

Hjort, Jonas

Working Paper

Ethnic Divisions and Production in Firms

CESifo Working Paper, No. 4449

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Hjort, Jonas (2013) : Ethnic Divisions and Production in Firms, CESifo Working Paper, No. 4449, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/89778>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Working Papers

www.cesifo.org/wp

Ethnic Divisions and Production in Firms

Jonas Hjort

CESIFO WORKING PAPER NO. 4449

CATEGORY 4: LABOUR MARKETS

OCTOBER 2013

Presented at CESifo Area Conference on Applied Microeconomics, March 2013

An electronic version of the paper may be downloaded

- *from the SSRN website:* www.SSRN.com
- *from the RePEc website:* www.RePEc.org
- *from the CESifo website:* www.CESifo-group.org/wp

CESifo

Center for Economic Studies & Ifo Institute

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. 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.

JEL-Code: D030, D220, D240, D610, D640, D740, F630, J240, J330, J710, O120, O140.

Jonas Hjort
Columbia University / USA
hjort@columbia.edu

October 17, 2013

Winner of the 2013 CESifo Prize for best paper in Applied Microeconomics. hjort@columbia.edu. I am very grateful for invaluable guidance and support from Ted Miguel, David Card, Stefano DellaVigna and Fred Finan. I also thank Miguel Almunia, Vladimir Asriyan, Ivan Balbuzanov, Oriana Bandiera, Oystein Daljord, Andreas Fagereng, Willa Friedman, Francois Gerard, Bryan Graham, Pat Kline, Ed Lazear, Rocco Macchiavello, Jamie McCasland, Justin McCrary, Enrico Moretti, Kaivan Munshi, Takeshi Murooka, Omar Nayeem, Paul Niehaus, Matthew Rabin, Gautam Rao, Imran Rasul, James Robinson, Ana Rocca, Antonio Rosato, Yona Rubinstein, Michel Serafinelli, Lowell Taylor, Rebecca Thornton, Victoria Vanasco, Chris Woodruff and seminar participants at UC Berkeley, Chicago Booth, Columbia GSB / Economics, Cornell, UC Davis ARE, Harvard / MIT, Kellogg, University of Michigan, Penn State, Stanford GSB, Stockholm IIES, Warwick, Wharton, the IGC Growth Week, the NBER workshop on Development and Organizational Economics, and the Working Group on African Political Economy for comments and suggestions. Jason Luk provided excellent research assistance. I thank CEGA, IBER and the Russel Sage Foundation for financial support. All errors are my own.

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, 1999). 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 13 months of the sample period, processors were paid a piece rate based on own output and suppliers a piece rate based on total team output. Low supply of roses to downstream workers of the rival ethnic group thus implied lower supplier pay.

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 effect of team-level ethnic diversity on output, and the majority of this paper’s results, can thus be identified by comparing the output of teams of different ethnic compositions. I include person-position fixed effects throughout for consistence – and thus focus on variation in teams’ ethnic composition, controlling for their worker productivity composition – and refer to the regressions without fixed effects where relevant.

Two natural experiments during the sample period allow me to go further. During the second

part of the sample period, 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 part of the sample period, 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 degree of conflict between groups; 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 ethnic 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), Chen and Li (2009), and others. The model predicts that discriminatory suppliers in mixed teams will lower total output by “misallocating” 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 intensified 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.² About 88 percent of the output gap between horizontally mixed and homogeneous teams is due to vertical misallocation and 12 percent

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.

²This paper attempts to estimate and explain output differences between homogeneous and diverse production units, not their welfare implications. With the data and variation used here I am not able to determine if workers are worse off, in utility terms, in diverse teams.

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 24 percent more than processors of the other ethnic group, a gap exceeding one standard deviation of processor output.

In the second main result of the paper, I find that the output gap between homogeneous and diverse teams nearly doubled when political conflict between the Kikuyu and Luo ethnic blocs intensified in early 2008. The reason appears to be an increase in suppliers' taste for ethnic discrimination: while the decrease in diverse teams' output is driven by the output of processors who are not of the supplier's ethnic group, the output of processors in mixed teams who are of the same ethnic group as the supplier in fact increased significantly in early 2008, as also predicted by the model. 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 due to a reduction in horizontal misallocation: a 30 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 across 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, and particularly large effects of conflict on discriminatory work-place behavior among young workers. 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 I study, biased upstream workers continued to derive less benefit from flowers supplied to pairs of processors that included non-coethnics under team pay, and output in vertically mixed teams was 16 percent lower than in homogeneous teams after team pay was introduced. I discuss the sample plant's chosen response to ethnic diversity distortions, and how other firms respond.

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.³ 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, 2009; Mas and Moretti, 2009). The difference between the findings of Bandiera, Barankay and Rasul (2009) 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 Europe.⁴ Burgess et al. (2013) 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,⁵ 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

³Fisman, Paravisini and Vig (2012) show that cultural proximity between loan officers and borrowers increases the efficiency of credit allocation at an Indian bank. There is also 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 (1999) provides an interesting theoretical discussion of the potential costs and benefits of diversity in joint production situations. Prat (2002) shows that managers optimally hire same-type agents when the team's payoff function is supermodular and different types of agents when the team's payoff function is submodular. Hamilton, Nickerson and Owan (2012) 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.

⁴Note, however, that Hedegaard and Tyran (2011) find that in Denmark “taste-based discrimination is common but remarkably responsive to the price of prejudice” (Hedegaard and Tyran, 2011, p. 1), through an experiment in which youths packed letters with either coethnic or non-coethnic neighboring co-workers. Hong, Karaca-Mandic and Maestas (2008) find the same “negative peer effect and positive leader effect” (in race) as this paper when studying promotions in the U.S. army.

⁵Alesina, Michalopoulos and Papaioannou (2012) find evidence that inequality across ethnic groups can in itself hamper development.

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) and Macchiavello and Morjaria (2013) 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 (Boisjoly et al., 2006; Rotemberg, 2006; Bauer, Cassar and Chytilová, 2011; Jakiela, Miguel and te Velde, 2011).⁶

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) provide 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 related to La Ferrara (2002) who shows that ethnically diverse Kenyan 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 what drives heterogeneity in distortionary attitudes towards non-coethnics, and section 7 how firms respond to ethnic diversity distortions. Section 8 concludes.

⁶The connection between this paper's results and those of Shayo and Zussman (2011), who find that judges in Israel display greater in-group favoritism during periods of increased conflict between groups, is particularly interesting.

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 tribe, the Luo. Most Kenyan tribes have aligned themselves with one of the two associated camps. (For example, in the 2007 presidential election, exit polls suggested that 94 percent of Kikuyus and 88 percent of Merus voted for the Kikuyu incumbent, while 98 percent of Luos and 75 percent of Luhyas voted for the Luo challenger (Gibson and Long, 2009)). 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).⁷

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). Flower farms are part of the fastest growing sub-sector of the Kenyan economy: around 50,000 Kenyans are employed in floriculture, and 500,000 in associated industries (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 a gated community. 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.

2.2 Data

The sample farm primarily produces roses. My primary data source is records of daily processor output from 2007 and 2008. The quantities produced were recorded on paper by the farm for remuneration purposes and subsequently converted to electronic format by the research team (after 2008).

Workers rotate between teams. 28281 different teams are observed during the sample period.

⁷I 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”. Focusing on individual tribes instead gives similar results – see section 4.

Individual workers are observed on 92 different teams on average, and nearly all workers are observed in both positions (supplier and processor). A given team is observed working together for 7 consecutive days on average, but there is substantial variation in the length of team spells. The same is true for individual work spells. On average, workers are observed working for 18 days followed by two leave days.

A survey provides additional information about workers' experience, ethnicity, birthplace and other background information. There are 924 packing plant workers in total. 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.

2.3 Organization of production at the plant

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.

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.⁸ 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.

2.4 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. When a worker takes leave, another worker returning from leave joins the two remaining workers. When the first worker returns to work, she then joins another team with an open position. By including person-position fixed effects in the regressions, I can therefore control for any differences in productivity between the types of workers that end up in homogeneous versus mixed teams. The effect of ethnic diversity on output is then identified off of variation in teams' ethnic composition, controlling for their worker productivity composition.

⁸Workers were additionally paid a small fixed component.

In fact it is ex ante difficult to see how the team assignment system in use at the plant could lead to systematic correlation between the characteristics of the workers in a team. 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.

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, 49.69 percent horizontally mixed, and 24.85 percent vertically mixed, if assignment was random. The percentages observed in the data are 25.64/49.61/24.76 during the pre-conflict period, 27.38/48.35/24.26 during the conflict period and 25.32/49.26/25.42 during the team pay period.⁹ It is clear that workers are not assigned to, or sort into, teams based on ethnicity.¹⁰

A possible concern is that the underlying 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,¹¹ and that the skills of the three workers in a team interact in ways that are not adequately captured by additive person-position fixed effects. A formal test of quasi-random assignment is in table 2.¹² 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. I thus use a periodical “snapshot” of data in the table: team compositions on the first day of every month.¹³ For 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,¹⁴ and little ability to do so given their limited knowledge of worker characteristics and the plant’s leave and rotation

⁹The 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).

¹⁰Appendix figure 1 displays the distribution of co-workers’ tribe (and other characteristics) across Kikuyu and Luo suppliers, during each of the three periods. The distributions are essentially identical.

¹¹Suppose for example that individuals are equally productive in homogeneous and diverse teams but prefer interacting with coethnics, as in Becker (1957). In that case it may be that supervisors assign well-liked, high-productivity workers to desirable homogeneous teams.

¹²Appendix 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.

¹³The 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 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.

¹⁴Supervisors were rarely, if ever, promoted, and their pay did not depend on performance.

system.¹⁵ 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).

In the analysis to follow I include person-position fixed effects wherever possible, for completeness. Results excluding person-position fixed effects are nearly identical and available from the author upon request.

2.5 Investigating the shape of the production function

In order to correctly model and 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 :

$$q_{p,d} = \alpha_i^p D_{i,d}^p + \beta_j^o D_{j,d}^o + \alpha_k^s D_{k,d}^s + \varepsilon_{p,d} \quad (1)$$

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.¹⁶

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), and (c) other processor permanent productivity (across panel A and B).¹⁷

¹⁵Team 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 (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).

¹⁶Two 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 et al., 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.

¹⁷ $\hat{\alpha}^p$ 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

The evidence in figure 4 is consistent with a range of possible production functions, which I do not attempt to distinguish between here. What is key for the purposes of this paper is to note that processor output is increasing in both processor and supplier productivity throughout the range, suggesting that processors are always better off with greater supply of intermediate flowers.¹⁸ Figure 4 also shows a small but negative effect of other processor’s productivity, indicating that upstream workers consider the benefits of supply to both downstream workers when making their supply decisions.

3 Discrimination: Theoretical Framework

3.1 Set-up

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 total output thereby lowered, if suppliers have discriminatory social preferences. I model the possibility of lower output in ethnically diverse teams as driven by tastes for three reasons. First, the simple, triangular structure of production at the plant 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 in the context considered. Second, taste-based misallocation is a particularly important (possible) source of lower output in ethnically mixed production units because the distortions implied would likely affect the aggregate economy considerably (Becker, 1957). Third, I will argue that the variation available during the period observed allows me to distinguish taste-based discrimination from other forms of diversity effects. I test the model’s predictions in the next section and consider the ability of non-taste-based ethnic diversity effects to explain the results in section 5.

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 herself. Let processor output depend on supplier effort and ability, e_{sp} 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

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.

¹⁸The evidence in figure 4 is for example not consistent with a Leontief production function or one in which the supplier can dictate work-speed regardless of the processor’s desired speed. One possible 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.

simultaneously.¹⁹ Finally, assume that the supplier attaches weight θ_p to the utility of processor p .²⁰ Suppliers with a different weight for coethnics and non-coethnics have discriminatory preferences.²¹

Note that θ_p can be either positive or negative. 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 $0 < \theta_1 < \theta_2$.²²

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

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

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:

$$\begin{aligned} \text{Max}_{e_{s1}, e_{s2}} w(f(e_{s1}, \alpha_s, e_1, \alpha_1) + f(e_{s2}, \alpha_s, e_2, \alpha_2)) - d(e_{s1} + e_{s2}) \\ + \theta_1(2wf(e_{s1}, \alpha_s, e_1, \alpha_1) - d(e_1)) + \theta_2(2wf(e_{s2}, \alpha_s, e_2, \alpha_2) - d(e_2)) \end{aligned} \quad (3)$$

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 (HM, C) or may

¹⁹In reality, supply and processing decisions take place continuously throughout the work-day, but when the (data) time-unit to which the model must be compared is a whole work-day a reasonable simplification is to assume a single decision on the part of a processor and two on the part of the supplier (one for each processor). Given that the model is "static", it makes sense to assume simultaneous moves as any co-dependence between a supplier's and a processor's effort level likely runs both ways. It turns out that a model with simultaneous moves in fact describes workers' response to the introduction of team pay better than would, for example, a model in which the supplier expected her effort to influence processors' effort levels, as discussed below.

²⁰This formulation follows Becker (1974), Charness and Rabin (2002), Chen and Li (2009), and others.

²¹I focus on discriminatory behavior on the part of the supplier in the model and the empirics for several reasons. First, while a model in which both the supplier and the processors had social preferences would generate similar predictions to the one presented here, such a model would be less tractable to solve and interpret. Second, while it is not possible to distinguish suppliers behaving differently across teams of different ethnic compositions from processors doing so by comparing average (individual or team) output in homogeneous versus diverse teams, the graphs in figure 4 and the optimal assignment exercise in section 7 suggest that suppliers' behavior is more important for output than processors. Third, while processors influence the supplier's pay only through effort, the supplier influences processors' pay also through simple, and presumably salient, differential supply decisions. It turns out that a model focused on supplier behavior describes the output response to the introduction of team pay better than would a model in which processors also (act as if they) have discriminatory preferences, as discussed below.

²²If 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".

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 do not distinguish between the two specific ethnic groups here; “homogeneous teams” may for example be either Kikuyu-Kikuyu-Kikuyu or Luo-Luo-Luo.²³

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 in homogeneous than in mixed teams: $Q_H > Q_{VM}$ and $Q_H > Q_{HM}$.²⁴

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

²³I highlight the additional cases to be considered if ability or taste for discrimination differs across the two ethnic groups and empirically test for these scenarios below.

²⁴Note 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 – as would appear less reasonable – 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.

$$\partial q_{VM}/\partial \alpha_s > \partial q_{HM,NC}/\partial \alpha_s.$$

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 because the relative benefits of supplying such processors go up. 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. While it is of course also possible that output in homogeneous teams is affected by conflict,²⁵ the focus here is on differences in output between teams of different ethnic compositions. To summarize, the models predicts that $\partial q_{HM,C}/\partial \theta_{NC} < 0 = q_H/\partial \theta_{NC} < \partial q_{VM}/\partial \theta_{NC} < \partial q_{HM,NC}/\partial \theta_{NC}$.

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:

$$\begin{aligned} \text{Max}_{e_{s1}, e_{s2}} \quad & w(f(e_{s1}, \alpha_s, e_1, \alpha_1) + f(e_{s2}, \alpha_s, e_2, \alpha_2)) - d(e_{s1} + e_{s2}) \\ & + (\theta_1 + \theta_2)w(f(e_{s1}, \alpha_s, e_1, \alpha_1) + f(e_{s2}, \alpha_s, e_2, \alpha_2)) - \theta_1 d(e_1) - \theta_2 d(e_2) \end{aligned} \quad (4)$$

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 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 the same 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: $Q_{HM}^{TP} \geq Q_{HM}$.²⁶

²⁵For example due to disruption effects (Ksoll, Macchiavello and Morjaria, 2010) or changes in individuals' weight on coethnics' utility (Eifert, Miguel and Posner, 2010).

²⁶It is not the case in this framework that more is supplied to non-coethnic processors in horizontally mixed teams

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

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 respond similarly to non-coethnic co-workers 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 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. Excluding person-position fixed effects has little influence on the results, as expected. The output of processors in vertically mixed teams is eight percent lower than that of proces-

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.

sors in homogeneous teams, an output gap that is also reflected in the total output of vertically mixed teams. It appears that upstream workers discriminate against non-coethnics downstream by undersupplying them, as predicted by the model. 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 eight 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. The output gap between coethnic and non-coethnic processors in horizontally mixed teams exceeds one standard deviation of processor output. That processor output is lower if the other processor is of the same ethnicity as the supplier points to horizontal favoritism. As Becker (1957) emphasized, favored workers benefit from discrimination against non-favored workers.

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 are 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 three 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. We can then decompose the output gap between homogeneous and horizontally mixed teams: 12 percent would be due to the effect of horizontal misallocation and 88 percent due to vertical misallocation.²⁷ While the magnitude of the “misallocation multiplier” associated with horizontal 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 $\partial q_{HM,C}/\partial \alpha_s > \partial q_{VM}/\partial \alpha_s >$

²⁷This decomposition is illustrative in that it ignores the convexity of effort costs, and it is not clear that the effect of vertical and horizontal misallocation is “additive”.

$\partial q_{HM,NC}/\partial\alpha_s$ and $\partial q_H/\partial\alpha_s > \partial q_{VM}/\partial\alpha_s > \partial q_{HM,NC}/\partial\alpha_s$. Note that 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 two other workers in the team. The reason appears to be that upstream workers undersupply non-coethnics and distort their supply of intermediate flowers to benefit coethnics downstream. Since such behavior also lowers the pay of the supplier, the results are consistent a willingness to pay to discriminate on the part of upstream workers.²⁸

In section 5 I consider alternative theories that predict negative ethnic diversity effects but do so for reasons unrelated to discriminatory preferences. I now consider the extent to which explanations other than a negative output effect of ethnic diversity 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, not the underlying mechanism.

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 plant to provide arguably comparable evidence. The analysis explores what happens when a worker is replaced by another worker of the same average 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.²⁹ 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 explain the estimates focusing on 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 six 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 eight 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

²⁸The contrast between these findings and those of Bandiera, Barankay and Rasul (2009) 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 (2009) 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.

²⁹The 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.

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 seven percent when a team goes from being homogeneous to vertically mixed, and by three 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.

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. In the three panels, I 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. 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.

A comparison of teams of different compositions as defined by other worker characteristics – such as for example gender – show minor output differences and no sign of systematic, differential behavior when working with “in-group” versus “out-group” co-workers (results available from the author upon request).

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 show that, 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. Interestingly, these results highlight that the relevant ethnic relations in Kenya are historical or political constructs rather than primordially defined—consistent with the response to intensified political conflict shown below. While the Luo are Nilotic, both the Luhya and Kikuyu are Bantu, but the Luo and Luhya are long-term political allies in rivalry with the Kikuyu (see also Posner, 2004; Shetler, 2010).

So far we have seen that in the context of factory production in Kenya, output is significantly higher in ethnically homogeneous teams, and a range of robustness checks documenting that the output gap is caused by ethnic diversity itself. 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.

4.2 Ethnic conflict and the impact of diversity 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,³⁰ but supervisors at the sample plant reported that logistics and worker absence at the farm were largely unaffected and that production continued as usual. Because the workers live on the farm in a gated community it was safest to remain on the farm.

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.³¹

In table 8, the difference in output between mixed and homogeneous teams before and after conflict began is compared.³² 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 – a

³⁰Dupas and Robinson (2012) document, for example, a dramatic fall in income and consumption for the rural poor in Western Kenya during the crisis.

³¹Eifert, Miguel and Posner (2010) show that ethnic identities can vary over time, and Charness and Rabin (2002) and others show that preferences towards others generally depend on their behavior. 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.

³²Data 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 .

(non-) response that may cover up countervailing effects of conflict on productivity in homogeneous teams,³³ but the focus here is on the difference in output between homogeneous and mixed teams. 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 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, and that the economic costs of ethnic diversity vary with the political environment. Note that, as seen figure 2, the increased output gap between homogeneous and mixed teams shows no sign of decay in the nine months after conflict ended (controlling for the effect of changes in suppliers' incentive to discriminate as discussed in the next section). Periods of increased antagonism may entail significant hidden economic costs if reversion in taste for discrimination is slow (or does not occur).

Firms may be forced to take measures to limit distortions that arise from internal, ethnic discrimination, especially in times of conflict. In section 7 I analyze the firm's chosen response to the 2008 political conflict in Kenya. 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.

4.3 Firm's response to ethnic diversity distortions and the impact 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. When a processors is paid in part based on the output of the other processor, freeriding is expected to negatively affect output in all teams, but the framework above predicts an offsetting positive effect in horizontally mixed teams. Under team pay, suppliers are unable to influence the relative pay of the two processors through

³³Disruption effects, as Ksoll, Macchiavello and Morjaria (2010) find for some flower farms during the crisis, would likely affect all teams. If Kenyans increasingly identify with coethnics during times of heightened political competition between groups, as the results of Eifert, Miguel and Posner (2010) would suggest, weight on the utility of coethnics may have increased.

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 through the remainder of 2008 as a single team pay period.³⁴ 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). Comparing the second and third periods, the figure clearly indicates that the introduction of team pay had a positive effect on output in horizontally mixed teams.

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.³⁵

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 30 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

³⁴In 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 noted above, the output gap between teams of difference ethnicity configurations showed no sign of reversion later in 2008.

³⁵The modest magnitude of this effect is noteworthy and interesting in itself. As 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.

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 output 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.³⁶ It is difficult to imagine a standard economic model of joint production that would predict an increase in output when team pay is introduced.

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. It thus appears that ethnic diversity affects how firms organize production. I discuss the sample farm's chosen response in section 7. In the next section I analyze 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 be *efficient* if individuals are simply more productive when collaborating with others of their own ethnic group. If Kenyans have a taste for discrimination, allocative efficiency in the broader economy is likely affected. The optimal policy and firm response will also depend on the underlying mechanism.

The sources of non-taste-based, negative ethnic diversity effects discussed in the literature can be classified into three broad categories:³⁷

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; Fisman, Paravisini and Vig, 2012). Such informational effects can lead to higher supply to downstream coethnics. Note that, in the context of the sample

³⁶Note 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.

³⁷Habyarimana, Posner and Weinstein (2007) uses a similar classification, but their focus is on explaining why public goods provision is lower in diverse societies.

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 and little communication is required, but complicated forms of peer effects could explain the results in table 4.
3. *“Cooperational” diversity effects* arise if coethnics are better able to sustain cooperation (Kandel and Lazear, 1992; Greif, 1993; 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). 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 see how a cooperational (or other non-taste-based) ethnic diversity mechanism can simultaneously explain a decrease in mixed teams’ output during conflict, an increase in the output of only those coethnic processors working alongside a non-coethnic other processor during conflict, equalization of processors’ output in horizontally mixed teams when team pay is introduced, and the other results of this paper taken together.

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.³⁸

³⁸As discussed above, a model in which both the supplier and the processors (act as if they) have social preferences would generate similar predictions to the one presented here. The fact that there is no significant difference in output between coethnic and non-coethnic processors in horizontally mixed teams is more consistent with the model presented here, however. If processors (act as if they) attach differential weight to the utility of coethnic and non-coethnic co-workers, in a horizontally mixed team we would under team pay expect the effort and output of a coethnic

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 workers respond to policies introduced by firms 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 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 Heterogeneity in Distortionary Discrimination

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. Analyzing differences in discriminatory behavior across workers, and how taste for discrimination responds to past experiences, will help us understand how distortionary preferences are shaped.

Figure 9 plots the distribution of suppliers' "discrimination coefficient" – the difference in output between homogeneous and (vertically and horizontally) mixed teams supplied by a specific worker, before and after conflict began. It appears that most suppliers discriminate against non-coethnic processors during the pre-conflict period. It is also clear that there is substantial heterogeneity in discriminatory behavior.

In column 1 of table 10 I investigate further by regressing individuals' discrimination coefficient on their characteristics and measures of past exposure to non-coethnics. The discrimination coefficient (outcome) is here constructed for the second half of 2007, while past experiences are measured during the first half of 2007. The results show that females are significantly less discriminatory than males, and young workers are less discriminatory than older workers. 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 (see e.g. Buckley, 1997).

The coefficient on "percent of workdays spent in mixed teams" (during the first half of 2007) is negative, while the coefficient on "average discrimination coefficient of non-coethnics the worker was supplied by" (during the first half of 2007) is positive. These signs are as we would expect if spending time with non-coethnics reduces discriminatory attitudes (Boisjoly et al., 2006), while being discriminated against exacerbates such attitudes. But the effects are imprecisely estimated

processor (who "works for" one coethnic and one non-coethnic, in addition to herself) to be higher than that of a non-coethnic processor (who works for two non-coethnics and herself). More 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.

and do not reach statistical significance.³⁹

In figure 9 we further see that intensified conflict between the ethnic groups in Kenya led to an increase in the discrimination coefficient of most upstream workers. In the empirical appendix, I take advantage of the worker rotation system to bound the magnitude of the average decrease in the weight attached to non-coethnics' utility. Figure 9 also shows a notable widening of the distribution of discrimination coefficients, however. Some suppliers responded more to conflict than other.

Column 2 in table 10 uses two observations per worker, one from (the second half of) 2007 and one from the conflict period, to investigate factors that influenced changes in an individual's propensity to discriminate when conflict began. We see that the increase in discriminatory behavior during conflict was (marginally) significantly greater among young workers. The results in table 10 thus suggest that youth start out relatively tolerant, but that the attitudes of the young towards non-coethnics may respond more negatively to conflict.

The signs of the coefficients on the "past exposure" variables suggest that those who had spent more time with non-coethnics in the past, and those who had been supplied by less discriminatory non-coethnics, may respond more strongly to conflict, but these results are not statistically significant.

The results discussed in this section paint a consistent albeit suggestive picture of how distortionary attitudes towards individuals of other ethnic groups are shaped. A serious episode of violent, political conflict between the Kikuyu and Luo blocs led to a significant shift in taste for ethnic discrimination in Kenya – a shift that in the data show no sign of decay or reversion in the nine months after conflict ended. Perhaps most importantly, it appears that younger individuals have more malleable social preferences that respond strongly to episodes of conflict. Conflict may thus entail significant hidden economic costs through its influence on distortionary social preferences.

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 Response to ethnic diversity distortions at the sample plant

Segregating workers of different ethnicities appears 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 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

³⁹As the unit of observation in table 10 is an individual, the sample size is limited.

associated with teams of each of 3 ethnicity configurations, 18 productivity configurations and 63 ethnicity-productivity configurations.⁴⁰ 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).⁴¹

The optimal assignments and associated expected output gains are shown in table 11.⁴² 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 – more than 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.⁴³ 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.

The fact that managers chose not segregate workers, even after conflict led to a dramatic drop in productivity in mixed teams, suggests that they expect there to be costs associated with segregation. It is possible that managers believe that interacting with individuals of other ethnic groups will in itself dampen discriminatory attitudes over time, as Boisjoly et al. (2006) find for the U.S.⁴⁴ But as we saw in table 10, I find no evidence of such an effect at the plant studied.

While the costs of ethnic segregation are likely incurred primarily by society at large,⁴⁵ the firm analyzed is of a size and form indicating that it would perhaps carry some of the costs of segregating workers itself. The optimal assignment to teams computed in table 11 is “out-of-sample” and therefore do not rule out the possibility that complete segregation between the two ethnic groups would have a negative effect on output. 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 “extreme” assignments. In the context of the

⁴⁰ $63 = [3 * ((3 * (3 + 1))/2)] + 3^3 + [3 * ((3 * (3 + 1))/2)]$. 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.

⁴¹Bhattacharya and Dupas (2012) 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.

⁴²The procedure used is explained in the empirical appendix. “Optimal” here means output-maximizing, as inferred from the data. The output-maximizing solution may be undesirable for other reasons discussed below.

⁴³An 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)}.

⁴⁴Boisjoly et al. (2006) find that white American college students become more friendly towards and supportive of African American students after spending time with a black roommate.

⁴⁵Alesina and Zhuravskaya (2011) find, for example, that more ethnically segregated countries have lower quality of government. The possible externalities of segregation is one justification for the hard-to-enforce laws against some forms of segregation that apply to firms in many countries.

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 antagonism outside of the production halls.

7.2 Response-inclusive implications of ethnic diversity distortions in the economy

On the basis of the data analyzed in this paper, it is ultimately surprising that a supposedly profit-maximizing firm chose to leave large productivity gains “on the table” by not segregating workers of different ethnicities (or allowing workers to self-segregate). In most scenarios, we expect firm responses and general equilibrium effects to cushion the impact of micro-level distortions on observed productivity and output. Consider the following possibilities.

Competition Becker (1957) pointed out that discriminatory employers should go out of business as their profits suffer. A priori, the same argument holds 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).⁴⁶ As the literature in macroeconomics on across-firm misallocation has highlighted, it is not necessarily the most productive firms that survive in poor countries’ economies (Banerjee and Moll, 2009; Hsieh and Klenow, 2009).⁴⁷

Composition of firms As for other agribusiness firms, the location of the sample farm was partly determined by availability of suitable land. The location where the farm is situated is ethnically mixed, partly for historical reasons and partly because the availability of jobs has attracted Kenyans from other parts of the country. It appears that other large firms in ethnically mixed parts of the country also do not engage in systematic ethnic discrimination in hiring. While it is difficult for large firms to do so, smaller firms likely suffer less from distortionary discrimination at the workplace, in part because they hire through social and ethnic networks.

Prices and technology As we have seen, managers at the sample plant did respond to the increase in distortionary discrimination when conflict began: the introduction of team pay for processors was likely motivated by the decrease in productivity in ethnically mixed teams in early 2008. It appears that managers preferred adjusting the “price” of intermediate flowers delivered to non-coethnic processors relative to that of flowers delivered to coethnic processors over a possible technological response (such as moving away from three-person joint production) or changing the composition of teams. In an different setting in urban Kenya, La Ferrara (2002) also observes a price/incentive system response to ethnic diversity: ethnically diverse cooperatives in Nairobi are more likely to adopt group-pay. Note, however, 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.

⁴⁶Entry- and exit- barriers are significant: large areas of land and expensive, complicated equipment is needed to produce roses and other cut flowers.

⁴⁷Futhermore, 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).

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 output effect of ethnic diversity in teams of co-workers in Kenya. I do so using two years of daily output data for 924 workers, almost equally drawn from two rival ethnic groups, 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 identify the effect of ethnic diversity on output off of variation in teams’ ethnic composition, controlling for their worker productivity composition. 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 model. The reason appears to be that workers’ taste for discrimination against non-coethnic co-workers increased.

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 (see also Rohner, Thoenig and Zilibotti, 2013).

9 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 $-\frac{1}{2} < \theta_p < \frac{1}{2}$.⁴⁸ 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:

$$\begin{aligned} \text{Max}_{e_1} \quad & 2w(e_1 \alpha_1)^\beta (e_{s1} \alpha_s)^\gamma - \frac{1}{2}e_1^2 \\ \text{s.t.} \quad & e_1 \geq 0 \end{aligned} \tag{5}$$

which gives

$$e_1 = \left(2w\beta(e_{s1} \alpha_s)^\gamma \alpha_1^\beta \right)^{\frac{1}{2-\beta}} \tag{6}$$

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:

$$\begin{aligned} \text{Max}_{e_{s1}, e_{s2}} \quad & w((e_1 \alpha_1)^\beta (e_{s1} \alpha_s)^\gamma + (e_2 \alpha_2)^\beta (e_{s2} \alpha_s)^\gamma) - \frac{1}{2}(e_{s1} + e_{s2})^2 \\ & + \theta_1 \left(2w(e_1 \alpha_1)^\beta (e_{s1} \alpha_s)^\gamma - \frac{1}{2}e_1^2 \right) + \theta_2 \left(2w(e_2 \alpha_2)^\beta (e_{s2} \alpha_s)^\gamma - \frac{1}{2}e_2^2 \right) \\ \text{s.t.} \quad & e_{s1} \geq 0 \text{ and } e_{s2} \geq 0 \end{aligned} \tag{7}$$

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

$$(e_{s1} + e_{s2}) = (1 + 2\theta_1) w(e_1 \alpha_1)^\beta \gamma (e_{s1} \alpha_s)^{\gamma-1} \tag{8}$$

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

$$e_{s1} = \left(\frac{(1 + 2\theta_1) w(e_1 \alpha_1)^\beta \gamma (\alpha_s)^{\gamma-1}}{1 + \left(\frac{1+2\theta_1}{1+2\theta_2} \right)^{\frac{1}{\gamma-1}} \left(\frac{e_1 \alpha_1}{e_2 \alpha_2} \right)^{\frac{\beta}{\gamma-1}}} \right)^{\frac{1}{2-\gamma}} \tag{9}$$

Because the supplier considers the pay-off (from own pay and processors' utility) of supply to

⁴⁸If this restriction is violated corner solutions arise.

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

$$q_1^* = \frac{k_q \alpha_s^{\frac{2\gamma}{2-\beta-\gamma}} \alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (1+2\theta_1)^{\frac{2\gamma}{2-\beta-2\gamma}}}{\left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} (1+2\theta_1)^{\frac{2-\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} (1+2\theta_2)^{\frac{2-\beta}{2-2\gamma-\beta}} \right)^{\frac{\gamma}{2-\beta-\gamma}}} \quad (10)$$

$$Q^* = \frac{k_q \alpha_s^{\frac{2\gamma}{2-\beta-\gamma}} \left(\alpha_1^{\frac{2\beta}{2-\beta-2\gamma}} (1+2\theta_1)^{\frac{2\gamma}{2-\beta-2\gamma}} + \alpha_2^{\frac{2\beta}{2-\beta-2\gamma}} (1+2\theta_2)^{\frac{2\gamma}{2-\beta-2\gamma}} \right)}{\left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} (1+2\theta_1)^{\frac{2-\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} (1+2\theta_2)^{\frac{2-\beta}{2-2\gamma-\beta}} \right)^{\frac{\gamma}{2-\beta-\gamma}}} \quad (11)$$

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.

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' 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:

$$\begin{aligned} \text{Max}_{e_1} \quad & w \left((e_1 \alpha_1)^\beta (e_{s1} \alpha_s)^\gamma + (e_2 \alpha_2)^\beta (e_{s2} \alpha_s)^\gamma \right) - \frac{1}{2} e_1^2 \\ \text{s.t.} \quad & e_1 \geq 0 \end{aligned} \tag{12}$$

which gives

$$e_1 = \left(w \beta (e_{s1} \alpha_s)^\gamma \alpha_1^\beta \right)^{\frac{1}{2-\beta}} \tag{13}$$

Under team pay the supplier solves

$$\begin{aligned} \text{Max}_{e_{s1}, e_{s2}} \quad & w((e_1\alpha_1)^\beta(e_{s1}\alpha_s)^\gamma + (e_2\alpha_2)^\beta(e_{s2}\alpha_s)^\gamma) - \frac{1}{2}(e_{s1} + e_{s2})^2 \\ & + w(\theta_1 + \theta_2)((e_1\alpha_1)^\beta(e_{s1}\alpha_s)^\gamma + (e_2\alpha_2)^\beta(e_{s2}\alpha_s)^\gamma) - \theta_1\frac{1}{2}e_1^2 - \theta_2\frac{1}{2}e_2^2 \\ \text{s.t.} \quad & e_{s1} \geq 0 \text{ and } e_{s2} \geq 0 \end{aligned} \quad (14)$$

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

$$e_{s1} + e_{s2} = w(1 + \theta_1 + \theta_2)(e_1\alpha_1)^\beta\gamma(e_{s1}\alpha_s)^{\gamma-1}\alpha_s \quad (15)$$

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

$$e_{s1} = \left(\frac{w(1 + \theta_1 + \theta_2)\gamma(e_1\alpha_1)^\beta\alpha_s^\gamma}{1 + \left(\frac{e_2\alpha_2}{e_1\alpha_1}\right)^{\frac{\beta}{1-\gamma}}} \right)^{\frac{1}{2-\gamma}} \quad (16)$$

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

$$q_1^{TP*} = \frac{k_q^{TP} \alpha_s^{\frac{\gamma}{2-\beta-\gamma}} \alpha_p^{\frac{2\beta}{2-\beta-2\gamma}} (1 + \theta_1 + \theta_2)^{\frac{\gamma}{2-\beta-\gamma}}}{\left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} \right)^{\frac{\gamma}{2-\beta-\gamma}}} \quad (17)$$

$$Q^{TP*} = k_q^{TP} \alpha_s^{\frac{\gamma}{2-\beta-\gamma}} \left(\alpha_1^{\frac{2\beta}{2-2\gamma-\beta}} + \alpha_2^{\frac{2\beta}{2-2\gamma-\beta}} \right)^{\frac{2-\beta-2\gamma}{2-\beta-\gamma}} (1 + \theta_1 + \theta_2)^{\frac{\gamma}{2-\beta-\gamma}} \quad (18)$$

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.

10 Empirical Appendix

10.1 Magnitude of the increase in taste for discrimination during conflict

By how much did suppliers' weight on the utility of non-coethnic downstream workers 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 if the model above holds, I can bound the impact of conflict on θ_{NC} by taking advantage of the plant's worker rotation system. 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 laid out in the theoretical appendix, the supplier ability term cancels out if we take the ratio of the two processors' output:

$$\frac{q_1}{q_2} = \left(\frac{\alpha_1}{\alpha_2} \right)^{\frac{2\beta}{2-\beta-2\gamma}} \left(\frac{1+2\theta_1}{1+2\theta_2} \right)^{\frac{2\gamma}{2-\beta-2\gamma}} \quad (19)$$

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 (in between dates $d = 0$ and $d = 1$). In the model above, the abilities of the two processors do not influence their relative output under one supplier compared to their relative output under another supplier:

$$\frac{q_{1,d=0}/q_{2,d=0}}{q_{1,d=1}/q_{2,d=1}} = \left(\frac{1+2\theta_C}{1+2\theta_{NC}} \right)^{\frac{4\gamma}{2-\beta-2\gamma}} \quad (20)$$

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

$$\frac{(q_{1,d=0}/q_{2,d=0})/(q_{1,d=1}/q_{2,d=1})}{(q_{1,d=0'}/q_{2,d=0'})/(q_{1,d=1'}/q_{2,d=1'})} = \left(\frac{1+2\theta_{NC}}{1+2\theta'_{NC}} \right)^{\frac{4\gamma}{2-\beta-2\gamma}} \quad (21)$$

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:

$$\log(q_{1,s=0}/q_{2,s=0}) = \lambda + \eta \log(q_{1,s=1}/q_{2,s=1}) + \varepsilon \quad (22)$$

The resulting $\hat{\lambda}$ 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 (22) on pre-conflict data gives $\hat{\lambda} = 0.27$. $\hat{\lambda}'$, from estimating (22) on data from the conflict period, is 0.4. Both estimates are significantly greater than zero at the 1% level.

Noting that $\hat{\theta}_C = \frac{1}{2} \left(\left(\exp(\hat{\lambda}) \right)^{\frac{2-\beta-2\gamma}{4\gamma}} (1 + 2\hat{\theta}_{NC}) - 1 \right)$, with $\hat{\lambda}$ 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. Suppose further that θ_C did not change when conflict began, as the results of table 8 suggest. Then,

$$1 = \frac{\hat{\theta}_C}{\hat{\theta}'_C} = \frac{\frac{1}{2} \left(\left(\exp(\hat{\lambda}) \right)^{\frac{2-\beta-2\gamma}{4\gamma}} (1 + 2\hat{\theta}_{NC}) - 1 \right)}{\frac{1}{2} \left(\left(\exp(\hat{\lambda}') \right)^{\frac{2-\beta-2\gamma}{4\gamma}} (1 + 2\hat{\theta}'_{NC}) - 1 \right)} \quad (23)$$

which gives

$$\theta'_{NC} = \frac{1}{2} \left(\frac{1 + 2\hat{\theta}_{NC}}{\left(\exp(\hat{\lambda}' - \hat{\lambda}) \right)^{\frac{2-\beta-2\gamma}{4\gamma}}} - 1 \right) = \frac{1}{2} \left(\frac{1 + 2\hat{\theta}_{NC}}{\left(\exp(0.13) \right)^{\frac{2-\beta-2\gamma}{4\gamma}}} - 1 \right) \quad (24)$$

In appendix figure 3, I plot θ'_{NC} against θ_{NC} for various combinations of β and γ . It is clear from the figure that the decrease in θ_{NC} – or put differently, the increase in taste for discrimination – required to explain the differential decrease in output in mixed teams during conflict is substantial.

10.2 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:

$$\text{Max}_{t_1, \dots, t_{3078}} Q = \sum_{i=1}^{3078} \bar{Q}_i t_i \quad (25)$$

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:

$$\begin{aligned} & \sum \{t_i | \text{there is 1 } w_j \text{ worker in } t_i\} \\ & + 2 \sum \{t_i | \text{there are 2 } w_j \text{ workers in } t_i\} \\ & + 3 \sum \{t_i | \text{there are 3 } w_j \text{ workers in } t_i\} = w_j \end{aligned} \quad (26)$$

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.

References

- Alesina, A., and Eliana La Ferrara.** 2005. "Ethnic Diversity and Economic Performance" *Journal of Economic Literature*, 43(3): 762–800.
- Alesina, A., and Enrico Spolaore.** 1997. "On The Number and Size of Nations" *Quarterly Journal of Economics*, 112(4): 1027–1056.
- Alesina, Alberto, and Ekaterina Zhuravskaya.** 2011. "Segregation and the Quality of Government in a Cross Section of Countries" *American Economic Review*, 101(5): 1872–1911.
- Alesina, Alberto F., Stelios Michalopoulos, and Elias Papaioannou.** 2012. "Ethnic Inequality" NBER wp 18512.
- Arcidiacono, P., Gigi Foster, Natalie Goodpaster, and Josh Kinsler.** 2011. "Estimating Spillovers using Panel Data, with an Application to the Classroom" *Quantitative Economics*.
- Bandiera, O., Iwan Barankay, and Imran Rasul.** 2005. "Social Preferences and the Response to Incentives: Evidence from Personnel Data" *Quarterly Journal of Economics*, 120(3): 917–962.
- Bandiera, O., Iwan Barankay, and Imran Rasul.** 2009. "Social Connections and Incentives in the Workplace: Evidence from Personnel data" *Econometrica*, 77: 1047–1094.
- Banerjee, A. V., and Benjamin Moll.** 2009. "Why does misallocation persist?" Mimeo, MIT.
- Barkan, J. D.** 2004. "Kenya After Moi" *Foreign Affairs*.
- Bauer, M., Alessandra Cassar, and Julie Chytilová.** 2011. "Warfare and Social Preferences in Children" *Mimeo, USF*.
- Becker, Gary S.** 1957. *The Economics of Discrimination* Univ. Chicago Pres.
- Becker, Gary S.** 1974. "A Theory of Social Interactions" *Journal of Political Economy*, 82(6): 1063–1093.
- Bhattacharya, D.** 2009. "Inferring Optimal Peer Assignment from Experimental Data" *Journal of the American Statistical Association*.
- Bhattacharya, Debopam, and Pascaline Dupas.** 2012. "Inferring welfare maximizing treatment assignment under budget constraints" *Journal of Econometrics*, 167(1): 168–196.
- Boisjoly, Johanne, Greg J. Duncan, Michael Kremer, Dan M. Levy, and Jacque Eccles.** 2006. "Empathy or Antipathy? The Impact of Diversity" *American Economic Review*, 96(5): 1890–1905.
- Buckley, Stephen.** 1997. "Youth in Kenya Now Feel Few Tribal Ties" *Washington Post*. <http://www.washingtonpost.com/wpsrv/inatl/longterm/africanlives/kenya/kenya.htm>.

- Burgess, Robin, Remi Jedwab, Edward Miguel, Ameet Morjaria, and Gerard Padró-i-Miquel.** 2013. “The Value of Democracy: Evidence from Road Building in Kenya” National Bureau of Economic Research Working Paper No. 19398.
- Carrell, S., B. Sacerdote, and J. West.** 2011. “From Natural Variation to Optimal Policy? The Importance of Endogenous Peer Group Formation” *forthcoming, Econometrica*.
- Charness, Gary, and Matthew Rabin.** 2002. “Understanding Social Preferences with Simple Tests” *Quarterly Journal of Economics*.
- Chen, Yan, and Sherry Xin Li.** 2009. “Group Identity and Social Preferences” *American Economic Review*, 99(1): 431–457.
- Dupas, Pascaline, and Jonathan Robinson.** 2012. “The (hidden) costs of political instability: Evidence from Kenya’s 2007 election crisis” *Journal of Development Economics*, 99(2): 314–329.
- Easterly, W., and Ross Levine.** 1997. “Africa’s Growth Tragedy: Policies and Ethnic Divisions” *Quarterly Journal of Economics*.
- EDRI.** 2008. “Ethiopian Development Research Institute: Micro Evidence on the Development of Cut-flower Industry in Ethiopia”
- Eifert, B., Edward Miguel, and Daniel N. Posner.** 2010. “Political Competition and Ethnic Identification in Africa” *American Journal of Political Science*, 54(2): 494–510.
- Fisman, Raymond, Daniel Paravisini, and Vikrant Vig.** 2012. “Cultural Proximity and Loan Outcomes” Mimeo, Columbia University.
- Gibson, C., and James Long.** 2009. “The presidential and parliamentary elections in Kenya, 2007” *Electoral Studies*, 28(3): 497–502.
- Graham, B, Guido Imbens, and Geert Ridder.** 2011. “Measuring the effects of segregation in the presence of social spillovers: A nonparametric approach” National Bureau of Economic Research Working Paper No. 16499.
- Greif, Avner.** 1993. “Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders’ Coalition” *American Economic Review*, 83(3): 525–548.
- Habyarimana, J., Daniel N. Posner, and Jeremy Weinstein.** 2007. “Why Does Ethnic Diversity Undermine Public Goods Provision?” *American Political Science Review*.
- Hamilton, Barton H, Jack A Nickerson, and Hideo Owan.** 2012. “Diversity and Productivity in Production Teams” *Advances in the Economic Analysis of Participatory & Labor-Managed Firms*, 13: 99–138.
- Hedegaard, Morten, and Jean-Robert Tyran.** 2011. “The Price of Prejudice” The Rockwool Foundation Research Unit Study Paper No. 32, University Press of Southern Denmark.

- Hong, Seo Yeon, Pinar Karaca-Mandic, and Nicole Maestas.** 2008. "Peer Groups and Employment Outcomes: Evidence Based on Conditional Random Assignment in the U.S. Army" Mimeo, RAND.
- Hsieh, Chang-Tai, and Peter J. Klenow.** 2009. "Misallocation and Manufacturing TFP in China and India" *Quarterly Journal of Economics*, 124(4): 1403–1448.
- 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*.
- Jones, Charles.** 2011. "Intermediate Goods and Weak Links in the Theory of Economic Development" *American Economic Journal: Macroeconomics*.
- Kandel, Eugene, and Edward Lazear.** 1992. "Peer Pressure and Partnerships" *Journal of Political Economy*, 100: 801–817.
- Kenya Flower Council.** 2011. "<http://www.kenyaflowercouncil.org/floricultureinkenya.php>"
- Krueger, A., and Alex Mas.** 2004. "Strikes, Scabs and Tread Separations: Labor Strife and the Production of Defective Bridgestone/Firestone Tires" *Journal of Political Economy*.
- Ksoll, C., Rocco Macchiavello, and Ameet Morjaria.** 2010. "The Effect of Ethnic Violence on an Export-Oriented Industry" *Mimeo, Warwick University*.
- La Ferrara, Eliana.** 2002. "Self-Help Groups and Income Generation in the Informal Settlements of Nairobi" *Journal of African Economies*, 11(1): 61–89.
- La Ferrara, Eliana.** 2003. "Kin Groups and Reciprocity: A Model of Credit Transactions in Ghana" *American Economic Review*, 93(5): 1730–1751.
- Lang, K.** 1986. "A Language Theory of Discrimination" *Quarterly Journal of Economics*, 101: 363–382.
- Lazear, Edward.** 1999. "Globalization and the Market for Teammates" *Economic Journal*, 109: 15–40.
- Macchiavello, Rocco, and Ameet Morjaria.** 2013. "The Value of Relationships: Evidence from a Supply Shock to Kenya Rose Exports" *American Economic Review*, forthcoming.
- Mas, Alexandre, and Enrico Moretti.** 2009. "Peers at Work" *American Economic Review*, 99(1): 112–145.
- Miguel, Edward.** 2004. "Tribe or Nation? Nation-Building and Public Goods in Kenya versus Tanzania" *World Politics*, 56: 327–362.
- Ndegwa, S.** 1997. "Citizenship and Ethnicity: An Examination of Two Transition Moments in Kenyan Politics" *American Political Science Review*, 91(3): 599–616.

- Noury, Valerie.** 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*.
- Posner, Daniel L.** 2004. “The Political Salience of Cultural Difference: Why Chewas and Tumbukas Are Allies in Zambia and Adversaries in Malawi” *American Political Science Review*, 98(4): 529–545.
- Prat, Andrea.** 2002. “How Homogeneous Should a Team Be?” *European Economic Review*, 46(7): 1187–1207.
- Prendergast, Canice, and Robert H. Topel.** 1997. “Favoritism in Organizations” *Journal of Economic Perspectives*, 104: 958–978.
- Rohner, Dominic, Mathias Thoenig, and Fabrizio Zilibotti.** 2013. “War Signals: A Theory of Trade, Trust, and Conflict” *Review of Economic Studies*, forthcoming.
- Rotemberg, Julio J.** 2006. “Altruism, Cooperation and Reciprocity in the Workplace” In *Handbook of the Economics of Giving, Altruism and Reciprocity*. Vol. 2, , ed. Serge-Christophe Kolm and Jean Mercier Ythier, Chapter 21, 13711435. Elsevier Science.
- Shayo, Moses, and Asaf Zussman.** 2011. “Judicial Ingroup Bias in the Shadow of Terrorism” *Quarterly Journal of Economics*, 126(3): 1447–1484.
- Shetler, Jan Bender.** 2010. “Historical memory as a foundation for peace: Network formation and ethnic identity in North Mara, Tanzania” *Journal of Peace Research*, 47(5): 639–650.

Tables

Table I
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 (0.50)	0.61 (0.49)
Age (average age)	34.63 (5.21)	34.45 (5.20)	34.78 (5.21)
Experience (average years of tenure)	5.49 (1.48)	5.62 (1.40)	5.39 (1.54)

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 II
Testing for systematic team assignment

Characteristics listed in the following order: Tribe (Kikuyu = 1), Gender (Female = 1), Productivity (Above median = 1). Top number in cell: observed proportion. Bottom number (in parenthesis): proportion expected under random assignment.

		Processor 1								
		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
S u p p l i e r	0,0,0	0.009 (0.011)	0.013 (0.011)	0.012 (0.015)	0.017 (0.017)	0.013 (0.011)	0.010 (0.010)	0.012 (0.014)	0.014 (0.012)	0.101
	0,0,1	0.010 (0.012)	0.012 (0.012)	0.017 (0.016)	0.017 (0.018)	0.012 (0.011)	0.010 (0.010)	0.017 (0.015)	0.014 (0.013)	0.108
	0,1,0	0.018 (0.017)	0.020 (0.017)	0.023 (0.023)	0.025 (0.026)	0.017 (0.016)	0.016 (0.015)	0.021 (0.022)	0.016 (0.019)	0.156
	0,1,1	0.019 (0.020)	0.020 (0.020)	0.027 (0.027)	0.029 (0.029)	0.018 (0.019)	0.017 (0.017)	0.029 (0.026)	0.020 (0.022)	0.179
	1,0,0	0.012 (0.011)	0.009 (0.011)	0.015 (0.015)	0.017 (0.016)	0.009 (0.010)	0.008 (0.009)	0.016 (0.014)	0.012 (0.012)	0.098
	1,0,1	0.010 (0.011)	0.010 (0.011)	0.015 (0.015)	0.018 (0.016)	0.009 (0.010)	0.009 (0.009)	0.014 (0.014)	0.011 (0.012)	0.097
	1,1,0	0.017 (0.015)	0.016 (0.015)	0.020 (0.021)	0.024 (0.023)	0.014 (0.015)	0.013 (0.013)	0.018 (0.020)	0.016 (0.017)	0.138
	1,1,1	0.015 (0.013)	0.011 (0.014)	0.021 (0.019)	0.019 (0.020)	0.013 (0.013)	0.012 (0.012)	0.015 (0.018)	0.017 (0.015)	0.123
	TOTAL	0.110	0.110	0.151	0.165	0.105	0.094	0.143	0.122	

p-values: **Whole sample period** **Pre-conflict** **Conflict** **Team pay**
0.27
0.29
0.43
0.63

(the table continues below)

		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	0.011 (0.011)	0.010 (0.011)	0.018 (0.017)	0.020 (0.019)	0.014 (0.012)	0.010 (0.010)	0.015 (0.016)	0.013 (0.014)	0.110
	0,0,1	0.011 (0.011)	0.011 (0.011)	0.018 (0.017)	0.020 (0.019)	0.011 (0.012)	0.011 (0.010)	0.016 (0.016)	0.012 (0.014)	0.110
	0,1,0	0.016 (0.015)	0.016 (0.015)	0.022 (0.023)	0.024 (0.027)	0.016 (0.017)	0.015 (0.014)	0.023 (0.021)	0.019 (0.019)	0.151
	0,1,1	0.016 (0.016)	0.017 (0.016)	0.028 (0.025)	0.028 (0.029)	0.017 (0.018)	0.012 (0.015)	0.026 (0.023)	0.020 (0.021)	0.165
	1,0,0	0.011 (0.011)	0.011 (0.011)	0.014 (0.016)	0.020 (0.019)	0.012 (0.012)	0.010 (0.010)	0.015 (0.015)	0.013 (0.013)	0.105
	1,0,1	0.010 (0.009)	0.011 (0.009)	0.016 (0.015)	0.016 (0.017)	0.008 (0.010)	0.007 (0.009)	0.012 (0.013)	0.014 (0.012)	0.094
	1,1,0	0.015 (0.014)	0.012 (0.014)	0.021 (0.022)	0.027 (0.025)	0.016 (0.016)	0.015 (0.013)	0.020 (0.020)	0.017 (0.018)	0.143
	1,1,1	0.011 (0.012)	0.013 (0.012)	0.017 (0.019)	0.021 (0.022)	0.016 (0.013)	0.011 (0.011)	0.017 (0.017)	0.017 (0.015)	0.122
	TOTAL	0.100	0.100	0.155	0.176	0.110	0.091	0.142	0.126	

p-values:	Whole sample period	Pre-conflict	Conflict	Team pay
	0.77	0.63	0.56	0.17

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 III
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 IV
Output by team ethnicity configuration

	Processor output (1)	Processor output (2)	Team output (3)	Log (Processor output) (4)	Log (Processor output) (5)	Log (Team output) (6)
Constant	3741.35*** (162.79)	3741.35*** (147.96)	7482.71*** (285.83)	8.23*** (0.05)	8.23*** (0.05)	8.93*** (0.05)
Horizontally mixed			-291.88*** (7.51)			-0.05*** (0.00)
Horizontally mixed, processor of supplier's ethnicity	233.86*** (4.68)	233.86*** (5.54)		0.07*** (0.00)	0.07*** (0.00)	
Horizontally mixed, processor not of supplier's ethnicity	-525.74*** (4.58)	-525.74*** (5.44)		-0.18*** (0.00)	-0.18*** (0.00)	
Vertically mixed	-256.99*** (4.69)	-256.99*** (5.47)	-513.99*** (8.95)	-0.08*** (0.00)	-0.08*** (0.00)	-0.08*** (0.00)
N	199810	199810	99905	199810	199810	99905
Person-position FE?	YES	YES	YES	YES	YES	YES
Clustering	One-way (processor- team)	Two-way (processor and team)	One-way (team)	One-way (processor- team)	Two-way (processor and team)	One-way (team)

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The omitted category is homogeneous teams / processors in homogeneous teams. Data from 2007 is used in these OLS regressions. The outcome variables are de-seasonalized, daily output quantities. Point estimates are almost identical with no person-position fixed effects included (results available from the author).

Table V
Supplier ability effect by team ethnicity configuration

	Processor output	Processor output
	(1)	(2)
Supplier permanent productivity	0.09*** (0.00)	0.09*** (0.00)
Supplier permanent productivity×Horizontally mixed, processor of supplier’s ethnicity	−0.01 (0.00)	−0.01 (0.01)
Supplier permanent productivity×Horizontally mixed, processor not of supplier’s ethnicity	−0.02*** (0.01)	−0.02*** (0.01)
Supplier permanent productivity×Vertically mixed	−0.03*** (0.01)	−0.03*** (0.01)
Constant	2955.63*** (94.23)	2955.63*** (66.72)
Horizontally mixed, processor of supplier’s ethnicity	274.27*** (31.46)	274.27*** (32.19)
Horizontally mixed, processor not of supplier’s ethnicity	−395.84*** (31.63)	−395.84*** (34.11)
Vertically mixed	−84.86** (37.65)	−84.86** (42.81)
N	197054	197054
Processor (person-position) FE?	YES	YES
Clustering	One-way (processor-team)	Two-way (processor and team)

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The omitted category is processors in homogeneous teams. 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 2. 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 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).

Table VI
Testing for ethnic diversity effects using pair fixed effects and 3rd worker switches

	Unswitched Processor output	Team output	Log (Unswitched Processor output)	Log (Team output)
	(1)	(2)	(3)	(4)
Worker switched				
No change	0.32 (14.56)	4.30 (25.90)	-0.00 (0.00)	0.00 (0.00)
Homogeneous to horizontally mixed	Either P	-323.09*** (49.75)		-0.05*** (0.01)
Homogeneous to horizontally mixed, processor of supplier's ethnicity	Other P	203.24*** (28.86)	0.06*** (0.01)	
Homogeneous to vertically mixed	Supplier	-237.76*** (46.51)	-0.08*** (0.02)	-0.07*** (0.02)
Horizontally mixed to vertically mixed	Coethnic P	-206.99*** (51.09)		-0.03*** (0.01)
Horizontally mixed, processor of supplier's ethnicity to horizontally mixed, processor not of supplier's ethnicity	Supplier	-777.99*** (36.55)	-0.25*** (0.01)	
Horizontally mixed, processor not of supplier's ethnicity to vertically mixed	Coethnic P	248.27*** (28.65)	0.09*** (0.01)	
Constant		3172.70*** (5.54)	8.04*** (0.00)	8.75*** (0.00)
N	19656	15438	19656	15438
Pair FE for unswitched workers?	YES	YES	YES	YES
Clustering	One-way (unswitched pair)	One-way (unswitched pair)	One-way (unswitched pair)	One-way (unswitched pair)

P = Processor, S = Supplier, Coethnic = of supplier's ethnicity. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered by unswitched pair. Fixed effects of unswitched pairs are included. 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. 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 omitted category is pairs of workers "before" the third worker switch. The non-standard definition of regressors in the table was used so as to ease comparison of the results with those in table 4.

Table VII
Output by tribe-specific team ethnicity configuration

	Kikuyu-Luo			Kikuyu-Luhya			Luo-Luhya		
	Processor output	Team output	Processor output	Processor output	Team output	Processor output	Processor output	Team output	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Constant	3812.12*** (191.89)	3812.12*** (31.03)	7624.24*** (255.20)	4047.07*** (353.64)	4047.07*** (332.16)	8094.14*** (838.45)	1886.85*** (487.54)	1886.85*** (212.27)	5186.97*** (865.90)
Horizontally mixed			-298.18*** (13.19)			-241.67*** (28.35)			33.74 (34.67)
Horizontally mixed, processor of supplier's ethnicity	237.18*** (8.21)	237.18*** (8.70)		258.23*** (17.28)	258.23*** (17.66)		24.07 (24.49)	24.07 (21.64)	
Horizontally mixed, processor not of supplier's ethnicity	-535.35*** (8.21)	-535.35*** (9.05)		-499.89*** (17.09)	-499.89*** (17.40)		9.68 (24.25)	9.68 (21.14)	
Vertically mixed	-264.47*** (8.45)	-264.47*** (8.63)	-528.95*** (15.62)	-226.99*** (19.23)	-226.99*** (18.63)	-453.97*** (33.23)	0.61 (27.85)	0.61 (23.79)	1.22 (42.48)
N	72188	72188	36094	34986	34986	17493	16876	16876	8438
Person-position FE?	YES	YES	YES	YES	YES	YES	YES	YES	YES
Clustering	One-way (processor-team)	Two-way (processor and team)	One-way (team)	One-way (processor-team)	Two-way (processor and team)	One-way (team)	One-way (processor-team)	Two-way (processor and team)	One-way (team)

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The omitted category is homogeneous teams / processor in homogeneous teams. Data from 2007 is used in these OLS regressions. The outcome variables are de-seasonalized, daily output quantities. Point estimates are almost identical with no person-position fixed effects included (result available from the author). In this paper, Luo and Luhya workers are categorized as belonging to the Luo tribal bloc and Kikuyu workers to the Kikuyu bloc.

Table VIII
Output by team ethnicity configuration before and after conflict

	Processor output	Processor output	Team output	Log (Processor output)	Log (Processor output)	Log (Team output)
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	3554.46*** (126.10)	3554.46*** (106.49)	6899.33*** (348.17)	8.16*** (0.04)	8.16*** (0.04)	8.82*** (0.06)
Horizontally mixed			-291.31*** (7.49)			-0.05*** (0.00)
Horizontally mixed, processor of supplier's ethnicity	234.15*** (4.65)	234.15*** (5.51)		0.07*** (0.00)	0.07*** (0.00)	
Horizontally mixed, processor not of supplier's ethnicity	-525.45*** (4.56)	-525.45*** (5.33)		-0.18*** (0.00)	-0.18*** (0.00)	
Vertically mixed	-257.94*** (4.64)	-257.94*** (5.47)	-515.87*** (8.92)	-0.08*** (0.00)	-0.08*** (0.00)	-0.08*** (0.00)
Conflict	8.08 (9.13)	8.08 (9.35)	16.16 (17.87)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
Horizontally mixed \times Conflict			-260.17*** (22.54)			-0.04*** (0.00)
Horizontally mixed, processor of supplier's ethnicity \times Conflict	69.78*** (13.06)	69.78*** (13.75)		0.02*** (0.00)	0.02*** (0.00)	
Horizontally mixed, processor not of supplier's ethnicity \times Conflict	-329.95*** (13.68)	-329.95*** (14.89)		-0.13*** (0.00)	-0.13*** (0.01)	
Vertically mixed \times Conflict	-211.21*** (13.20)	-211.21*** (14.06)	-422.41*** (25.84)	-0.07*** (0.00)	-0.07*** (0.00)	-0.07*** (0.00)
N	224730	224730	112365	224730	224730	112365
Person-position FE?	YES	YES	YES	YES	YES	YES
Clustering	One-way (processor- team)	Two-way (processor and team)	One-way (team)	One-way (processor- team)	Two-way (processor and team)	One-way (team)

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The omitted category is homogeneous teams / processor in homogeneous teams. 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. Point estimates are almost identical with no person-position fixed effects included (results available from the author).

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

	Processor output (1)	Processor output (2)	Team output (3)	Log (Processor output) (4)	Log (Processor output) (5)	Log (Team output) (6)
Constant	3096.97*** (137.10)	3096.97*** (143.41)	6441.81*** (69.10)	8.01*** (0.05)	8.01*** (0.05)	8.76*** (0.01)
Horizontally mixed			-562.70*** (20.97)			-0.09*** (0.00)
Horizontally mixed, processor of supplier's ethnicity	297.47*** (12.22)	297.47*** (12.25)		0.09*** (0.00)	0.09*** (0.00)	
Horizontally mixed, processor not of supplier's ethnicity	-861.90*** (12.72)	-861.90*** (12.64)		-0.32*** (0.00)	-0.32*** (0.00)	
Vertically mixed	-480.38*** (12.29)	-480.38*** (11.98)	-957.52*** (24.06)	-0.16*** (0.00)	-0.16*** (0.00)	-0.16*** (0.00)
Team pay	-40.40*** (9.07)	-40.40*** (8.99)	-80.48*** (17.80)	-0.01*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)
Horizontally mixed \times Team pay			259.30*** (22.40)			0.04*** (0.00)
Horizontally mixed, processor of supplier's ethnicity \times Team pay	-422.54*** (13.02)	-422.54*** (13.33)		-0.13*** (0.00)	-0.13*** (0.00)	
Horizontally mixed, processor not of supplier's ethnicity \times Team pay	682.44*** (13.45)	682.44*** (13.26)		0.26*** (0.00)	0.26*** (0.00)	
Vertically mixed \times Team pay	-4.31 (13.09)	-4.31 (12.78)	-9.36 (25.85)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)
N	204107	204107	102074	204107	204107	102074
Person-position FE?	YES	YES	YES	YES	YES	YES
Clustering	One-way (processor- team)	Two-way (processor and team)	One-way (team)	One-way (processor- team)	Two-way (processor and team)	One-way (team)

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The omitted category is homogeneous teams / processor in homogeneous teams. Data from 2008 is used in these OLS regressions. The outcome variables are de-seasonalized, daily output quantities. Point estimates are almost identical with no person-position fixed effects included (results available from the author).

Table X
Heterogeneity in discriminatory behavior

	Discrimination coefficient	Discrimination coefficient
	(1)	(2)
Female	-102.41*** (30.35)	-102.41*** (33.94)
Young	-145.77*** (29.77)	-145.77*** (33.29)
Percent of workdays spent in mixed teams	-41.51 (119.74)	-41.51 (133.91)
Average discrimination coefficient of non-coethnics worker was supplied by	84.10 (195.85)	84.10 (219.03)
Conflict		163.25 (293.54)
Conflict × Female		-95.74 (71.04)
Conflict × Young		124.10* (70.46)
Conflict × Percent of workdays spent in mixed teams		100.59 (280.49)
Conflict × Average discrimination coefficient of non-coethnics worker was supplied by		-106.44 (467.47)
Constant	493.01*** (126.34)	493.01*** (141.29)
N	675	880

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. A worker's discrimination coefficient is given by the difference between the mean outputs of homogeneous and ethnically mixed teams supplied by that worker. In column 1 discrimination coefficients during the second half of 2007 are used as outcomes, while in column 2 discrimination coefficients during the second half of 2007 and the conflict period are used. "Percent of workdays spent in mixed teams" and "Average discrimination coefficient of non-coethnics worker was supplied by" refer to the previous period.

Table XI
Output gains from optimal assignment to teams 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	99.67%	s3p1p3	62.17%	Homogeneous,s3p3p3	50.00%
	Horizontally mixed	0.33%	s2p2p3	37.17%	Homogeneous,s3p1p2	31.58%
			s1p3p3	0.33%	Homogeneous,s3p2p2	10.86%
			s3p2p2	0.33%	Homogeneous,s2p2p2	6.91%
					Homogeneous,s3p2p3	0.33%
Output gains relative to:				Horiz. mixed,s2p2p2	0.33%	
observed assignment	4.39%		3.60%		11.54%	
output-minimizing assignment	8.60%		8.88%		20.47%	
		Period: Conflict				
Output-maximizing assignment by:	Ethnicity		Productivity as P and S		Ethnicity and productivity as P and S	
Assignment	Homogeneous	99.67%	s3p2p2	50.00%	Homogeneous,s3p2p3	62.83%
	Horizontally mixed	0.33%	s1p3p3	33.88%	Homogeneous,s3p1p2	17.43%
			s3p1p1	16.12%	Homogeneous,s2p3p3	10.86%
					Homogeneous,s3p3p3	7.57%
Output gains relative to:				Homogeneous,s3p1p1	0.99%	
observed assignment	8.16%		2.74%		15.09%	
output-minimizing assignment	16.96%		10.43%		29.91%	
		Period: Team pay				
Optimal assignment by:	Ethnicity		Productivity as P and S		Ethnicity and productivity as P and S	
Assignment	Homogeneous	99.67%	s3p1p3	62.17%	Homogeneous,s3p2p3	36.51%
	Horizontally mixed	0.33%	s2p2p3	37.83%	Homogeneous,s3p1p2	31.25%
					Homogeneous,s3p3p3	24.67%
					Homogeneous,s2p3p3	7.24%
Output gains relative to:				Horiz. mixed,s3p1p2	0.33%	
observed assignment	6.34%		2.70%		11.76%	
output-minimizing assignment	17.22%		10.17%		26.49%	

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. Only team compositions observed in the data (i.e., for which average output can be computed) are allowed. The procedure is described in the empirical appendix.

Figures

Figure I.a
Organization of team production

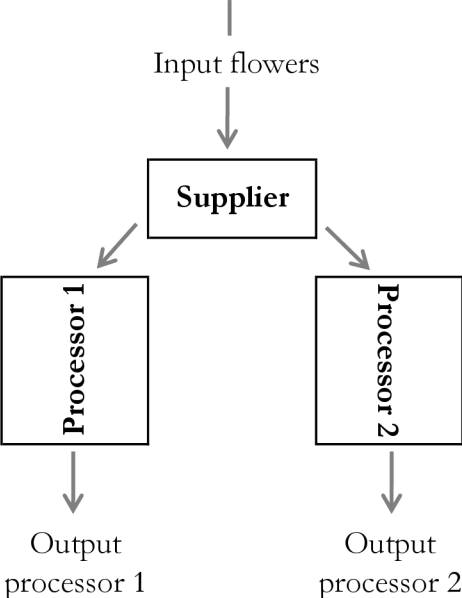


Figure I.b
Team ethnicity configuration categories

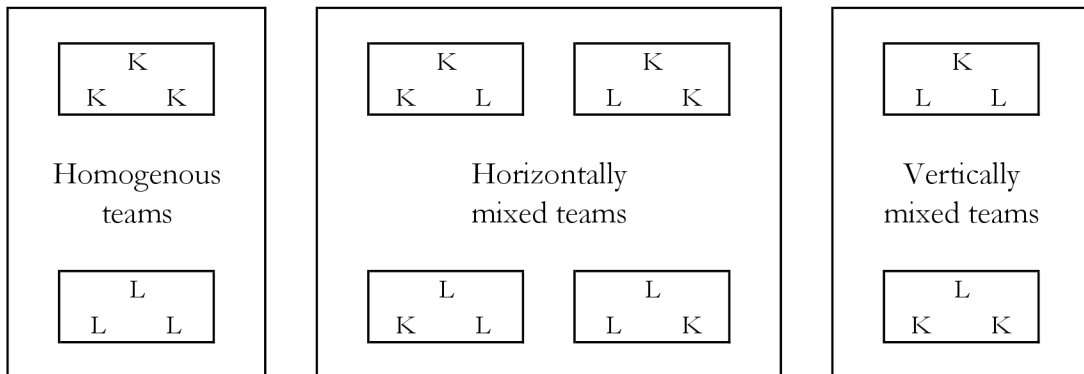


Figure II
Output in homogeneous and mixed teams across time

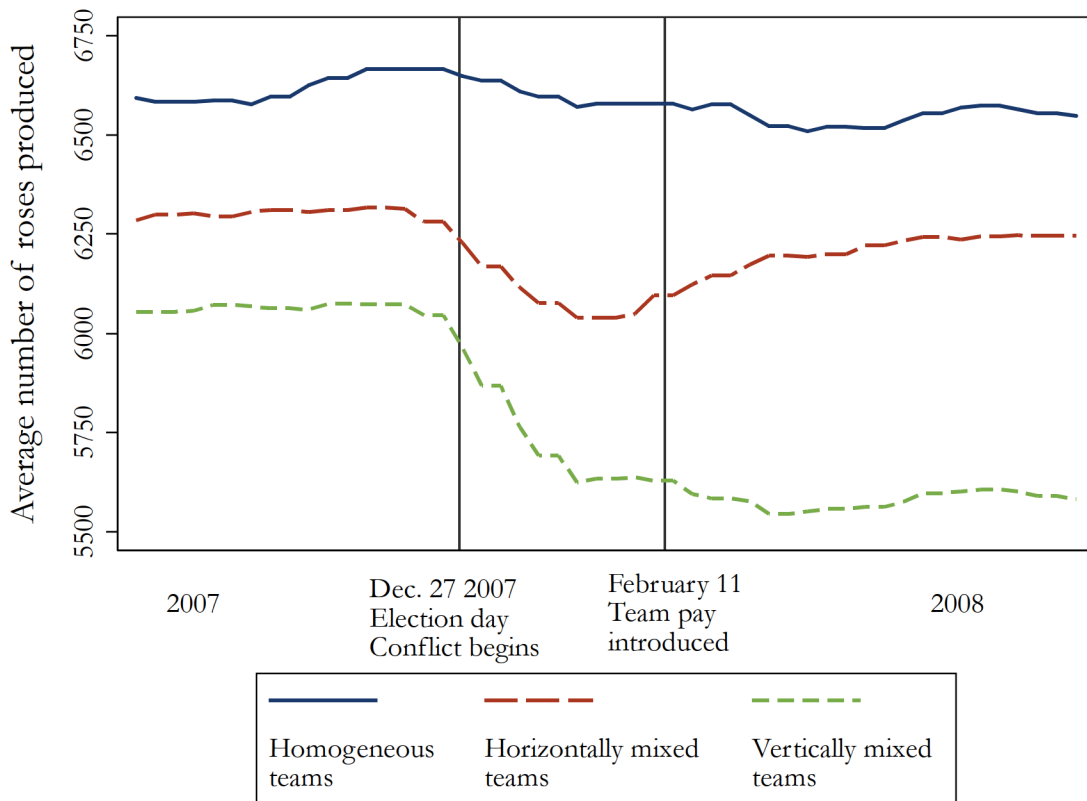


Figure III
Model predictions

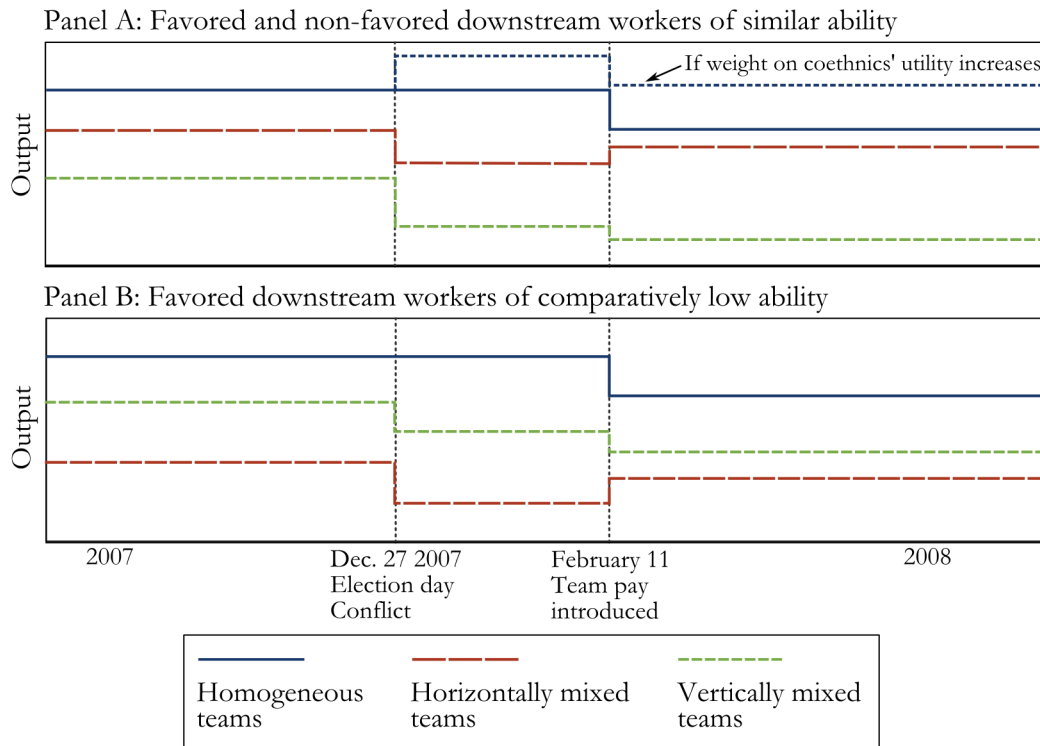
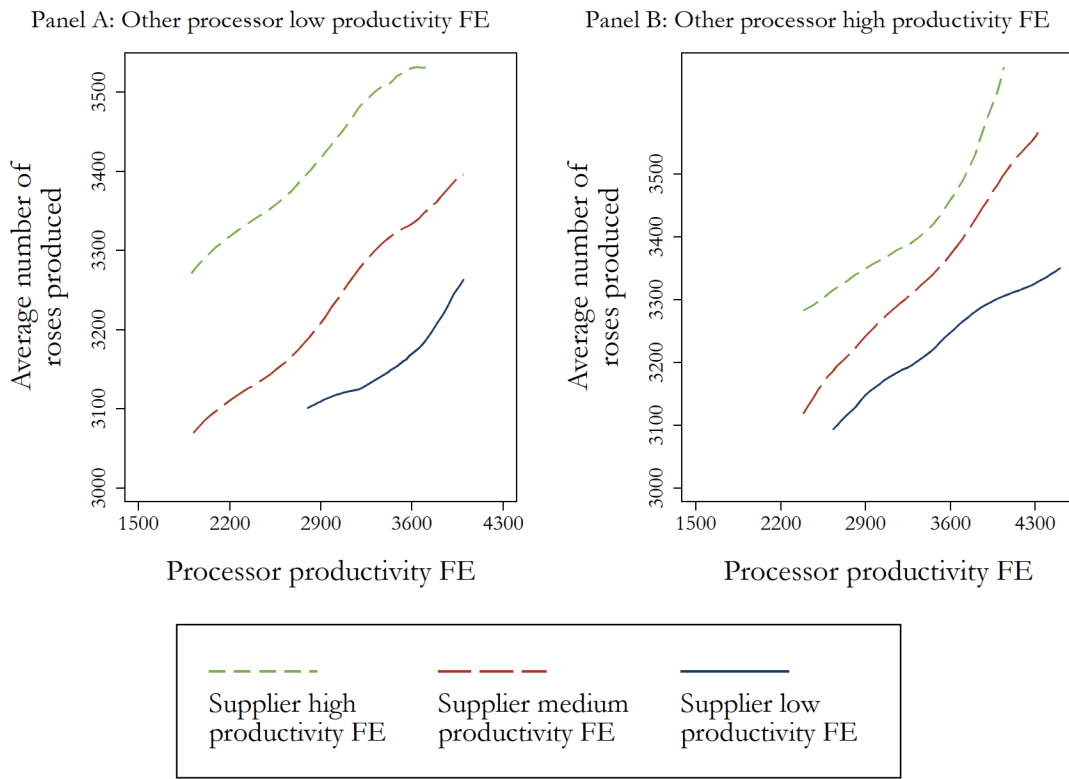
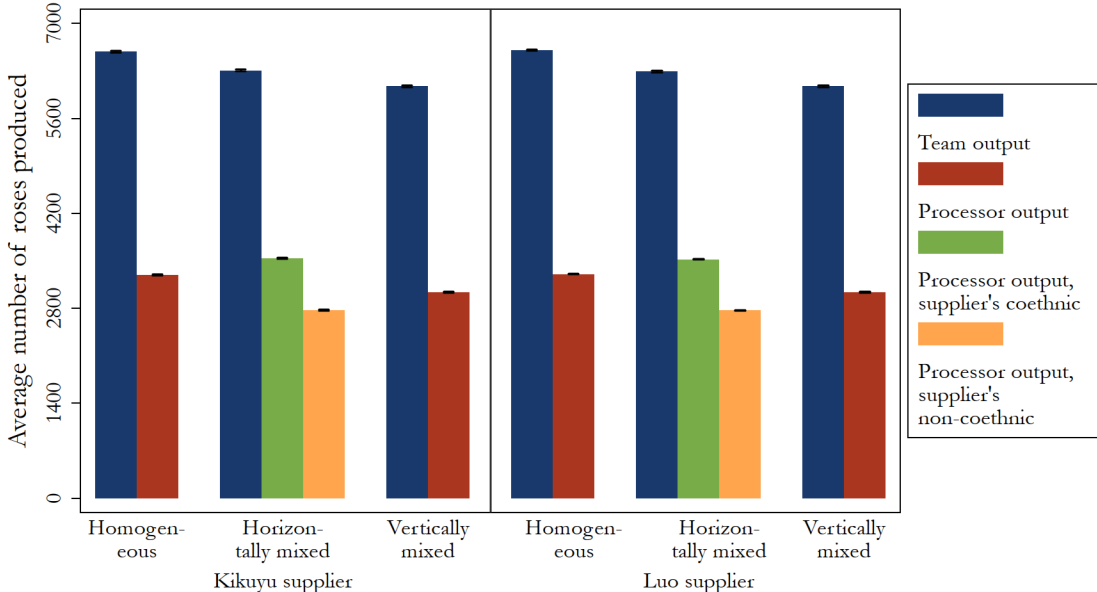


Figure IV
Investigating the shape of the production function



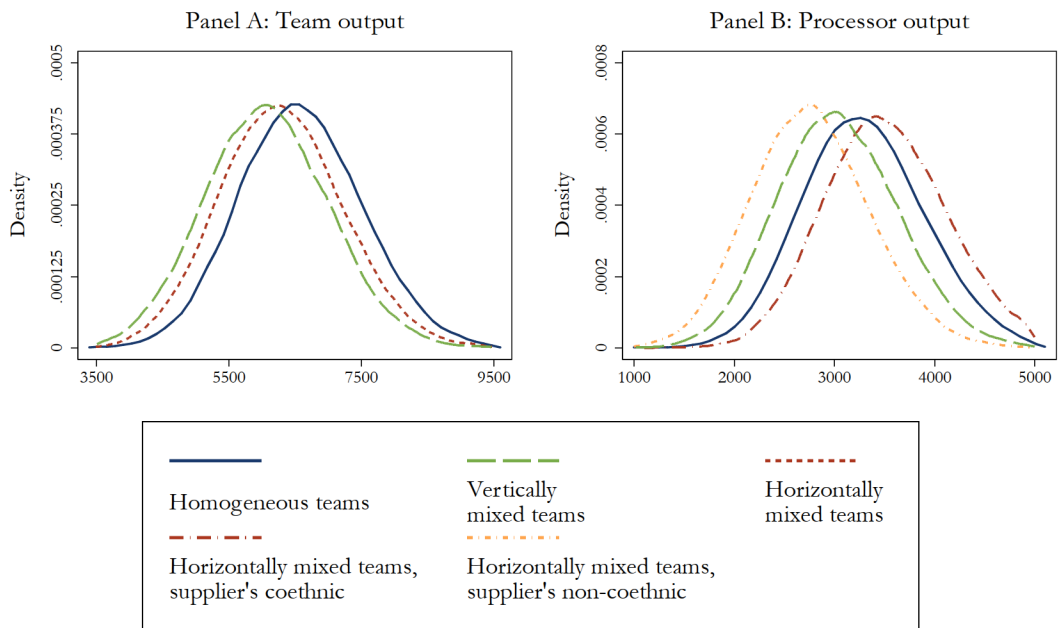
Data from 2007. Outliers (top and bottom percentile) excluded. Local polynomial plots, bandwidth = 350. The processor productivity FE is normalized to have the mean and standard deviation of processor output, and the supplier productivity FE the mean and standard deviation of team output.

Figure V
Output by team ethnicity configuration



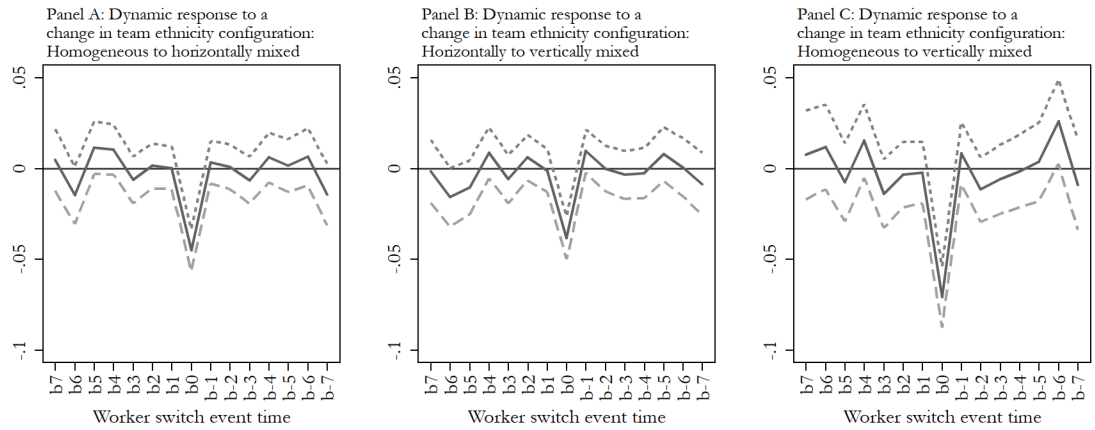
95% confidence intervals are shown but narrow. No significant differences between the quantities in the left side of the figure and the corresponding quantities in the right side of the figure (e.g. comparing average output in homogeneous teams led by Kikuyu vs Luo suppliers, the average output of a Kikuyu processor working with a Kikuyu supplier and a Luo other processor to that of a Luo processor working with a Luo supplier and a Kikuyu other processor, etc)

Figure VI
 Distribution of output by team ethnicity configuration



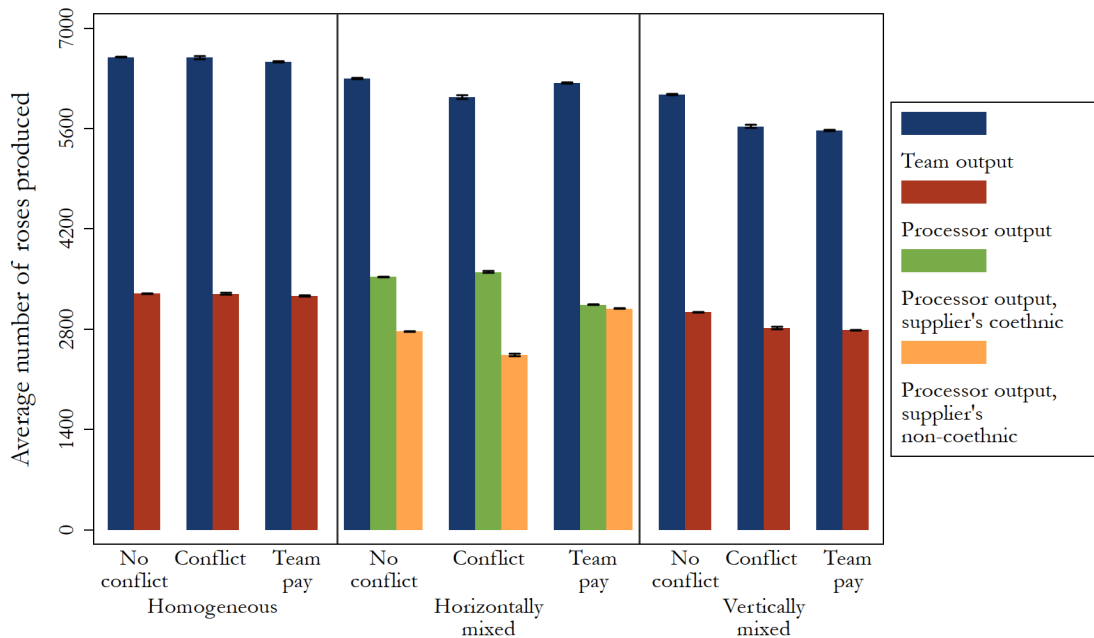
Data from 2007. Local polynomial plots, bandwidth = 100.

Figure VII
 Team output responses to changes in team ethnicity configuration



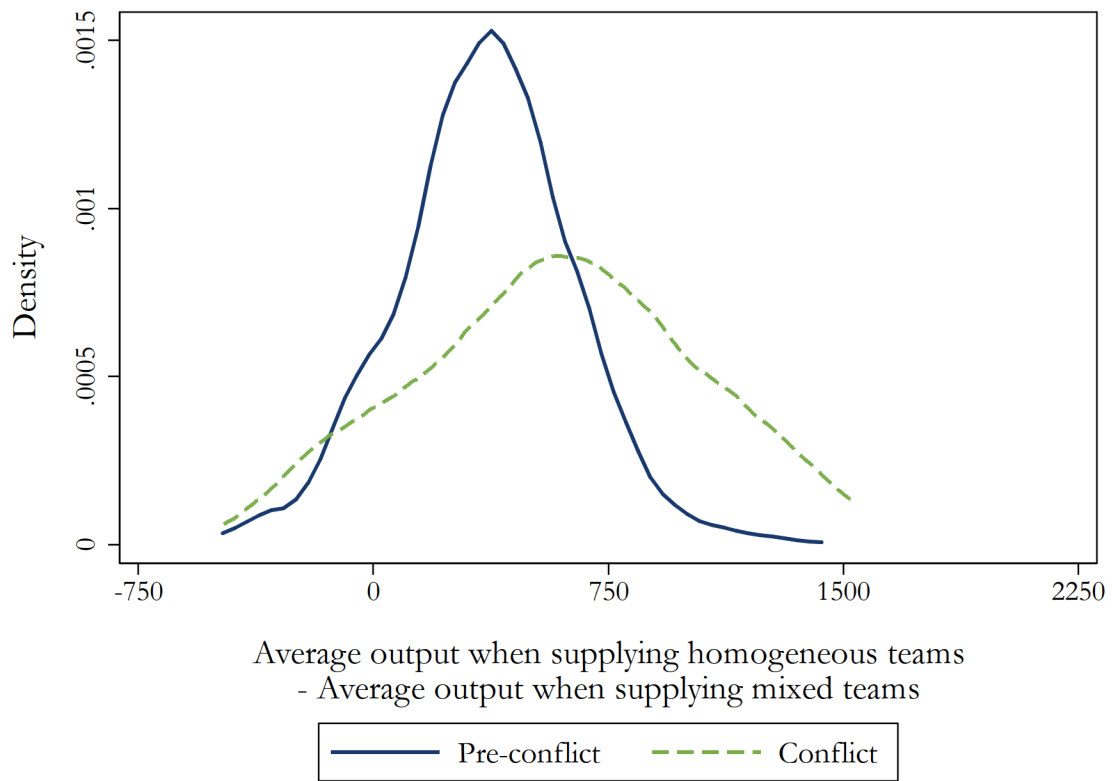
Data from 2007. The estimated coefficients from a regression of the first difference (across days) in output on an indicator for a worker switch entailing a change in team ethnicity configuration and its lead and lag terms (the other two workers on the team are unchanged) are plotted. For example, b_1 is the coefficient on the 7th lead term. The dotted lines represent 95% confidence intervals.

Figure VIII
 Output by team ethnicity configuration before and after conflict, and under team pay



95% confidence intervals are shown but narrow. "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 IX
Heterogeneity in distortionary discrimination



An observation is the output differential of a given supplier across homogeneous and mixed teams. Outliers (top and bottom percentile) excluded.

Appendix Figures

Figure A.1
Distribution of co-worker characteristics for Kikuyu and Luo suppliers

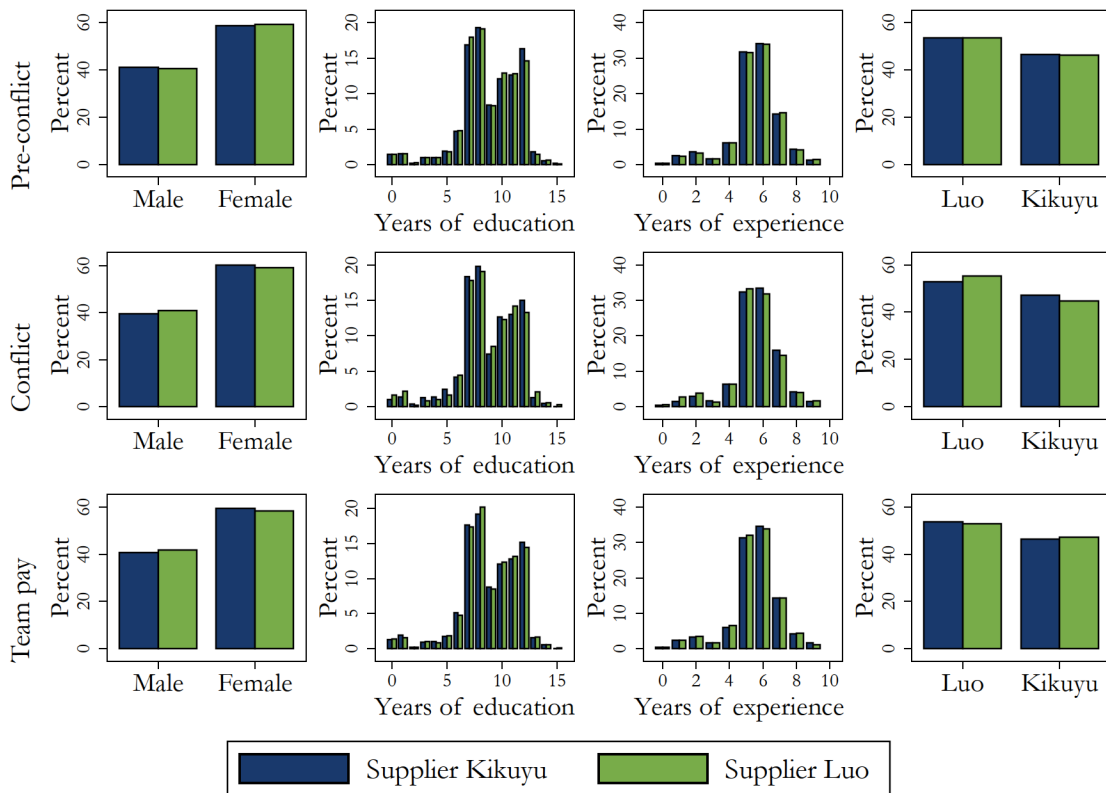


Figure A.2
 Distribution of worker characteristics in teams of different ethnicity configurations

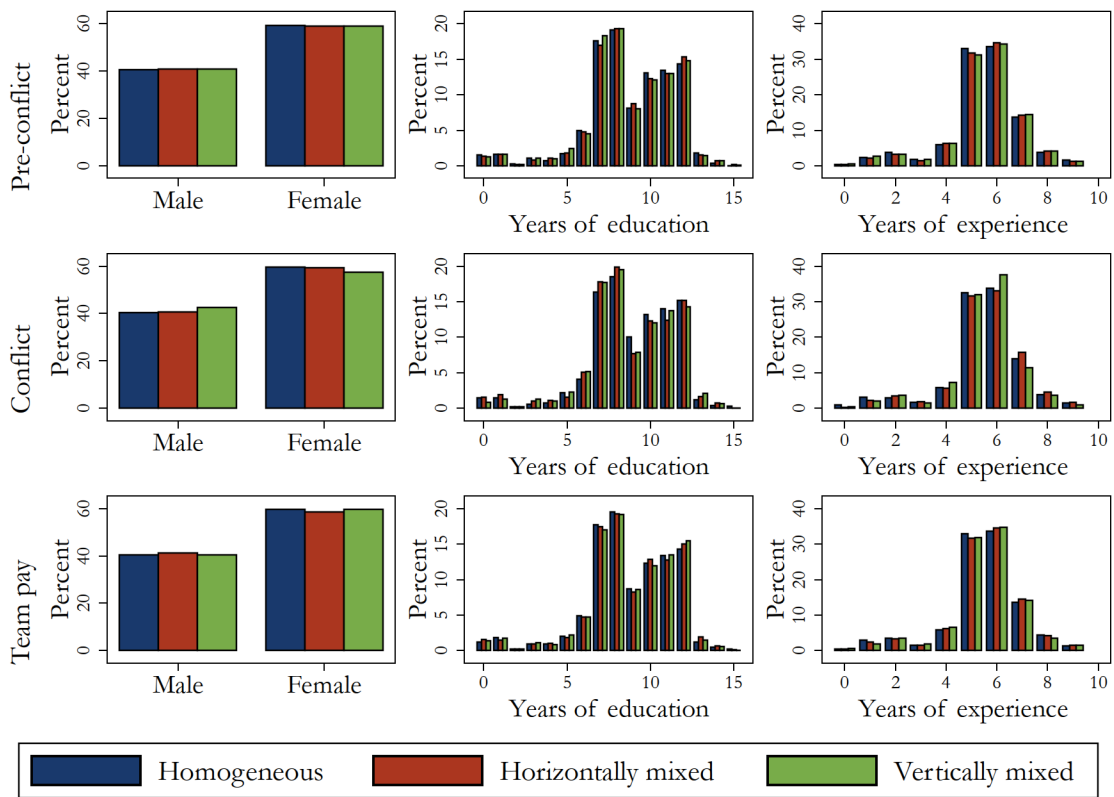



Figure A.3

Bounding the magnitude of the increase in taste for discrimination during conflict



$\lambda(\text{pre-conflict}) = 0.27, \lambda(\text{conflict}) = 0.40$