against his/her spouse, which improves the abuser’s self-image while worsening the victim’s. Let $v \in [0, \dots, g, \dots, 1]$ and $r \in [0, \dots, g, \dots, 1]$ indicate the intensity of male and female violence, respectively, where $v \wedge r < g$ represents moderate violence and $v \wedge r \geq g$ extreme violence. Violence carries costs for both spouses, $c_s(v, r)$, and the costs are increasing in intensity.¹⁴

If spouse s ’s self-image is $i_s = i_s(j_s, \frac{y_{-s}}{Y}, v, r)$ we can write

$$U_s = u_s(Y, i_s(j_s, \frac{y_{-s}}{Y}, v, r), c_s)$$

where U_s is increasing and quasi-concave in Y , increasing and quasi-concave in i_s for $i_s < a$, and decreasing and strictly concave in c_s . We assume that the wife never initiates violence, and that there is a probability q that she retaliates if her husband abuses her. While we treat q as fixed for simplicity, in reality it is reasonable to assume that q is influenced by the woman’s attitude towards domestic violence. Since husbands are likely to know their wives’ attitudes well, we assume that q is known to the husband. If the wife retaliates, she does so with the intensity of the violence she was subjected to - $r \equiv v \vee 0$ - which brings both spouses’ self-images back to their initial levels. Because of the possible humiliation of having his dominance challenged, if the wife retaliates we assume that there is a probability p that the husband “loses control” and responds with extreme physical assaults - even if his costs of doing so are higher than his expected benefits. If the wife retaliates and the husband does not lose control, he abstains from further violence. Denote by v^{ic} the level of violence

¹⁴The costs to the abuser could for example arise from increased tension between the spouses or loss of intimacy.

chosen by the husband when “in control”, and v^{lc} the level if he lost control. We assume that $v^{lc} > v^{ic*}$ and $v^{lc} > g$.¹⁵ Due to the physical disparity and risks involved, we assume that the wife does not retaliate if the husband loses control. The timeline of the game can be summarized as follows:

Stage 1. Husband chooses v^{ic} . If $v^{ic*} = 0$, the game ends. If $v^{ic*} > 0 \rightarrow$ stage 2

Stage 2. Wife chooses r . If $r^* = 0$, the game ends. If $r^* > 0 \rightarrow$ stage 3

Stage 3. Husband “chooses” $v^{lc} \geq g$ with probability p , and $v^{lc} = 0$ with probability $(1 - p)$

We begin with a situation where the wife is not employed. Let the reservation utility of a spouse, the minimum utility level required to remain in the marriage, be denoted \bar{U}_s^0 . The husband’s choice of v^{ic} is then given by maximization of U_h s.t. $U_h \geq \bar{U}_h^0$ and $U_w \geq \bar{U}_w^0$. Let v^{ic*} denote an interior solution. The wife decides whether to retaliate if the husband abuses her: if $v^{ic*} > 0$ her choice of r is given by maximization of U_w s.t. $U_h \geq \bar{U}_h^0$ and $U_w \geq \bar{U}_w^0$. Let r^* denote an interior solution. The four possible outcomes are:

1. (U_w^0, U_h^0) . No violence: the husband chooses $v^{ic*} = 0$ at stage 1
2. (U_w^{ic}, U_h^{ic}) . Violence at stage 1: $v^{ic*} > 0$. No retaliation at stage 2
3. (U_w^r, U_h^r) . Violence at stage 1: $v^{ic*} > 0$. Retaliation at stage 2: $r^* > 0$. No violence at stage 3
4. (U_w^{lc}, U_h^{lc}) . Violence at stage 1: $v^{ic*} > 0$. Retaliation at stage 2: $r^* > 0$. Violence at stage 3: $v^{lc*} > 0$

The wife always prefers no violence. If the husband acts violently at stage 1, the wife is better off retaliating compared to not retaliating only if there is no violence at stage 3, so $U_w^0 > U_w^r > U_w^{ic} > U_w^{lc}$. If a husband with $i_h < a$ acts violently at stage 1 and the wife does not retaliate, then this leaves the husband better off as compared to the no-violence alternative. But if the wife retaliates, then the husband would have been better off not abusing her in the first place, so that $U_h^{ic} > U_h^o > U_h^r$. If the husband loses control and engages in extreme violence at stage 3 he gets U_h^{lc} where $U_h^{lc} < U_h^{ic}$.

We solve the game by backward induction. At stage 3, the wife has retaliated; the husband then loses control and engages in extreme violence with probability p . At stage 2, the wife’s expected payoff of retaliating will thus be $pU_w^{lc} + (1 - p)U_w^r$, so she retaliates if $pU_w^{lc} + (1 - p)U_w^r > U_w^{ic}$. At stage 1, the husband can secure U_h^0 by refraining from violence. If $i_h < a$ and the husband abuses the wife at stage 1, he gets his highest possible outcome, U_h^{ic} , if there is no retaliation. But the wife retaliates with probability q , in which case his expected utility is $pU_h^{lc} + (1 - p)U_h^r$. So the husband will be violent at stage 1 if

¹⁵We allow for the possibility that the husband rationally chooses extreme violence, i.e. $v^{ic*} > g$.

$q(pU_h^{lc} + (1 - p)U_h^r) + (1 - q)U_h^{ic} > U_h^o$. In households in which the wife is close to her reservation utility, or where the husband has a strong self-image, or where the spouses are highly risk averse, there will be no violence.

Although shown here primarily to illustrate that self-image can be incorporated into domestic violence models in a reasonable fashion, this framework in fact predicts the results in tables 3 and 4. Several additional predictions also find empirical support.

(A) When women get employed and their relative incomes increase, violence becomes more prevalent. For reasonable assumptions about the impact of violence on utility across the intensity range, the proportional increase in moderate violence will be greater than the proportional increase in extreme violence (since U_h is increasing and quasi-concave in i_h for $i_h < a$, the optimal level of violence is increasing in y_w/Y . Female employment essentially makes violence beneficial for some men, but those newly violent men will generally exert comparatively moderate violence).

(B) The increase in the incidence of both moderate and extreme violence when women get employed will be greater the lower the ex ante relative income of husbands (the greater the initial relative income and self-image of a husband, the more likely that he will continue to have $i_h \geq a$ when his wife's employment leads to a worsening of his self-image, in which case his behavior will be unaffected).

(C) When women get employed and their relative incomes increase, the incidence of female resistance to violence will increase. The increase in the incidence of women resisting violence when they get employed will be greater among women with low ex ante relative incomes. If q , the probability of a wife resisting, varies, for example with employment itself or with the extent to which the wife disapproves of the use of violence, then the increase in retaliation will also vary with those factors (some of the women who have a husband that turns violent after employment will prefer to resist. Wives who already had a high relative income and therefore a sound self-image before employment further improves their standing may continue to have $i_w \geq a$, even if abuse from their husbands increases. And if, for example, women who generally consider abuse to be unacceptable have higher qs , then the increase in resistance following employment will be greater among such women).

This framework was inspired by our primary finding – the increase in violence when women get employed – and by anecdotal evidence and stories of male “backlashes” from Ethiopian women. As mentioned, the framework’s auxiliary predictions are also empirically supported. Resistance to violence increases significantly following female employment. Both the increase in the incidence of male violence and in that of female resistance is significantly greater the lower the baseline relative income of the relevant spouse. The increase in the incidence of resistance is significantly greater among women that at baseline believed domestic violence to be unjustified.

Bibliography

- AIZER, A. (2010): “The gender wage gap and domestic violence,” *American Economic Review*, forthcoming.
- AKERLOF, G. A., AND R. E. KRANTON (2010): *Identity Economics*. Princeton University Press.
- ALESINA, A., AND E. L. FERRARA (2005): “Ethnic diversity and economic performance,” *Journal of Economic Literature*, 43(3), 762–800.
- ALESINA, A., AND E. SPOLAORE (1997): “On the number and size of nations,” *Quarterly Journal of Economics*, 112(4), 1027–1056.
- ARCIDIACONO, P., G. FOSTER, N. GOODPASTER, AND J. KINSLER (2011): “Estimating spillovers using panel data, with an application to the classroom,” *Quantitative Economics*.
- BANDIERA, O., I. BARANKAY, AND I. RASUL (2005): “Social preferences and the response to incentives: Evidence from personnel data,” *Quarterly Journal of Economics*, 120(3), 917–962.
- BANERJEE, A. V., AND B. MOLL (2009): “Why does misallocation persist?,” Mimeo, MIT.
- BAR, T., AND K. BASU (2009): “Children, education, labor, and land: In the long run and short run,” *Journal of the European Economic Association*, 7(2–3), 487–497.
- BARKAN, J. D. (2004): “Kenya after moi,” *Foreign Affairs*.
- BASU, K., S. DAS, AND B. DUTTA (2010): “Child labor and household wealth: Theory and empirical evidence of an inverted-U,” *Journal of Development Economics*, 91(1), 8–14.
- BAUER, M., A. CASSAR, AND J. CHYTILOVÁ (2011): “Warfare and social preferences in children,” *Mimeo, USF*.
- BECKER, G. (1957): *The Economics of Discrimination*. Univ. Chicago Pres.
- (1974): “A theory of social interactions,” *Journal of Political Economy*.

BIBLIOGRAPHY

- BHATTACHARYA, D. (2009): “Inferring optimal peer assignment from experimental data,” *Journal of the American Statistical Association*.
- BLOCH, F., AND V. RAO (2002): “Terror as a bargaining instrument: A case study of dowry violence in rural India,” *American Economic Review*, 92(4), 1029–43.
- BOBONIS, G. J., R. CASTRO, AND M. GONZÁLEZ-BRENES (2009): *Public Transfers and Domestic Violence: The Roles of Private Information and Spousal Control*. University of Toronto.
- BOISJOLY, J., G. DUNCAN, M. KREMER, D. LEVY, AND J. ECCLES (2006): “Empathy or antipathy? The consequences of racially and socially diverse peers on attitudes and behaviors,” *American Economic Review*.
- BONKE, J., AND M. BROWNING (2009): “The distribution of financial well-being and income within the household,” *Review of Economics of the House*, 7(1).
- BOOZER, M., AND T. SURI (2001): “Child labor and schooling decisions in Ghana,” *Manuscript, Yale University*.
- BROWNING, M., AND P. A. CHIAPPORI (1998): “Efficient intra-household allocations: A general characterization and empirical tests,” *Econometrica*, 66(6), 1241–1278.
- BUCKLEY, S. (1997): “Youth in Kenya now feel few tribal ties,” *Washington Post*, <http://www.washingtonpost.com/wpsrv/inatl/longterm/africanlives/kenya/kenya.htm>.
- BURGESS, R., R. JEDWAB, E. MIGUEL, A. MORJARIA, AND G. P. I MIQUEL (2011): “Ethnic favoritism,” *unpublished working paper*.
- CARD, D., AND G. DAHL (2010): “Family violence and football: The effect of unexpected emotional cues on violent behavior,” *Quarterly Journal of Economics*, forthcoming.
- CARRELL, S., B. SACERDOTE, AND J. WEST (2011): “From natural variation to optimal policy? The Lucas critique meets peer effects,” *Mimeo, UC Davis*.
- CHARNESS, G., AND M. RABIN (2002): “Understanding social preferences with simple tests,” *Quarterly Journal of Economics*.
- CHETTY, R. (2009): “Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods,” *Annual Review of Economics*.
- CHIAPPORI, P. A. (1988): “Nash-bargained households decisions: A comment,” *International Economic Review*, 29(4), 791–796.
- (1992): “Collective labor supply and welfare,” *Journal of Political Economy*, 100(3), 437–467.

BIBLIOGRAPHY

- CHWE, M. (1990): “Why were workers whipped? Pain in a principal-agent model,” *Economic Journal*, 100, 1109–1121.
- CLARK, S. (2000): “Son preference and sex composition of children: Evidence from India,” *Demography*, 37(1), 95–108.
- CONNELLY, R., D. S. DEGRAFF, AND D. LEVISON (1996): “Women’s employment and child care in Brazil,” *Economic Development and Cultural Change*, 44(3), 619–656.
- DAUPHIN, A., A.-R. EL LAHGA, B. FORTIN, AND G. LACROIX (2008): “Are children decision-makers within the household?,” working paper 3728, Institute for the Study of Labor (IZA).
- DORAN, K. (2008): “Are adults in demand when children leave the land? Evidence from rural Mexico,” .
- DUPAS, P., AND D. BHATTACHARYA (2011): “Inferring welfare maximizing treatment assignment under budget constraints,” *Mimeo, Stanford*.
- DUPAS, P., AND J. ROBINSON (2011): “The (hidden) costs of political instability: Evidence from Kenya’s 2007 election crisis,” *Mimeo, Stanford*.
- EASTERLY, W., AND R. LEVINE (1997): “Africa’s growth tragedy: Policies and ethnic divisions,” *Quarterly Journal of Economics*.
- ECONOMIST INTELLIGENCE UNIT (2008): “Economist intelligence unit Ethiopia report,” .
- ECSA, AND ORC MACRO (2006): *Ethiopia Demographic and Health Survey*. Ethiopia Central Statistical Agency and ORC Macro, Addis Ababa, Ethiopia and Calverton, USA.
- EDMONDS, E. V. (2008): “Child labor,” vol. 4 of *Handbook of Development Economics*, chap. 57, pp. 3607 – 3709. Elsevier.
- EDMONDS, E. V., AND N. PAVCNIK (2005): “The effect of trade liberalization on child labor,” *Journal of International Economics*, 65(2), 401–419.
- EDRI (2008): “Ethiopian development research institute: Micro evidence on the development of cut-flower industry in Ethiopia,” Discussion paper.
- EIFERT, B., E. MIGUEL, AND D. N. POSNER (2010): “Political competition and ethnic identification in Africa,” *American Journal of Political Science*, 54(2), 494–510.
- ESWARAN, M., AND N. MALHOTRA (2009): *Domestic Violence and Women’s Autonomy: Evidence from India*. University of British Columbia.
- FRANCAVILLA, F., G. C. GIANNELLI, AND L. GRILLI (2010): “Mothers’ employment and their children’s schooling: A joint multilevel analysis for India,” *Working Papers Series*.

BIBLIOGRAPHY

- GARCIA-MORENO, C., H. A. F. M. JANSSEN, M. ELLSBERG, L. HEISE, AND C. H. WATTS (2006): "Prevalence of intimate partner violence: Findings from the WHO multi-country study on women's health and domestic violence," *The Lancet*, 368(9543), 1260–1269.
- GEBREEYESUS, M., AND M. IIZUKA (2010): "Discovery of the flower industry in Ethiopia: Experimentation and coordination," working paper 2010–025, UNU-MERIT.
- GIBSON, C., AND J. LONG (2009): "The presidential and parliamentary elections in Kenya," *Electoral Studies*.
- GONZALEZ-BRENES, M. (2004): "Domestic violence and household decisionmaking: Evidence from East Africa," UC Berkeley, http://www.sscnet.ucla.edu/polisci/wgape/papers/7_Gonzalez.pdf.
- GRAHAM, B., G. IMBENS, AND G. RIDDER (2011): "Measuring the effects of segregation in the presence of social spillovers: A nonparametric approach," *Mimeo, UC Berkeley*.
- HABYARIMANA, J., D. N. POSNER, AND J. WEINSTEIN (2007): "Why does ethnic diversity undermine public goods provision?," *American Political Science Review*.
- HAMILTON, B. H., J. A. NICKERSON, AND H. OWAN (2005): "Diversity and productivity in production teams," *working paper, Washington University*.
- HJORT, J. (2011): "Parent's employment and child outcomes: Experimental evidence from Ethiopia," *Mimeo, UC Berkeley*.
- HSIEH, C.-T., AND P. J. KLENOW (2009): "Misallocation and manufacturing TFP in China and India," *Quarterly Journal of Economics*, 124(4), 1403–1448.
- INTERNATIONAL LABOR ORGANIZATION (2005): "Characteristics and determinants of youth unemployment, underemployment and inadequate employment in Ethiopia," .
- JAKIELA, P., E. MIGUEL, AND V. TE VELDE (2011): "Combining field and lab experiments to estimate the impact of human capital on social preferences," *Mimeo, UC Berkeley*.
- JOHNSON, M. P. (2009): "Differentiating among types of domestic violence: Implications for healthy marriages," in *Marriage and Family: Perspectives and Complexities*, ed. by H. E. Peters, and C. M. K. Dush. Columbia University Press, New York.
- JONES, C. (2011): "Intermediate goods and weak links in the theory of economic development," *American Economic Journal: Macroeconomics*.
- KANDEL, E., AND E. LAZEAR (1992): "Peer pressure and partnerships," *Journal of Political Economy*, 100, 801–817.

BIBLIOGRAPHY

- KENYA FLOWER COUNCIL (2011): “<http://www.kenyaflowercouncil.org/floricultureinkenya.php>,”
- KRUEGER, A., AND A. MAS (2004): “Strikes, scabs and tread separations: Labor strife and the production of defective Bridgestone/Firestone tires,” *Journal of Political Economy*.
- KSOLL, C., R. MACCHIAVELLO, AND A. MORJARIA (2010): “The effect of ethnic violence on an export-oriented industry,” *Mimeo, Oxford*.
- LA FERRARA, E. (2002): “Self-help groups and income generation in the informal settlements of Nairobi,” *Journal of African Economies*.
- (2003): “Kin groups and reciprocity: A model of credit transactions in Ghana,” *American Economic Review*.
- LANG, K. (1986): “A language theory of discrimination,” *Quarterly Journal of Economics*, 101, 363–382.
- LAZEAR, E. P. (1998): “Globalization and the market for teammates,” working paper 6579, National Bureau of Economic Research.
- MACMILLAN, R., AND R. GARTNER (1999): “When she brings home the bacon: Labor force participation and the risk of spousal violence against women,” *Journal of Marriage and the Family*, 61(4), 927–958.
- MAS, A., AND E. MORETTI (2009): “Peers at work,” *American Economic Review*.
- MELESE, A. T., AND A. H. J. HELMSING (2010): “Endogenisation or enclave formation? The development of the Ethiopian cut flower industry,” *Journal of Modern African Studies*.
- MIGUEL, E. (2004): “Tribe or nation? Nation-building and public goods in Kenya versus Tanzania,” *World Politics*.
- NDEGWA, S. (1997): “Citizenship and ethnicity: An examination of two transition moments in Kenyan politics,” *American Political Science Review*, 91(3), 599–616.
- NOURY, V. (2011): “Kenya’s cut flower tax swindle,” *African Business*, <http://62.128.151.219/A1rwqk/AB0511/resources/64.htm>.
- OYUGI, W. (1997): “Ethnicity in the electoral process: The 1992 general elections in Kenya,” *African Journal of Political Science*.
- PRENDERGAST, C., AND R. H. TOPEL (1997): “Favoritism in organizations,” *Journal of Economic Perspectives*, 104, 958–978.

BIBLIOGRAPHY

- RAVALLION, M., AND Q. WODON (1999): “Does child labor displace schooling? Evidence on behavioral responses to an enrollment subsidy,” Policy Research Working Papers 2116, World Bank.
- TAUCHEN, H. V., A. D. WITTE, AND S. K. LONG (1991): “Domestic violence: A non-random affair,” *International Economic Review*, 32, 491–511.
- UNICEF (2011): “<http://www.unicef.org/infobycountry>,” .
- WASHINGTON, E. L. (2008): “Female socialization: How daughters affect their legislator fathers’ voting on women’s issues,” *The American economic review*, 98(1), 311–332.
- WORLD HEALTH ORGANIZATION (2005): “Multi-country study on women’s health and domestic violence against women,” Summary report.