

ETHNIC DIVISIONS AND PRODUCTION IN FIRMS*

JONAS HJORT

A body of literature suggests that ethnic heterogeneity limits economic growth. This article 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: (i) ethnically homogeneous teams, and teams in which (ii) one or (iii) 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 (i) and (iii) teams—as predicted by standard incentive models—but an increase in output in (ii) 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 interethnic 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* Codes: D03, D22, D24, D61, D64, D74, F63, J24, J33, J71, 012, 014.

*Winner of the 2013 CESifo Prize for best paper in Applied Microeconomics.

I am grateful to Lawrence Katz (editor), Andrei Shleifer (co-editor), and three anonymous referees for insightful comments that significantly improved the article. I also thank Ted Miguel, David Card, Stefano DellaVigna and Fred Finan for invaluable guidance and support, and 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, LSE/UCL, University of Michigan, Oslo University, Penn State, Stanford GSB, Stockholm IIES, UPF, 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 and Morgan Gibson provided excellent research assistance. I thank CEGA, IBER, and the Russell Sage Foundation for financial support. All errors are my own.

© The Author(s) 2014. Published by Oxford University Press, on behalf of President and Fellows of Harvard College. All rights reserved. For Permissions, please email: journals.permissions@oup.com

The Quarterly Journal of Economics (2014), 1899–1946. doi:10.1093/qje/qju028.

Advance Access publication on October 7, 2014.

I. 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 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 article 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. 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 who package flowers at a plant in Kenya. 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 antagonistic ethnic blocs—the Kikuyu (and allied tribes) and the Luo (and allied tribes).

Production takes place in triangular packing units. One upstream “supplier” prepares roses that are passed on to two downstream “processors” who assemble the flowers into bunches, as illustrated in Figure I. 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 were paid a piece rate based on total team output. Low supply of roses to downstream workers of the rival ethnic group thus implied lower own pay for suppliers.




FIGURE I
Organization of Team Production

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 article'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 isolate variation in teams' ethnic composition controlling for their worker productivity composition.

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. 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 (i) the source of

productivity effects of ethnic diversity in the context of plant production in Kenya; (ii) how the economic costs of ethnic diversity vary with the degree of conflict between groups; and (iii) how workplace behavior changed as a consequence of the policies implemented by managers in response to ethnic diversity distortions at the plant.

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.¹ 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, by shifting flowers from non-coethnic to coethnic downstream workers.² 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 under team pay, suppliers can no longer influence the relative pay of the two processors through relative supply.

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. In the first main result, I find that vertically mixed teams were 8% less productive and horizontally mixed teams 5% 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 to lower

1. Unless 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 use *upstream worker* and *supplier* synonymously, and *downstream worker* and *processor* synonymously.

2. Note that both vertical and horizontal discrimination refer to behavior by the upstream worker.

the pay of non-coethnic co-workers.³ Because Kikuyu and Luo workers are of similar productivity on average, horizontal misallocation has little effect 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% more than processors of the other ethnic group.

In the second main result, 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 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, 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. The increase was due to a reduction in horizontal misallocation: a 30% 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 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 II illustrates the evolution of output in teams of different ethnicity configurations across the three sample periods observed.

This article's findings have important implications for theory and policy. 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. Distortionary, taste-based discrimination in production appears to be the primary

3. This article 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.




FIGURE II
Output in Homogeneous and Mixed Teams across Time

explanation behind my results. Discriminatory preferences should lead to 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 findings of this article also suggest that relatively brief episodes of conflict can have a long-lasting impact on distortionary attitudes toward individuals of other groups. I find no reversion in the output gap in ethnically mixed teams in the nine months after conflict ended and particularly large effects of conflict on the workplace behavior of young workers working with non-coethnics. 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, output in vertically mixed teams was 16% lower than in homogeneous teams after team pay was introduced. I discuss the sample plant's chosen response to ethnic diversity distortions and compare the effect on output to that expected from optimally assigning workers to teams.

This article contributes to and ties together several areas of research. To my knowledge, the results herein are 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 are rarely available.⁴ 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 discrimination initiated by Becker (1957). The natural experiments analyzed help me distinguish taste-based discrimination from other forms of discrimination that would have different efficiency implications, such as statistical discrimination (Phelps 1972; Arrow 1973).

I add to the recent literature on social preferences at work (Bandiera, Barankay, and Rasul 2005, 2009; Mas and Moretti 2009) and show that in the Kenyan context, upstream workers discriminate against out-group downstream workers also when doing so is costly to themselves.⁵ Burgess et al. (2013) and La Ferrara (2002) show that Africans belonging to a different ethnic group than "upstream" decision makers have less access to economic resources in other contexts,⁶ suggesting that distortionary discrimination may be a common phenomenon in Africa.

4. Fisman, 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 (see e.g. Lazear 1999; Prat 2002; Hamilton, Nickerson, and Owan 2012). See Alesina and La Ferrara (2005) for a survey of the literature.

5. Bandiera, Barankay, and Rasul (2009) find that, at a fruit farm in the United Kingdom, upstream supervisors, in their allocation of own effort and assignments, discriminate against downstream workers to whom they are not socially connected only when doing so is costless to the supervisor. Note, however, that Kranton et al. (2013) find that Americans in the lab are on average willing to lower their own income by 4.8x% to lower the income of out-group members by x% even across minimal (arbitrarily defined) groups (see also Hedegaard and Tyrann 2011; Hong, Karaca-Mandic, and Maestas 2008).

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

If individuals have discriminatory preferences, output is likely to be lower in diverse production units in most production situations in which co-workers affect each other's income. I begin to address how the productivity effects of ethnic diversity are likely to vary across time and space by studying how workplace behavior responds to increased ethnic conflict in society. Several recent papers have analyzed how the extent of trade varies with relations between countries (Michaels and Zhi 2010; Fisman, Hamao, and Wang 2013). 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. Though I focus on the effect of conflict on individual behavior, I follow Ksoll, Macchiavello, and Morjaria (2010) and Macchiavello and Morjaria (2014) in studying Kenyan flower farms during the political crisis of 2008. The authors find that flower farms' export volumes dropped by 38% on average during the conflict.⁷ By analyzing how discrimination varies with relations between groups, and across individuals with varying degrees of past exposure to and experiences with non-coethnics, this article also adds to an emerging literature investigating how social preferences are shaped (Fershtman and Gneezy 2001; Boisjoly et al. 2006; Rotemberg 2006; Bauer, Cassar, and Chytilová 2011; Jakiela, Miguel, and te Velde 2011; Shayo and Zussman 2011; Rao 2013).

How distortions due to ethnic diversity and discriminatory worker attitudes affect firms and their organization of production is an exciting venue for future research.⁸ Prendergast and Topel (1996) 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 article is particularly related to La Ferrara (2002),

7. Note, however, that Ksoll, Macchiavello, and Morjaria (2010) find that the primary reason was worker absence, which was not a problem at the farm studied here.

8. There are interesting connections between this study'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 2010; Hsieh and Klenow 2009). Some of the distortionary policies studied by macroeconomists may exist in part as a means for politicians to skew the distribution of resources toward their own ethnic groups and thus ultimately arise from biased preferences upstream. Firms whose output suffers from internal misallocation due to ethnic diversity distortions may survive due to macro-level misallocation of capital.

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.

The article is organized as follows. In Section II, 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 III, and its predictions for the three sample periods observed tested in Section IV. Section V explores the farm's response to ethnic diversity distortions. Section VI concludes. All appendix material can be found in the Online Appendix.

II. THE SETTING

II.A. 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 the 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 the postindependence era. The other major tribes have therefore typically defined themselves politically in opposition to the Kikuyu. In recent elections the opposition has been led by another big tribe, the Luo. Although political alliances have varied over time, other tribes have typically aligned themselves with one of the two associated camps. For example, in the 2007 presidential election, exit polls suggested that 94% of Kikuyus and 88% of Merus voted for the Kikuyu incumbent, whereas 98% of Luos and 75% 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).⁹

9. I 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. Although the same categorization is used by other researchers, there are two tribes for which it is not entirely clear-cut: the small Kisii tribe, whose vote was split in the 2007 election, and the Kalenjin tribe, which could be seen as a third bloc but sided with the Luo challenger in the 2007 election. Kisii and Kalenjins make up 0.5% and 4% of my sample, respectively; excluding them does not affect my results. As seen in Section IV, focusing the analysis on individual tribes gives results reaffirming the

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% of flowers imported into Europe (Noury 2011). 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 on which I focus, 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.

II.B. Data

The sample farm mainly 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; 28,281 different teams are observed during the sample period. 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 seven 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 2 leave days. Note that only workers who are observed working in all three sample periods analyzed are included in the analysis.¹⁰

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

political alliance-based categorization used here and advocated by others (see, e.g., Posner 2004a).

10. As is clear from Table I, there was very little turnover at the plant in 2007 and 2008.

statistics are in Table I. Of workers, 59% are female and 46% are Kikuyu. The average worker is 35 years old and has 5 years of tenure at the factory. These figures are similar for Kikuyu and Luo workers.

II.C. 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 I. 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 $2w$.¹¹ 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 IV.

II.D. 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. The plant I study is ideal for analyzing the impact of ethnic diversity on productivity because of its team rotation system. When a worker takes leave, another worker returning from leave joins the two remaining workers. 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.

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

11. Workers were additionally paid a small fixed component.

TABLE I
SAMPLE SUMMARY STATISTICS

	Whole sample (<i>N</i> = 924)	Kikuyu (<i>N</i> = 426)	Luo (<i>N</i> = 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)
Percent of days worked, preconflict	0.90 (0.02)	0.90 (0.02)	0.90 (0.02)
Percent of days worked, conflict	0.90 (0.05)	0.90 (0.05)	0.90 (0.05)
Percent of days worked, team pay	0.90 (0.02)	0.90 (0.02)	0.90 (0.02)
Average work spell, preconflict	18.38 (1.38)	18.38 (1.42)	18.39 (1.34)
Average work spell, conflict	19.34 (2.98)	19.37 (2.98)	19.32 (2.98)
Average work spell, team pay	18.18 (1.47)	18.17 (1.45)	18.18 (1.49)

Notes. 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.

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.

With 46.10% Kikuyu and 53.90% Luo workers, 25.46% of teams should be ethnically homogeneous, 49.69% horizontally mixed, and 24.85% vertically mixed, if assignment was random. The percentages observed in the data are 25.64%, 49.61%, and 24.76% during the preconflict period; 27.38%, 48.35%, and 24.26% during the conflict period; and 25.32%, 49.26%, and 25.42% during the team pay period.¹² It is clear that workers are not assigned to or sort into teams based on ethnicity.¹³

12. The preconflict 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 IV).

13. Online Appendix 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.

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,¹⁴ 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 II.¹⁵ The matrixes 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 were used. I thus use a periodical "snapshot" of data in the table: team compositions on the first day of every month.¹⁶ 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,¹⁷ and little ability to do so given their limited knowledge of worker

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

15. Online Appendix Figure 2 displays the distribution of workers' gender, years of education, and years of experience across homogeneous, horizontally mixed, and vertically mixed teams during each of the three sample periods. The distributions are essentially identical.

16. The tests are insignificant if data from other days are used instead. Note that the table uses three binary worker characteristics 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, and the chi-squared statistics are insignificant for that pair of workers.

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

TABLE II
TESTING FOR SYSTEMATIC TEAM ASSIGNMENT

	Processor 1							Total	
	0,0,0	0,0,1	0,1,0	0,1,1	1,0,0	1,0,1	1,1,0		1,1,1
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	
<i>p</i> -values:	Whole sample period		Preconflict		Conflict		Team pay		
	.27		.29		.43		.63		

TABLE II
(CONTINUED)

		Processor 2								Total
		0,0,0	0,0,1	0,1,0	0,1,1	1,0,0	1,0,1	1,1,0	1,1,1	
	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
P	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
r	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
o	0,1,1	0.016 (0.016)	0.017 (0.016)	0.028 (0.028)	0.028 (0.029)	0.017 (0.018)	0.012 (0.015)	0.026 (0.023)	0.020 (0.021)	0.165
e	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
s	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
o	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
I	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				Preconflict		Conflict	Team pay	
		.77				.63		.56	.17	

Notes. 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 parentheses): proportion expected under random assignment. 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 2nd characteristics listed in row i and the worker in position j has the 2nd characteristics listed in column j . The bottom number is the expected proportion under the null hypothesis of independence. χ^2 -values for Pearson's chi-squared 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-squared tests would be violated if all data were 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-squared 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 $t - 2$.

characteristics and the plant's leave and rotation system.¹⁸ 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 low proportion of flower farms' total costs (EDRI 2008).

In the analysis that follows 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.

III. DISCRIMINATION: THEORETICAL FRAMEWORK

In this section, I present a simple framework in which output gaps that arise in ethnically diverse teams do so because suppliers behave as if they attach a lower weight to the utility of non-coethnic than coethnic processors. The starting framework is reduced form in the sense that several underlying mechanisms could drive such behavior. In deriving further predictions that depend on the particular mechanism at play, I interpret the differential weight attached to coethnics' well-being as arising from suppliers' (discriminatory) social preferences, for two reasons. First, taste-based misallocation is a particularly important possibility to consider due to the implied distortions in the aggregate economy (Becker 1957). Second, I 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 framework's predictions in the next section.

18. Team 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 turnover).

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 their effort simultaneously.¹⁹

Finally, assume that the supplier attaches weight θ_p to the utility of processor p .²⁰ Let $\theta_i = \theta_C$ if processor i is of the supplier's ethnic group and $\theta_i = \theta_{NC}$ if not. (I do not distinguish between the two specific ethnic groups here.²¹) Suppliers with a different weight for coethnics and non-coethnics have discriminatory preferences. I focus on discriminatory behavior on the part of suppliers because the effect of suppliers' behavior on processors' pay is presumably more salient than that of processors' effort on the

19. In reality, supply and processing decisions take place continuously throughout the workday, but when the (data) time-unit to which the model must be compared is a whole workday, 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). Although simultaneous moves are assumed here for simplicity, and because any codependence between a supplier's and a processor's effort level likely runs both ways, a Stackelberg version of the model in which the supplier moves first, taking the processors' expected response to her effort into account, gives very similar predictions. (Propositions 2–5, 6.iii, and 6.v in the Online Appendix are unchanged in the Stackelberg scenario, while 6.ii and 6.iv differ in ways that are noted below. It turns out that the simultaneous moves version of the model in fact describes workers' behavior better than the Stackelberg version.)

20. This formulation follows Becker (1974), Charness and Rabin (2002), Chen and Li (2009), and others. θ_p can be either positive or negative, but even if positive, that is, 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 to lower the income of processor 1 relative to processor 2 if $0 < \theta_1 < \theta_2$. If we abstract from the supplier's cost of effort for purposes of illustration, the analogy between the specification here and Becker's (1957) 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.

21. Homogeneous teams may, for example, be either Kikuyu-Kikuyu-Kikuyu or Luo-Luo-Luo. I 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.

supplier's pay.²² Although a model in which processors also (behaved as if they) had social preferences would be less tractable, such a model would give generally similar predictions but predict a response to the introduction of team pay that does not match the observed patterns, as discussed later.

A processor maximizes her utility of pay minus her cost of effort, $2wf(e_{sp}, \alpha_s, e_p, \alpha_p) - d(e_p)$, and the supplier her utility of pay minus her cost of effort plus the additional utility (or disutility) she derives from the well-being of processor 1 and processor 2:

$$\begin{aligned} \text{Max}_{e_{s1}, e_{s2}} & w(f(e_{s1}, \alpha_s, e_1, \alpha_1) + f(e_{s2}, \alpha_s, e_2, \alpha_2)) - d(e_{s1} + e_{s2}) \\ & + \theta_1(2wf(e_{s1}, \alpha_s, e_1, \alpha_1) - d(e_1)) + \theta_2(2wf(e_{s2}, \alpha_s, e_2, \alpha_2) - d(e_2)). \end{aligned} \quad (1)$$

A full model with output a Cobb-Douglas function of its arguments is developed in the Online Appendix, and the predictions it implies are shown and proved. Here I lay out the intuition of the framework. The model predicts that a processor's output is increasing in the weight the supplier attaches to her utility and decreasing in the weight of the other processor. If the supplier has discriminatory preferences, processor output is thus expected to be higher (i) when working with a coethnic supplier, and (ii) when working with another processor who is not of the supplier's ethnicity. Similarly, the framework predicts that biased high-ability suppliers allocate more of their additional capacity to supplying coethnic processors.

Biased suppliers are predicted to in effect discriminate both vertically, undersupplying processors of the other ethnic group, and horizontally, additionally shifting roses from non-coethnic to coethnic processors when possible. Vertical and horizontal misallocation of roses is predicted to lower total team output so that output is higher in homogeneous than in both vertically and horizontally mixed teams.²³ Although total supply will be lower in vertically mixed teams than in horizontally mixed teams, horizontal misallocation will occur only in horizontally mixed

22. While processors influence the supplier's pay only through effort, the supplier influences processors' pay also through simple, differential supply decisions.

23. Note 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.

teams. If on average flowers are shifted toward 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.²⁴ Otherwise output is expected to be lowest in vertically mixed teams.

It is possible that the period of ethnic conflict in Kenya in early 2008 led to a change in attitudes toward co-workers of the other ethnic group, which I model as a change in θ_{NC} .²⁵ 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 being greatest for non-coethnic processors in horizontally mixed teams where the relative benefits of supplying a non-coethnic processor also decrease.

Six weeks into the conflict period the plant began paying processors for their combined output (the supplier's pay system did not change). Under such a pay system, processor 1's utility from pay is $w(q_1 + q_2)$, rather than $2wq_1$, so that the supplier's problem becomes:

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

In scenarios in which the two downstream workers are of the same ethnic group—homogeneous and vertically mixed teams—the supplier's problem now 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.²⁶ Because the two processors in a team are paid the same under

24. In general the impact of horizontal misallocation on the average output of horizontally mixed teams will thus depend on (i) the ethnic make-up of the population of workers and (ii) the relative productivity of individuals of different ethnic groups.

25. It is also possible that output in homogeneous teams is affected by conflict, for example, due to disruption effects (Ksoll, Macchiavello, and Morjaria, 2010) or changes in individuals' weight on coethnics' utility (Eifert, Miguel, and Posner, 2010), but the focus here is on differences in output between teams of different ethnic compositions.

26. In a Stackelberg version of the model, output in homogeneous and VM teams could decrease or increase when team pay is introduced.

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. 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 free-riding on each other.²⁷ 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.

In the next section I interpret the results in light of the model presented here and then discuss the ability of non-taste-based mechanisms to explain the results.

IV. THE EFFECT OF ETHNIC DIVERSITY ON PRODUCTIVITY

IV.A. Productivity in Homogeneous and Ethnically Diverse Teams

To correctly interpret observed ethnic diversity effects, it is useful to first investigate the shape of the production function. In the Online Appendix I explore how average processor output varies with the productivity of the worker in each of the three positions in a team nonparametrically. 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. There is also evidence of 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.²⁸

27. Unlike the Stackelberg scenario, it is not the case in this framework that more is supplied to non-coethnic processors in HM teams under team pay. This is because assuming simultaneous moves means that the supplier does not take processors' cost of effort into account when making her supply decisions.

28. The evidence is consistent with a range of possible production functions, which I do not attempt to distinguish between. What is important for the purposes of this article is that the evidence is not consistent with some *ex ante* plausible production function "shapes" which would give different predictions for how biased social preferences are expected to influence supplier behavior, such as Leontief or a production function in which the supplier can dictate work speed regardless of the processor's desired speed. One possible reason supplier

I begin by focusing on the first year of the sample period, when processors were paid based on their own output, before conflict began. The histogram in Figure III displays mean output by team ethnicity configuration in 2007. Note first that there are no significant differences between teams with Kikuyu and Luo suppliers.²⁹ This finding enables a more concise presentation of the evidence to follow. In the remainder of the article, 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 III 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 Online Appendix Figure 4. 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.

Regression results corresponding to Figure III are presented in columns (1) and (2) of Table III.³⁰ In column 1 I regress individual output on dummies for the processor's ethnicity in relation to that of her two co-workers. From the perspective of a processor, a team has one processor p , one other-processor o and one

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.

29. In the preconflict period, the average output of all-Kikuyu teams was 6,586 and that of all-Luo teams 6,606 (p -value on the difference = .1231). The average output in HM teams with a Kikuyu supplier was 6,307 and that in HM teams with a Luo supplier 6,290 (p = .0503). The average output in VM teams with a Kikuyu supplier was 6,073 and that in HM teams with a Luo supplier 6,075 (p = .9764). Note that even if the first two differences are nearly significant, it is clear from the magnitude of the output differences (relative to that found comparing homogeneous, HM, and VM teams) that productivity (or propensity to discriminate) differences across ethnicities do not drive the results discussed in the remainder of the article.

30. Throughout the analysis, data is deseasonalized 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 $\frac{\bar{m}}{m_i}$.




FIGURE III
Output by Team Ethnicity Configuration

Data from 2007.

supplier s . A worker $w \in W$ in position $m \in \{p, o, s\}$ on date d is of tribe $t_{wd}^m \in \{K, L\}$. The output of processor $i \in W$ on date d , q_{id}^p , is then specified as

$$q_{id}^p = \alpha + \beta^{HM,C} \mathbb{1}\{t_{id}^p = t_{kd}^s, t_{id}^p \neq t_{jd}^o\} + \beta^{HM,NC} \mathbb{1}\{t_{id}^p \neq t_{kd}^s, t_{id}^p \neq t_{jd}^o\} + \beta^{VM} \mathbb{1}\{t_{id}^p \neq t_{kd}^s, t_{id}^p = t_{jd}^o\} + \delta_w^m + \theta_d + \varepsilon_{id}^p \tag{3}$$

where $j \in W$ is other-processor and $k \in W$ is supplier. $\hat{\beta}^{HM,C}$ is the resulting estimate of the additional output associated with a processor in a horizontally mixed (HM) team who is a coethnic (C) of the supplier, relative to that of a processor in a homogeneous teams (the omitted category), $\hat{\beta}^{HM,NC}$ the estimate of the additional output associated with a processor in a horizontally mixed team who is not of the supplier's ethnicity (NC), and $\hat{\beta}^{VM}$ the estimate of the additional output associated with a processor in a vertically mixed (VM) team. θ_d is a date fixed effect and δ_w^m the set of person-position fixed effects so that the effect of ethnic diversity on output is identified from variation in teams' ethnic composition, controlling for their worker productivity composition.

TABLE III
OUTPUT BY TEAM ETHNICITY CONFIGURATION

Sample	Preconflict		Preconflict, incoming/outgoing worker of same productivity tercile	
	(1) Log (processor output)	(2) Log (team output)	(3) Log (unswitched processor output)	(4) Log (team output)
Constant (H)	8.153*** (0.024)	8.846*** (0.029)	8.018*** (0.056)	8.729*** (0.065)
Horizontally mixed (HM)		-0.046*** (0.001)		
Horizontally mixed, processor of supplier's ethnicity (HM,C)	0.070*** (0.002)			
Horizontally mixed, processor not of supplier's ethnicity (HM,NC)	-0.181*** (0.002)			
Vertically mixed (VM)	-0.084*** (0.002)	-0.083*** (0.001)		
H to HM				-0.051*** (0.005)
H to HM,C			0.064*** (0.006)	
H to VM			-0.076***	-0.076***

TABLE III
(CONTINUED)

Sample	Preconflict		Preconflict, incoming/outgoing worker of same productivity tercile	
	(1) Log (processor output)	(2) Log (team output)		(3) Log (unswitched processor output)
HM to VM			(0.008)	
			-0.038***	
			(0.006)	
Coethnic P			(0.007)	
HM,C to HM,NC			-0.243***	
			(0.006)	
Supplier			0.086***	
			(0.007)	
Coethnic P				
HM,NC to VM				
N	199,810	99,905	18,504	13,168
Person-position FE?	Yes	Yes	No	No
Pair FE for unswitched workers?			Yes	Yes
Date FE?	Yes	Yes	Yes	Yes
Clustering	Two-way (processor and team)	One-way (team)	One-way (unswitched pair)	One-way (unswitched pair)

Notes. Standard errors in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$. Data from 2007 are used in these OLS regressions. The outcome variables are deseasonalized output quantities. The omitted category is homogeneous teams/processors in homogeneous teams in columns (1) and (2), and pairs of workers before a third worker switch in columns (3) and (4). The outcome variables in columns (3) and (4) are output quantities averaged within three-person teams. In column (3) the output of the processor *not being switched* is on the left-hand side. The coefficient on "no change" in ethnicity configuration is not shown. 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 a t for pair A and a $t-1$ for pair B. For switches involving no change in ethnicity configuration, the dummy is (arbitrarily) turned on at the later of the two periods. The nonstandard definition of regressors in columns (3) and (4) was used to ease comparison of the results with those in columns (1) and (2).

In column (2) I regress team output on dummies for the ethnicity configuration of the team as a whole. From the perspective of a team as a whole, a team has one processor 1, one processor 2, and one supplier s . As the processor positions are “interchangeable,” $m \in \{p, s\}$. Team T 's output on date d , Q_{Td} , is then specified as

$$(4) \quad Q_{Td} = \lambda + \gamma^{HM} \mathbb{I}\{t_{id}^1 \neq t_{jd}^2\} + \gamma^{VM} \mathbb{I}\{t_{id}^1 \neq t_{kd}^s, t_{jd}^2 \neq t_{kd}^s\} \\ + \eta_w^m + \theta_d + \epsilon_{Td},$$

where worker i is processor 1, j processor 2 and k supplier s . $\hat{\gamma}^{HM}$ is the resulting estimate of the additional output associated with an HM team, relative to that of a homogeneous team (the omitted category), and $\hat{\gamma}^{VM}$ the estimate of the additional output associated with a VM team. θ_d is a date fixed effect and η_w^m the set of person-position fixed effects.

To account for possible within-team and within-processor correlation in output, I cluster the standard errors at the team level in the team output regressions and at both the processor and team level in the individual output regressions. Unless otherwise noted, analogous specifications are used in all tables to follow.

The effects are very precisely estimated. Excluding person-position fixed effects has little influence on the results, as expected.³¹ The output of processors in VM teams is 8% lower than that of processors in homogeneous teams, an output gap that is also reflected in the total output of VM 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 III also indicate that suppliers discriminate horizontally. The output of the non-coethnic processor in an HM team is 18% lower than that of processors in homogeneous teams, and 8% lower than that of processors in VM teams. The output of the coethnic processor in an HM team is 7% higher than that of processors in homogeneous teams. The output gap between coethnic and non-coethnic processors in HM teams exceeds 1 standard deviation of processor output. As Becker (1957)

31. Note also that all results go through if the outcome variables are specified in levels instead of logs (results available on request).

predicted, favored workers benefit from discrimination against nonfavored workers.

Recall that the output loss from horizontal discrimination will depend on the relative productivity of favored and nonfavored downstream workers. In the context of the farm, the two ethnic groups are similarly sized, and we saw already that Kikuyu and Luo workers are of similar ability on average. In such a situation, the output of VM teams is expected to be lower than that of HM teams, which is what we see in Table III. Although VM teams are in aggregate 3% less productive than HM teams, the lowest output processors are found in HM teams. The distribution of output across downstream workers is significantly affected by horizontal discrimination.

That the absolute magnitude of the coefficients on the indicators for the coethnic and non-coethnic processor in HM teams are significantly different from each other highlights that the low output of non-coethnic processors in HM teams is due not just to horizontal discrimination. For purposes of illustration, suppose that in the absence of misallocation of roses across the two processors in a team, the output of a coethnic processor in an HM 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 an HM team would be equal to that of a processor in a VM team. We can then decompose the output gap between homogeneous and HM teams: 12% would be due to the effect of horizontal misallocation and 88% due to vertical misallocation.³² Although 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 Online Appendix Table I, I regress processor output on a dummy for the supplier being above median in average output (when

32. This 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.

working as a supplier), interacted with team ethnicity configuration dummies. The results show that higher supplier productivity benefits non-coethnic processors less than coethnic processors in homogeneous teams, as predicted by the model.³³

In light of the model, the results 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. Because such behavior also lowers the pay of the supplier, the results are consistent with a willingness to pay to discriminate on the part of upstream workers. Although I am not aware of other studies of individuals' willingness to pay to lower the incomes of out-group members in natural settings, the estimated magnitudes are comparable to those found in the lab. For example, Kranton et al. (2013) find that across "minimal" (arbitrarily defined) groups in the United States, subjects are on average willing to lower their own income by 4.8x% to lower the income of out-group members by x%.³⁴

In Section IV.C, I consider alternative theories that predict negative ethnic diversity effects, but I 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 columns (1) and (2) of Table III. 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 columns (3) and (4) of Table III 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

33. The estimates suggest that relative to processors in homogeneous teams, processors in VM teams benefit significantly less from higher supplier productivity. The differential effect for non-coethnic processors in HM teams is negative but insignificant. The greater positive effect of higher supplier productivity for coethnic processors in HM teams relative to those in homogeneous teams predicted by the model does not find support in the results.

34. The authors find that 22% of subject pay to lower the incomes of out-group members, and those who do so are on average willing to pay 22x% to lower the income of out-group members by x%. (This interpretation is subject to the usual caveats of lab studies, such as narrow bracketing.)

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.³⁵ There is no significant change in output when the outgoing and incoming worker are of the same ethnic group. The estimates for team output show output falling by 5% when a team goes from being homogeneous to HM due to a worker switch, by 8% when a team goes from being homogeneous to VM, and by 4% when a team goes from being HM to VM. In column (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. The estimates for individual output are also very similar to those found in column (4). 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 III represent the causal effect of ethnic diversity.

Figure IV depicts the temporal response of team output to the event of a worker substitution leading to a change in a team's ethnicity configuration. 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 gender—show minor output differences and no sign of systematic, differential behavior when working with in-group versus out-group co-workers (results available on request).

Recall that this article 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, for example, the Luhya tribe being categorized as belonging to the Luo ethnic group. Of the sample, 86% belongs to three tribes: the Kikuyu (41%), Luo (30%), and Luhya (15%). I now consider subsamples of teams in which

35. The 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.

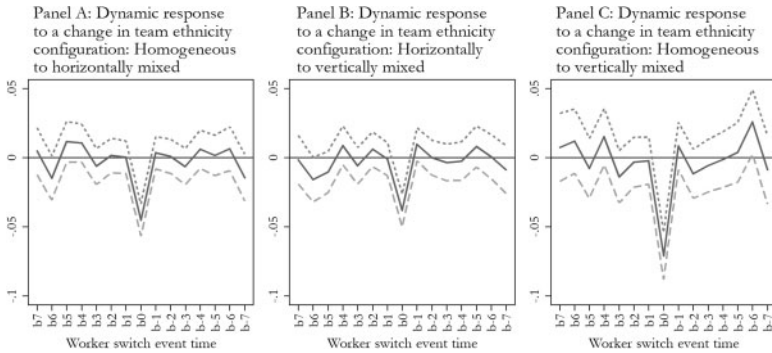


FIGURE IV

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 seventh lead term. The dotted lines represent 95% confidence intervals.

workers belong to two specific tribes, focusing on the Kikuyu–Luo, Kikuyu–Luhya, and Luo–Luhya subsamples. The estimates in Table IV show that in a subsample of teams consisting of workers from two different tribes categorized as belonging to the same tribal bloc, output differences across teams of different ethnic configurations are minor. The output of VM teams, for example, is not significantly different from that of homogeneous teams in the Luo–Luhya subsample. But within two different subsamples of teams consisting of workers of two specific tribes categorized as belonging to different tribal blocs, output differences across homogeneous and mixed teams are pervasive and of an extent similar to that seen in the full-sample analysis in Table III.³⁶ Interestingly, these results highlight that the relevant ethnic relations in Kenya are social or political constructs rather than primordially defined—consistent with the response to intensified political conflict found in the next subsection. While the Luo are Nilotic, both the Luhya and Kikuyu are Bantu, but the Luo and

36. There are only minor differences across the Kikuyu–Luo and the Kikuyu–Luhya subsamples, analyzed in columns (1)–(2) and (3)–(4) of Table IV, respectively.

TABLE IV
OUTPUT BY TRIBE-SPECIFIC TEAM ETHNICITY CONFIGURATION

	(1) Kikuyu-Luo	(2) Kikuyu-Luo	(3) Kikuyu-Luhya	(4) Kikuyu-Luhya	(5) Luo-Luhya	(6) Luo-Luhya
Sample	Log (processor output)	Log (team output)	Log (processor output)	Log (team output)	Log (processor output)	Log (team output)
Constant	8.109*** (0.085)	8.807*** (0.085) -0.048*** (0.002)	8.138*** (0.051)	8.837*** (0.061) -0.041*** (0.005)	8.033*** (0.115)	8.727*** (0.111) -0.000 (0.006)
Horizontally mixed						
Horizontally mixed, processor of supplier's ethnicity	0.072*** (0.003)		0.076*** (0.006)		-0.000 (0.007)	
Horizontally mixed, processor not of supplier's ethnicity	-0.185*** (0.003)		-0.175*** (0.006)		-0.003 (0.007)	
Vertically mixed	-0.087*** (0.003)	-0.086*** (0.003)	-0.077*** (0.006)	-0.075*** (0.005)	-0.007 (0.009)	-0.006 (0.007)
N	72,188	36,094	34,986	17,493	16,876	8,438
Person-position FE?	Yes	Yes	Yes	Yes	Yes	Yes
Date FE?	Yes	Yes	Yes	Yes	Yes	Yes
Clustering	Two-way (processor and team)	One-way (team)	Two-way (processor and team)	One-way (team)	Two-way (processor and team)	One-way (team)

Notes: Standard errors in parentheses. * $p < .05$, ** $p < .01$, *** $p < .001$. The omitted category is homogeneous teams/processor in homogeneous teams. Data from 2007 are used in these OLS regressions. The outcome variables are deseasonalized, daily output quantities. In this paper, Luo and Luhya workers are categorized as belonging to the Luo tribal bloc and Kikuyu workers to the Kikuyu bloc.

Luhya were political allies in rivalry with the Kikuyu in 2007 (see also Posner 2004a,b; 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 on private sector output to vary with factors that influence taste for discrimination, such as the political climate and relations between groups. In the next subsection, I analyze differences in output between homogeneous and mixed teams during the period of ethnically based political conflict in Kenya in early 2008.

IV.B. *Ethnic Conflict and the Impact of Diversity on Productivity*

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 counterattacks 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, 2008, a peace agreement was reached, although violence continued in many areas, and it was not until after April 3 that the political crisis ebbed when the two sides reached an agreement on the composition of a power-sharing government. The conflict period significantly disrupted life in parts of Kenya,³⁷ but supervisors at the sample plant reported that logistics and worker absence at the farm were largely unaffected and production continued as usual.³⁸

I interpret a possible increase in taste for discrimination when conflict began as a decrease in the weight attached to the

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

38. Because the workers live on the farm in a gated community, it was safest to remain on the farm.

well-being of non-coethnics, in which case the model predicts an increase in the gap between the average output of homogeneous and mixed teams. In columns (1) and (2) of Table V, the difference in output between mixed and homogeneous teams before and after conflict began is compared. Data from 2007 and the first six weeks of 2008 (when processors were still paid based on own output) are used. There was no significant change in the output of homogeneous teams when conflict began—a (non)response that may cover up countervailing effects of conflict on productivity in homogeneous teams,³⁹ but the focus here is on the difference in output between homogeneous and mixed teams. The output gap between homogeneous and VM teams nearly doubled in early 2008. Output in VM teams decreased by 7% when conflict began. The results in Table V thus indicate that upstream workers undersupply non-coethnic downstream workers to a greater extent during times of ethnic conflict.

Output in horizontally mixed teams decreased by 4% when conflict began, but there was a small but significant increase in the output of coethnic processors in HM 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.

In light of the model, the results for the conflict period thus suggest that discriminatory attitudes toward 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 II, the increased output gap between homogeneous and mixed teams shows no sign of decay in the nine months after conflict ebbed (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.

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

39. Disruption 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.

TABLE V
 OUTPUT BY TEAM ETHNICITY CONFIGURATION BEFORE AND AFTER CONFLICT, AND UNDER
 TEAM PAY

Sample	(1)	(2)	(3)	(4)
	Preconflict/conflict		Conflict/team pay	
	Log (processor output)	Log (team output)	Log (processor output)	Log (team output)
Constant	8.148*** (0.023)	8.840*** (0.027)	8.053*** (0.028)	8.755*** (0.030)
Horizontally mixed		-0.046*** (0.001)		-0.092*** (0.003)
Horizontally mixed, processor of supplier's ethnicity	0.070*** (0.002)		0.087*** (0.004)	
Horizontally mixed, processor not of supplier's ethnicity	-0.181*** (0.002)		-0.317*** (0.005)	
Vertically mixed	-0.084*** (0.002)	-0.084*** (0.001)	-0.163*** (0.004)	-0.161*** (0.004)
Conflict	-0.009 (0.013)	-0.010 (0.012)		
Horizontally mixed × Conflict		-0.044*** (0.004)		
Horizontally mixed, processor of supplier's ethnicity × Conflict	0.017*** (0.004)			
Horizontally mixed, processor not of supplier's ethnicity × Conflict	-0.131*** (0.005)			
Vertically mixed × Conflict	-0.074*** (0.005)	-0.073*** (0.004)		
Team pay			-0.007 (0.013)	-0.010 (0.013)
Horizontally mixed × Team pay				0.044*** (0.004)
Horizontally mixed, processor of supplier's ethnicity × Team pay			-0.127*** (0.005)	
Horizontally mixed, processor not of supplier's ethnicity × Team pay			0.258*** (0.005)	
Vertically mixed × Team pay			-0.003 (0.005)	-0.003 (0.004)
<i>N</i>	224,730	112,365	204,148	10,2074
Person-position FE?	Yes	Yes	Yes	Yes
Date FE?	Yes	Yes	Yes	Yes
Clustering	Two-way (processor and team)	One-way (team)	Two-way (processor and team)	One-way (team)

Notes. Standard errors in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$. The omitted category is homogeneous teams/processor in homogeneous teams. Data from 2007 and the first six weeks of 2008 are used in the OLS regressions in columns (1) and (2). Data from 2008 are used in the OLS regressions in columns (3) and (4). The outcome variables are deseasonalized, daily output quantities.

to the 2008 political conflict in Kenya. In the next subsection, I analyze how the gap in output between homogeneous and mixed teams was affected six weeks into the conflict period when the plant changed the pay system for processors and thereby altered the incentives faced by biased upstream workers.

IV.C. *Firm's Response to Ethnic Diversity Distortions and the Impact on Productivity*

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.⁴⁰ When a processor is paid in part based on the output of the other processor, free-riding is expected to negatively affect output in all teams, but the framework here predicts an offsetting positive effect in HM teams. Under team pay, suppliers are unable to influence the relative pay of the two processors through relative supply. If the higher output for processors of the supplier's ethnic group observed under individual pay is driven by suppliers' taste for discrimination, a decrease in the output gap between coethnic and non-coethnic processors in HM 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.⁴¹ Figure V displays team and individual output during the three sample periods: preconflict (2007), conflict (the first six weeks of 2008), and the team pay period (February 11 through the rest of 2008). Comparing the second and third periods, the figure clearly indicates that the introduction of team pay had a positive effect on output in HM teams.

Corresponding regression results are in columns (3) and (4) of Table V. I find suggestive evidence of inextensive free-riding among processors: the coefficient on team pay in homogeneous and VM teams is negative but insignificant (the coefficient is slightly bigger and significant if the person-position fixed effects are left out).⁴² Output in HM teams is 4% higher under team pay,

40. According to the supervisors, no other changes were made concurrently.

41. In principle, we could distinguish between a team pay/conflict period and a team pay/postconflict period. But it is unclear exactly when conflict effectively ended, and as already noted, the output gap between teams of difference ethnicity configurations showed no sign of reversion later in 2008.

42. The limited extent of freeriding is noteworthy and interesting in itself. As is clear from Figure I, processors can easily monitor each others' effort. A triangular organization of production may thus be a situation in which free-riding can be

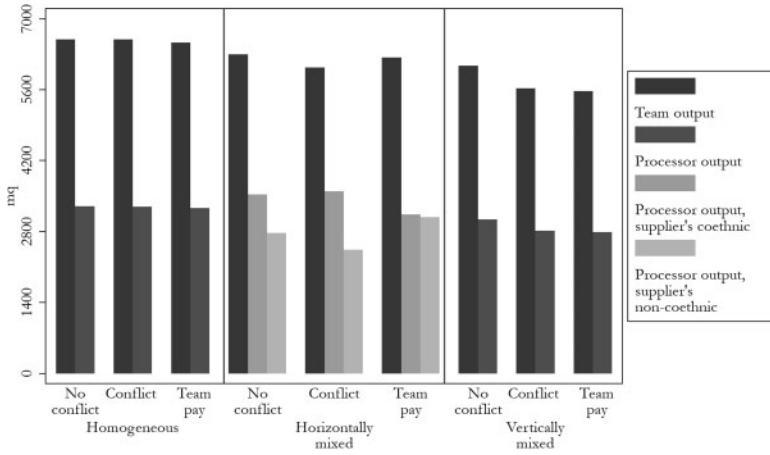


FIGURE V

Output by Team Ethnicity Configuration before and after Conflict and under Team Pay

Data from 2007 and 2008. “Conflict” signifies the first six 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.

as seen in columns (3) and (4) in Table V. The difference in output between HM and homogeneous teams thus decreased significantly when team pay was introduced. The introduction of team pay returned the difference in output between homogeneous and HM teams to preconflict levels.

The increase in HM 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. An output gap of 30% between processors of the supplier’s ethnicity and processors who are not of the supplier’s ethnicity in HM teams disappeared when team pay was introduced.

The positive impact on output in HM teams, which make up half of all teams, led to an overall increase in output when team

effectively dampened through comonitoring. 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 HM and VM teams. Such time-varying factors should not influence the comparison of different types of teams.

pay was introduced.⁴³ However, output in HM teams remains lower than in homogeneous teams under team pay, and output in VM teams still lower. Under team pay a biased supplier continues to derive greater benefit from flowers supplied the more downstream workers belong to her own ethnic group. 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.

In combination with the results in Table III and columns (1) and 2 of Table V, the results in columns (3) and (4) of Table V provide strong support for a taste-based interpretation of the lower output levels observed in ethnically diverse teams. If output was higher in homogeneous teams for informational or technological reasons, there is no obvious reason output in mixed teams would fall differentially during conflict, or why the output of the two processors in HM teams would be equalized under team pay. Some forms of cooperational diversity effects could explain the observed increase in the output of non-coethnic processors in HM teams under team pay. Coethnic processors who can exert effective social pressure on the upstream worker may, for example, induce the supplier to supply more to non-coethnic processors in HM teams under team pay because processors derive benefits from the output of the other processor. However, it is difficult to see how a specific 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 HM teams when team pay is introduced, and the other results of this article taken together, as the framework laid out above can. Though I cannot rule out that other forms of ethnic diversity effects also play a role, I 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.⁴⁴

43. Note that after the conflict period, the farm also hired more workers, perhaps partly to make up for lost capacity due to the decrease in productivity. The newly hired workers worked on other types of flowers than roses and are therefore not observed in the data used in this article.

44. As already discussed, 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

IV.D. *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 greater taste for discrimination than others. Online Appendix Figure 5 plots the distribution of suppliers' discrimination coefficient—the difference in output between homogeneous and mixed teams supplied by a specific worker. It appears that most suppliers discriminate against non-coethnics but there is substantial heterogeneity in the degree of discriminatory behavior. (In the Online 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.)

In column (1) of Table VI, I investigate further by regressing individuals' discrimination coefficient, measured during the second half of 2007, on their characteristics and measures of past exposure to non-coethnics. First, the results show that females are significantly less discriminatory than males, and young workers are less discriminatory than older workers. Second, the coefficient on percent of workdays spent in mixed teams in the past (i.e., during the first half of 2007) is negative, whereas the coefficient on average discrimination coefficient of non-coethnics the worker was supplied by is positive. While imprecisely estimated and not statistically significant,⁴⁵ the signs of the coefficients on measures of past exposure are as we would expect if spending time with non-coethnics reduces discriminatory attitudes (Boisjoly et al. 2006; Rao 2013), while being discriminated against exacerbates such attitudes.

In column (2) of Table VI, I investigate factors that influenced changes in an individual's propensity to discriminate when conflict began. We see that the increase in young workers'

coethnic and non-coethnic processors in HM teams under team pay 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 an HM team under team pay we would expect the effort and output of a coethnic 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 may also explain this study's results.

45. The unit of observation in Table VI is an individual, so the sample size is limited.

TABLE VI
HETEROGENEITY IN DISCRIMINATORY BEHAVIOR

	(1) Discrimination coefficient	(2) Discrimination coefficient
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)
<i>N</i>	675	880

Notes. Standard errors in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$. 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, whereas in column (2) discrimination coefficients during the second half of 2007 and the first six weeks of 2008 are used. "Percent of workdays spent in mixed teams" and "Average discrimination coefficient of non-coethnics worker was supplied by" refer to the first half of 2007.

discrimination coefficient was (marginally) significantly greater than that of older workers. The results in Table VI thus suggest that youth start out relatively tolerant,⁴⁶ but that the attitudes of the young toward non-coethnics respond more negatively to conflict.⁴⁷

46. 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).

47. 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 subsection paint a consistent albeit suggestive picture of how distortionary attitudes toward 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, especially among young workers—a shift that in the data show no sign of decay or reversion in the nine months after conflict ended. Conflict may thus entail significant hidden costs through its influence on the (distortionary) social preferences of those who will determine future economic outcomes.

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

V. THE FIRM'S CHOSEN RESPONSE TO ETHNIC DIVERSITY DISTORTIONS

The data suggest that segregating workers of different ethnicities would be the profit-maximizing response to distortionary discrimination. The results in Tables III and V indicate that segregation would have increased plant productivity by 4% before conflict and by 8% after conflict began, relative to the status quo of arbitrary assignment to teams. Are these expected benefits of a magnitude that is likely to be salient to supervisors? Consider the output increase expected from optimally assigning workers to teams and positions based on ethnicity, productivity, or both. If we view a worker as having three characteristics—the tercile to which she belongs in the distribution of processor productivity, the tercile to which she belongs in the distribution of supplier productivity, and her ethnicity—then an average output will be associated with teams of each of 3 ethnicity configurations, 18 productivity configurations, and 63 ethnicity-productivity configurations.⁴⁸ In theory, supervisors can then solve the linear programming problem of maximizing total output subject to the

48. $63 = \left[3 * \frac{3*(3+1)}{2} \right] + 3^3 + \left[3 * \frac{3*(3+1)}{2} \right]$. 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 $\left[3 * \frac{3*(3+1)}{2} \right]$ homogeneous types of teams and $\left[3 * \frac{3*(3+1)}{2} \right]$ VM types of teams. In HM 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.

expected output associated with a given type of team and the “budget set” of workers available (see Bhattacharya 2009; Graham, Imbens, and Ridder 2010).⁴⁹

The optimal assignments and associated expected output gains are shown in Table VII.⁵⁰ 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. 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 probably infeasible for supervisors to solve. 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 to 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 United States.⁵¹ But as we saw in Table VI, I find no evidence of such habituation affecting behavior at the plant studied.

Although the costs of ethnic segregation are likely incurred primarily by society at large,⁵² 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

49. Bhattacharya and Dupas (2012), Carrell, Sacerdote, and West (2013), and Garlick (2014) 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.

50. The procedure used is explained in detail in the Online Appendix. “Optimal” here means output-maximizing, as inferred from the data. The output-maximizing solution may be undesirable for other reasons discussed later.

51. Boisjoly 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.

52. Alesina 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.

TABLE VII
OUTPUT GAINS FROM OPTIMAL ASSIGNMENT TO TEAMS BY ETHNICITY, PRODUCTIVITY, OR BOTH

Output-maximizing assignment by:	Ethnicity	Productivity as P and S	Ethnicity and productivity as P and S
Preconflict			
Assignment	Homogeneous	100.00%	Homogeneous,s3p2p3 28.95%
			Homogeneous,s3p2p2 25.66%
			Homogeneous,s2p3p3 17.11%
			Homogeneous,s3p3p3 15.79%
			Homogeneous,s2p1p1 6.58%
			Homogeneous,s3p1p3 5.92%
Output gains relative to:			
observed assignment	4.41%	3.93%	9.60%
output-minimizing assignment	8.62%	8.51%	16.47%
Conflict period			
Assignment	Homogeneous	100.00%	Homogeneous,s3p2p3 59.54%
			Homogeneous,s2p2p2 17.11%
			Homogeneous,s2p3p3 13.16%
			Homogeneous,s3p1p3 4.93%
			Homogeneous,s2p1p1 4.61%
			Homogeneous,s2p2p3 0.33%

TABLE VII
(CONTINUED)

Output-maximizing assignment by:	Ethnicity	Productivity as P and S	Ethnicity and productivity as P and S
Output gains relative to:			
observed assignment	8.20%	3.83%	Horiz. mixed,s3p3p1 0.33%
output-minimizing assignment	17.05%	10.17%	15.27%
Team pay period			
Assignment	Homogeneous	100.00%	31.76%
Output gains relative to:			
observed assignment		s3p2p3 2.96%	Homogeneous,s3p2p2 30.92%
output-minimizing assignment		s2p2p3 9.36%	Homogeneous,s3p3p3 30.26%
		s3p2p2 9.87%	Homogeneous,s3p1p3 29.61%
		s1p3p3 9.87%	Homogeneous,s2p3p3 5.26%
			Homogeneous,s2p1p2 3.95%

Notes. X = supplier productivity of tercile X. pX analogous (only productivity tercile in assigned position is shown). For horizontally mixed teams, the first processor position listed refers to coethnics of the supplier (and the second to non-coethnics). 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 Online Appendix.

computed in Table VII is out of sample and therefore does not rule out the possibility that complete segregation between the two ethnic groups would have a negative effect on output. Carrell, Sacerdote, and West (2013) 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 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 lead to increased antagonism outside of the production halls, for example.

It appears that managers preferred adjusting the price of intermediate flowers delivered to non-coethnic processors over a possible technological response (such as moving away from three-person joint production) or changing the composition of teams. In a 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. But it is difficult 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.

VI. CONCLUSION

Although the possibility of a direct negative effect of ethnic diversity on micro-level productivity has long been recognized, corresponding evidence is largely absent. In this article, I began by identifying a sizable negative output effect of ethnic diversity in teams of co-workers in Kenya. I did so using two years of daily output data for 924 workers at a flower-packing plant. 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, upstream workers 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 took 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. 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 biased upstream workers were unable to increase the relative pay of favored downstream workers by distorting relative supply. As a result, horizontal misallocation of flowers was eliminated and total output in teams in which the two downstream workers were of different ethnic groups increased.

I show that less distortionary, non-taste-based ethnic diversity effects are unlikely to explain this study'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). But in most scenarios, we expect firm responses and general equilibrium effects to cushion the impact of micro-level distortions on observed productivity and output. As also pointed out by Becker (1957), discriminatory employers (those who allow workplace discrimination to influence productivity) should go out of business as their profits suffer. However, the Kenyan floriculture business is not very competitive (as evidenced by high entry barriers and profit margins; Noury 2011), nor is it necessarily the most productive firms that survive in poor countries' economies (Banerjee and Moll 2010; Hsieh and Klenow 2009). The location of businesses is partly determined by availability of suitable land, and it appears that other large firms in ethnically mixed parts of Kenya also do not engage in systematic ethnic discrimination in hiring. Although 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.

This article's results do indicate, however, that ethnic diversity affects how firms organize production—including larger firms. If taste for discrimination is high enough, firms may be forced to adopt second best policies to limit the distortions caused by such discrimination. But entirely removing workers'

incentives to discriminate 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 over time.

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 toward 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).

COLUMBIA UNIVERSITY, GRADUATE SCHOOL OF BUSINESS,
AND BREAD

SUPPLEMENTARY MATERIAL

An Online Appendix for this article can be found at QJE online (qje.oxfordjournal.org).

REFERENCES

- Alesina, A., and Eliana La Ferrara, "Ethnic Diversity and Economic Performance," *Journal of Economic Literature*, 43 (2005), 762–800.
- Alesina, Alberto F., Stelios Michalopoulos, and Elias Papaioannou, "Ethnic Inequality," NBER Working Paper No. 18512, 2012.
- Alesina, A., and Enrico Spolaore, "On the Number and Size of Nations," *Quarterly Journal of Economics*, 112 (1997), 1027–1056.
- Alesina, Alberto, and Ekaterina Zhuravskaya, "Segregation and the Quality of Government in a Cross Section of Countries," *American Economic Review*, 101 (2011), 1872–1911.
- Arrow, Kenneth, "The Theory of Discrimination," in *Discrimination in Labor Markets*, Orley Ashenfelter, and Albert Rees, eds. (Princeton, NJ: Princeton University Press, 1973), 3–33.
- Bandiera, O., Iwan Barankay, and Imran Rasul, "Social Preferences and the Response to Incentives: Evidence from Personnel Data," *Quarterly Journal of Economics*, 120 (2005), 917–962.

- , “Social Connections and Incentives in the Workplace: Evidence from Personnel Data,” *Econometrica*, 77 (2009), 1047–1094.
- Banerjee, Abhijit V., and Benjamin Moll, “Why Does Misallocation Persist?,” *American Economic Journal: Macroeconomics*, 2 (2010), 189–206.
- Barkan, Joel D., “Kenya after Moi,” *Foreign Affairs*, 83 (2004), 87–100.
- Bauer, M., Alessandra Cassar, and Julie Chytilová, “Warfare and Social Preferences in Children,” Mimeo, USF.
- Becker, Gary S. *The Economics of Discrimination* (Chicago: University of Chicago Press, 1957).
- , “A Theory of Social Interactions,” *Journal of Political Economy*, 82, no. 6, 1063–1093.
- Bhattacharya, D., “Inferring Optimal Peer Assignment from Experimental Data,” *Journal of the American Statistical Association*, 104 (2009), 486–500.
- Bhattacharya, Debopam, and Pascaline Dupas, “Inferring Welfare Maximizing Treatment Assignment under Budget Constraints,” *Journal of Econometrics*, 167 (2012), 168–196.
- Boisjoly, Johanne, Greg J. Duncan, Michael Kremer, Dan M. Levy, and Jacque Eccles, “Empathy or Antipathy? The Impact of Diversity,” *American Economic Review*, 96 (2006), 1890–1905.
- Buckley, Stephen, “Youth in Kenya Feel Few Tribal Ties,” *Washington Post* 1997, <http://www.washingtonpost.com/wp-srv/inatl/longterm/africanlives/kenya/kenya.htm>.
- Burgess, Robin, Remi Jedwab, Edward Miguel, Ameet Morjaria, and Gerard Padró-i-Miquel, “The Value of Democracy: Evidence from Road Building in Kenya,” NBER Working Paper No. 19398, 2013.
- Carrell, Scott E., Bruce I. Sacerdote, and James E. West, “From Natural Variation to Optimal Policy? The Importance of Endogenous Peer Group Formation,” *Econometrica*, 81 (2013), 855–882.
- Charness, Gary, and Matthew Rabin, “Understanding Social Preferences with Simple Tests,” *Quarterly Journal of Economics*, 117 (2002), 817–869.
- Chen, Yan, and Sherry Xin Li, “Group Identity and Social Preferences,” *American Economic Review*, 99, no. 1, 431–457.
- Dupas, Pascaline, and Jonathan Robinson, “The (Hidden) Costs of Political Instability: Evidence from Kenya’s 2007 Election Crisis,” *Journal of Development Economics*, 99 (2012), 314–329.
- Easterly, W., and Ross Levine, “Africa’s Growth Tragedy: Policies and Ethnic Divisions,” *Quarterly Journal of Economics*, 112 (1997), 1203–1250.
- EDRI, “Ethiopian Development Research Institute: Micro Evidence on the Development of Cut-Flower Industry in Ethiopia,” 2008.
- Eifert, B., Edward Miguel, and Daniel N. Posner, “Political Competition and Ethnic Identification in Africa,” *American Journal of Political Science*, 54 (2010), 494–510.
- Fershtman, Chaim, and Uri Gneezy, “Discrimination in a Segmented Society: An Experimental Approach,” *Quarterly Journal of Economics*, 116 (2001), 351–377.
- Fisman, Raymond, Yasushi Hamao, and Yongxiang Wang, “The Impact of Interstate Tensions on Economic Exchange: Evidence from Shocks to Sino-Japanese Relations,” Mimeo, Columbia University, 2013.
- Fisman, Raymond, Daniel Paravisini, and Vikrant Vig, “Cultural Proximity and Loan Outcomes,” NBER Working Paper No. 18096, 2012.
- Garlick, Robert, “Academic Peer Effects with Different Group Assignment Rules: Residential Tracking versus Random Assignment,” Mimeo, Duke University, 2014.
- Gibson, C., and James Long, “The Presidential and Parliamentary Elections in Kenya, 2007,” *Electoral Studies*, 28 (2009), 497–502.
- Graham, B., Guido Imbens, and Geert Ridder, “Measuring the Effects of Segregation in the Presence of Social Spillovers: A Nonparametric Approach,” NBER Working Paper No. 16499, 2010.
- Hamilton, Barton H., Jack A. Nickerson, and Hideo Owan, “Diversity and Productivity in Production Teams,” *Advances in the Economic Analysis of Participatory & Labor-Managed Firms*, 13 (2012), 99–138.

- Hedegaard, Morten, and Jean-Robert Tyran, "The Price of Prejudice," Rockwool Foundation Research Unit Study Paper No. 32, University Press of Southern Denmark, 2011.
- Hong, Seo Yeon, Pinar Karaca-Mandic, and Nicole Maestas, "Peer Groups and Employment Outcomes: Evidence Based on Conditional Random Assignment in the U.S. Army," Mimeo, RAND, 2008.
- Hsieh, Chang-Tai, and Peter J. Klenow, "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 124 (2009), 1403–1448.
- Jakiela, P., E. Miguel, and V. te Velde, "Combining Field and Lab Experiments to Estimate the Impact of Human Capital on Social Preferences," Mimeo, UC Berkeley, 2011.
- Jones, Charles, "Intermediate Goods and Weak Links in the Theory of Economic Development," *American Economic Journal: Macroeconomics*, 3 (2011), 1–28.
- Kenya Flower Council, "The Flower Industry in Kenya," <http://www.kenyaflowercouncil.org/index.php/2013-03-24-08-12-08/floriculture-in-kenya>, 2011.
- Kranton, Rachel, Matthew Pease, Seth Sanders, and Scott Huettel, "Identity, Groups, and Social Preferences," Mimeo, Duke University, 2013.
- Krueger, A., and Alex Mas, "Strikes, Scabs and Tread Separations: Labor Strife and the Production of Defective Bridgestone/Firestone Tires," *Journal of Political Economy*, 112 (2004), 253–289.
- Ksoll, C., Rocco Macchiavello, and Ameet Morjaria, "The Effect of Ethnic Violence on an Export-Oriented Industry," Mimeo, Warwick University, 2010.
- La Ferrara, Eliana, "Self-Help Groups and Income Generation in the Informal Settlements of Nairobi," *Journal of African Economies*, 11 (2002), 61–89.
- , "Kin Groups and Reciprocity: A Model of Credit Transactions in Ghana," *American Economic Review*, 93 (2003), 1730–1751.
- Lang, K., "A Language Theory of Discrimination," *Quarterly Journal of Economics*, 101 (1986), 363–382.
- Lazear, Edward, "Globalization and the Market for Teammates," *Economic Journal*, 109 (1999), 15–40.
- Macchiavello, Rocco, and Ameet Morjaria, "The Value of Relationships: Evidence from a Supply Shock to Kenya Rose Exports," *American Economic Review* forthcoming (2014).
- Mas, Alexandre, and Enrico Moretti, "Peers at Work," *American Economic Review*, 99 (2009), 112–145.
- Michaels, Guy, and Xiaojia Zhi, "Freedom Fries," *American Economic Journal: Applied Economics*, 2 (2010), 256–281.
- Miguel, Edward, "Tribe or Nation? Nation-Building and Public Goods in Kenya versus Tanzania," *World Politics*, 56 (2004), 327–362.
- Ndegwa, S., "Citizenship and Ethnicity: An Examination of Two Transition Moments in Kenyan Politics," *American Political Science Review*, 91 (1997), 599–616.
- Noury, Valerie, "Kenya's Cut Flower Tax Swindle," *African Business* (2011), <http://62.128.151.219/A1rwqk/AB0511/resources/64.htm>.
- Oyugi, W., "Ethnicity in the Electoral Process: The 1992 General Elections In Kenya," *African Journal of Political Science*, 2 (1997), 41–69.
- Phelps, Edmund, "The Statistical Theory of Racism and Sexism," *American Economic Review*, 62 (1972), 659–661.
- Posner, Daniel L., "Measuring Ethnic Fractionalization in Africa," *American Journal of Political Science*, 48 (2004a), 849–863.
- , "The Political Salience of Cultural Difference: Why Chewas and Tumbukas Are Allies in Zambia and Adversaries in Malawi," *American Political Science Review*, 98 (2004b), 529–545.
- Prat, Andrea, "Should a Team Be Homogeneous?," *European Economic Review*, 46 (2002), 1187–1207.
- Prendergast, Canice, and Robert H. Topel, "Favoritism in Organizations," *Journal of Political Economy*, 104 (1996), 958–978.
- Rao, Gautam, "Familiarity Does Not Breed Contempt: Diversity, Discrimination and Generosity in Delhi Schools," Mimeo, UC Berkeley, 2013.

- Rohner, Dominic, Mathias Thoenig, and Fabrizio Zilibotti, "War Signals: A Theory of Trade, Trust, and Conflict," *Review of Economic Studies*, 80 (2013), 1114–1147.
- Rotemberg, Julio J., "Altruism, Cooperation and Reciprocity in the Workplace," in *Handbook of the Economics of Giving, Altruism and Reciprocity*, vol. 2, Serge-Christophe Kolm, and Jean Mercier Ythier, eds. (Amsterdam: Elsevier Science, 2006), 1371–1435.
- Shayo, Moses, and Asaf Zussman, "Judicial Ingroup Bias in the Shadow of Terrorism," *Quarterly Journal of Economics*, 126 (2011), 1447–1484.
- Shetler, Jan Bender, "Historical Memory as a Foundation for Peace: Network Formation and Ethnic Identity in North Mara, Tanzania," *Journal of Peace Research*, 47 (2010), 639–650.