Discrimination as a Self-Fulfilling Prophecy: Evidence from French Grocery Stores

Dylan Glover
Sciences Po

Amanda Pallais
Harvard and NBER

William Pariente*
Université Catholique de Louvain

January 20, 2015

Abstract

Examining the performance of cashiers in a French grocery store chain, we find that discrimination is a self-fulfilling prophecy. In these stores, cashiers work with different managers on different days and their schedules are determined quasi-randomly. We find that minorities (workers of African and North African descent) perform more poorly when they work with managers biased against them (as determined by an implicit association test or IAT). When scheduled to work with biased managers, minorities are absent more often, spend less time at work, scan items more slowly, and take more time between customers. Minority workers do not appear to perform more poorly with biased managers because they dislike them or because biased managers assign them to more unpleasant tasks. Rather, it appears that biased managers spend less time monitoring minority workers and, as a result, minority workers exert less effort. Manager bias has consequences for the average performance of minority and non-minority workers: while overall these two groups perform equivalently, on days where the manager is unbiased, minorities perform substantially better. This can explain how statistical discrimination can persist even though minority and majority workers have the same average capabilities.

*We would like to thank David Autor, Bruno Crepon, Esther Duflo, and Lawrence Katz for their helpful comments and Frédéric Cochinard, Elizabeth Mishkin, and Adrien Pawlik for superb research assistance. We would also like to thank Thomas Le Barbanchon, Eric Cediey, and ISM CORUM for their advice and help with categorizing names. We gratefully acknowledge funding from the Fonds d'Expérimentation pour la Jeunesse from the Ministère des Sports, de la Jeunesse, de l'Education populaire et de la Vie associative.

1 Introduction

Existing empirical literature has found that employers are less likely to hire and promote minority workers than observably-equivalent workers from majority groups (e.g., Bertrand and Mullianathan 2004; Ritter and Taylor, 2011; Lang and Lehmann, 2012). Theoretical literature has suggested that this may result from statistical discrimination (e.g., Phelps, 1972; Arrow, 1973). Yet, if minority and majority workers have the same average productivity (conditional on observable characteristics), how can differential beliefs about their abilities persist? One explanation is that signals about minority workers are less precise than those of majority workers and, as a result, risk-averse employers are likely to hire minorities (e.g., Phelps, 1972; Aigner and Cain, 1977). Another explanation is that employers' beliefs that minorities are less productive on average are self-reinforcing (e.g., Arrow, 1973; Coate and Loury, 1993). If minority and majority workers make different skill investment or effort choices in the face of these beliefs, minority workers may perform more poorly than majority workers with similar inherent abilities. In this way, these biased beliefs about minority workers can be self-fulfilling.

To our knowledge, this paper provides the first empirical evidence that discrimination is a self-fulfilling prophecy. We examine cashiers in a French grocery store chain. In these stores, there is a sizable minority of workers of North African and Subsaharan African descent and we assess whether these minority cashiers perform worse on the days when they work with managers who are biased against their minority group. There are two main benefits from utilizing this setting. First, cashiers in these stores work with different managers on different days. Because we observe the same worker working with different managers, we do not have to worry that minority workers selected to work with biased managers are more or less productive than minorities selected to work with unbiased managers. Second, the workers we analyze have virtually no control over their schedules. We consider "contrat pro" (CP) workers, new hires on six month contracts, who unlike more senior workers are not allowed to submit schedule preferences. Their schedules are determined by a computer

program which assign shifts to meet predicted demand, taking into account the preferences of the more senior cashiers. Thus, it is not the case that minority workers only work with biased managers on special days when their performance would differ for other reasons.

In our empirical analysis, we include worker fixed effects to examine how a worker's own performance changes when paired with different managers. We do not simply want to compare the performance of minority workers under more- and less-biased managers and attribute any difference in performance to manager bias. There are reasons that workers might perform differently under more- and less-biased managers other than the manager's bias. For example, biased managers may simply be less skilled as managers or they may work under different conditions. To eliminate these confounds, we utilize a difference-in-difference methodology, comparing the change in minority workers' performance under biased and non-biased managers with the change in non-minority performance.

We find strong evidence that minority workers perform worse with biased managers. We measure managers' bias towards African workers using an implicit association test (IAT). Because it is illegal in France to ask workers their ethnicities, workers are categorized into minority and non-minority status based on their names.[1] Using daily administrative data, we find that on days when they are scheduled to work with biased managers, minority workers are more likely to be absent. When they do come to work, they spend less time at work: in particular, they are much less likely to stay after their shift ends.[2] Next, we look at the main metric of performance the store uses, the number of articles per minute scanned. When working with biased managers, minority workers scan articles more slowly. Manager bias does not appear to induce minorities to perform extremely poorly (in the bottom 15%), but otherwise it effects the entire distribution of performance: from making minorities more likely to perform poorly (in the bottom 25%) to making them less likely to perform extremely well. Minorities also take longer between customers when working with biased managers.

[1] ISM CORUM, an expert in discrimination testing in France, did the categorization. We gave ISM CORUM separate lists of first and last names, so that ISM CORUM would not be able to identify any individual in the study.

[2] Workers are allowed to leave when their shift ends, but managers can ask them to stay late.

The differential effects of working with biased managers are neither explained by the fact that unbiased managers are more likely to be minorities themselves or by any other manager characteristics we have.

After showing that minority workers perform worse when working with biased managers, we turn our attention to understanding the mechanism behind the effect. We conducted a phone survey of cashiers, asking them to rank their managers on different dimensions, aimed at testing three potential hypotheses. We find no evidence for the first hypothesis: that minority workers dislike working with biased managers, they believe that biased managers dislike them, or biased managers make them feel less confident in their abilities. We also find no evidence that biased managers assign minority workers to less favorable tasks (as in Lehmann, 2013).[3] If anything, minority cashiers rate biased managers as providing them with better registers and breaks and fewer cleaning duties than do majority workers. We do, however, find evidence for the third hypothesis: biased managers simply put less effort into managing minority workers. Minority workers report that biased managers were less likely to come over to their cashier stations and that biased managers demanded less effort from them. Consistent with this, we find that the effect of manager bias grows during the contract, perhaps as workers may learn that they are not being monitored by biased managers.

Finally, we compare average minority and non-minority performance. Overall, minority and non-minority workers perform similarly. However, on days they work with unbiased managers, minority workers perform substantially better than do non-minority workers. This suggests there are real consequences of manager bias on workers' expected performance and thus, in theory, hiring and promotion rates. One concern is that in this context, minority workers can intertemporally substitute effort from days when they work with biased managers to days with unbiased managers. If minorities exclusively worked with unbiased managers,

[3]In an innovative paper, Lehmann (2013) shows that conditional on observable characteristics, black lawyers are more likely to be hired into top law firms than white lawyers. However, once they are there, they are assigned to worse tasks and less likely to be promoted, though promotion rates are similar conditional on task assignment. This is consistent with the predictions of a model where initial affirmative action and subseqent statistical discrimination lead to worse task assignment and lower promotion rates for blacks.

this shifting would not be possible. However, we do not find evidence in favor of this hypothesis. Minorities are no less responsive to bias when they happen to be scheduled to work with biased managers for longer periods of time. (We can examine periods of up to two weeks.)

The paper proceeds as follows. Section 2 describes the empirical context, while Section 3 describes the data and provides descriptive statistics on the sample. Section 4 tests the identifying assumption, showing that minority and majority workers have similar schedules and work with biased and unbiased managers under similar conditions. Section 5 provides the main results of the paper, showing that minorities perform worse when working with biased managers across a variety of outcomes. Section 6 uses performance and survey data to shed light on the mechanisms behind this effect. Section 7 shows the effect of manager bias on the comparison of minority and non-minority performance and that the effect of manager bias does not decrease when workers are exposed to manager bias for longer periods. Section 8 concludes.

2 Setting

We study entry-level cashiers in a large French grocery store chain. These cashiers are hired on a contrat pro: a six month contract subsidized by the government. In return for the subsidy, the firm trains the employee, both about being a cashier and about the retail sector in general. Apart from the direct subsidy, these contracts are advantageous to firms because they include a week-long trial period before the official start date in which employees are trained without pay. During this week, either party can walk away from the contract without penalty.

CP cashiers complete the same job (running a cash register) as other workers. However, there are two special aspects of their employment. First, one day each week CPs attend training, during which they are not on the store floor. Training days are not included in

our data. Second, CP workers have no control over their schedules. All other cashiers are allowed to submit schedule preferences. A computer system assigns shifts by matching predicted demand to the available workforce, taking the preferences of non-CP workers into account. The computer system is constrained to ensure that workers have the requisite number of days off and that no worker may have more than two split shifts per week, open the store more than twice per week, or close the store more than twice per week. Schedules are determined three weeks at a time and, once determined, the schedule for the next three weeks is publicly posted in the store. The head cashier can, in theory, revise the schedules assigned by the computer system. However, this happens very rarely.

CP workers are hired in waves: approximately twice a year each store has a promotion, where six to ten new CP workers are hired. The managers we study are rarely involved in the hiring process, which is conducted by a central office in Paris and the chief cashier (the managers' boss) at the store. Recruited workers who successfully complete a training week are offered a six-month contract. During the CP, workers have a fixed monthly salary and, in general, are expected to work the same number of hours every month. However, if workers end up working more or less, their paychecks get adjusted. Workers are not paid for days they are absent.[4] However, after three sick days (and a doctor's authorization), the government pays workers 70% of their pay for the duration of their sick leave.

In each store, the cashier manager on duty sits in a special station in the middle of the cashiers. The stores have about five cashier managers (henceforth managers) and 100 to 250 cashiers. There are 30 to 80 registers in each store, though it is rare that all the registers are open. When a cashier arrives at the store, she first badges in near the manager station. She typically has a brief conversation with the manager, who gives her the day's news and assigns her to a register. Some workers are assigned to special cash registers, such as the 10 items or fewer line, though this is rare for CPs. Then, a worker gets her till (cash box) from the safe, sets it up at her station, and starts receiving customers. During the shift, the manager roams

[4] CPs earn vacation time, but cannot use it for days off: they are paid their vacation days after the end of the contract.

the store, monitoring workers at their registers. Managers decide when workers can go on break. They also manage the lines, opening and closing new lines and directing customers to short lines. Managers assign cashiers whose lines are closed to other tasks such as aisle arrangement, the welcome desk, or assisting managers. Workers are allowed to leave at the end of their scheduled shifts, but they are sometimes asked by the manager on duty to stay later. Before leaving the store, cashiers confer with the manager, return their tills to the safe, and badge out near the manager station.

The firm considers showing up to work, showing up on time, and having the correct amount of money in the cash register to be important performance metrics. During the CP, workers can only be fired for misconduct, which includes having more than three unexcused absences or being late to work more than three times. Workers can, of course, leave voluntarily. Having one large till deviation, violent conduct, or more than three warnings for poor behavior can also count as misconduct. If misconduct occurs, the chief cashier decides whether to fire the CP worker, relying on the advice of the managers. Aside from the above, the most important indicator of cashier performance is the number of articles scanned per minute. Each week, a list of workers' average articles per minute scanned is publicly posted where employees can see it.

There are two other determinants of the line speed: (1) inactivity time: the time between finishing one customer's transaction and beginning to scan another customer's items and (2) transaction time: the time between when a customer's last item is scanned and the completion of the customer's transaction, during which the customer is paying. In these stores, customers bag their items themselves and can do so while the cashier is scanning the next customer's items. While the firms tracks both of these metrics, they are not emphasized as key measures of performance or measures that affect workers' position in the stores.

Less than half of workers (about 30% to 40%) are offered another contract after their initial CP. The chief cashier decides whether to offer subsequent contracts to each worker based on the worker's performance, manager evaluations, and the number of available positions at

the store. These subsequent contracts are of longer duration and pay higher salaries.

3 Data and Descriptive Statistics

3.1 Data Sources

We utilize three sources of data: store administrative data, manager survey data, and worker survey data. The store administrative data provide information on worker and manager schedules and time worked as well as worker performance. The manager surveys provide our measure of manager bias. Finally, we use the worker surveys to learn about the on-the-job interactions between workers and managers.

3.1.1 Administrative Store Data

We collected daily data for each CP in a given store over a six-week period from July 2011 to August 2012.[5] Because we wanted variation of the timing of the observations during the contract, we asked some stores for weeks three through eight of the contract, while we asked others for weeks 18 to 23.[6] These data include information on workers' and managers' schedules (the precise times at which they were supposed to enter and leave) as well as badge data (the precise times they actually entered and left). Both managers and CPs must badge in and out for breaks as well as at the beginning and end of shifts, so we have actual working times to the minute.

We also have daily CP performance data for the same time period.[7] In particular, we have data on the average articles per minute scanned, average inactivity time (number of seconds between customers), and the average transaction time (seconds between scanning

[5] We asked for data from each store in France who had CPs and could give us data. Of those, these 33 stores were the only ones to agree and provide usable data.

[6] Which stores provided which weeks was determined by the timing of the CP promotion. Because stores only kept data for one year, stores with early promotions provided weeks 18 to 23, while stores with late promotions provided weeks three through eight. Some stores provided data on one early promotion and one late promotion.

[7] The store does not collect performance data at a higher frequency than the day level.

the last item and finishing the transaction).[8]

In addition to this data, the chain provided a few other worker and manager characteristics. Most importantly, it provided their names. In France, it is illegal to ask people about their ethnicity. Thus, we utilize workers' names as an indicator of their minority status.[9] ISM CORUM (Inter Service Migrants, Centre d'Observation et de Recherche sur L'Urbain et ses Mutations), a leading specialist in discrimination testing in France, performed the categorization.[10] We provided ISM CORUM with separate lists of first and last names, so that ISM CORUM did not know the name of any individual in our study, much less any information about the workers it classified. Each first and last name was categorized into one of five possible origin types: (1) European, (2) Arab (North African), (3) Sub-Saharan African, (4) Mixed or undetermined, and (5) Other (including names of Turkish and Asian origins). We consider workers with African names as the minorities in this context and categorize as African any worker whose first or last name is of Arab or Sub-Saharan African origins. In the appendix we show results with other definitions of minority status.

Using workers' and managers' names, we were also able to determine their gender. The chain also provided managers' rank (position) within the store and managers' date of birth.

3.1.2 Manager Survey and IAT

To determine managers' bias towards minority workers, we asked them to take an implicit association test (e.g., Greenwald et al., 1998; Nosek et al., 2007). The IAT is widely used, particularly in psychology, to measure unconscious bias. The test involves categorizing two sets of words to the left- and right-hand side of a computer screen. In our case, subjects were presented with (1) names of traditionally French origin (e.g., Jean) or traditionally Arab origin (e.g., Muhammad) and (2) adjectives that describe good employees (e.g., hardworking)

[8] We also asked for till deviations, but the firm does not collect this information at a daily frequency.

[9] We received permission to do this from the CNIL, an independent administrative body in France that has the responsibility for reviewing sensitive data processing.

[10] For a list of ISM CORUM projects and studies please visit, http://www.ismcorum.org. We are very thankful to Thomas le Barbanchon for his valuable insights on name categorization in a the French context.

or bad employees (e.g., lazy). In all rounds, one word at a time (either a name or adjective) comes on to the screen and subjects are told how to categorize it (e.g., positive adjectives to the left, negative adjectives to the right). Subjects are instructed to categorize the words as quickly as possible. In the rounds used for scoring, the names and the adjectives are interspersed. In one of these rounds, subjects are tasked with categorizing French names and negative adjectives to the same side of the screen, while in the other, they are tasked with categorizing Arab names and negative adjectives to the same side. The idea of the test is that if a subject has an implicit association between two concepts (e.g., Arabs and unproductive workers), it should be easier (and therefore quicker) to categorize the words when they are placing those words on the same side of the screen. The test produces a measure of bias that compares the time taken to categorize items when Arab names and negative adjectives are categorized on the same side, relative to when French names and negative adjectives are categorized on the same side.[11]

IAT scores have been found to be correlated with judgements, choices, and psychological responses (Bertrand et al., 2005). For example, IAT scores are correlated with voting behavior (Friese et al., 2007), callback rates of minority job applicants (Rooth, 2010), and doctors providing differential medical treatments by race (Green et al., 2007). Moreover, because each word in IAT is categorized so quickly, it may be harder for subjects to appear unbiased on this test than if consciously asked their opinions. For example, Kim (2003) finds that when asked to fake IAT scores, subjects were unable to do so.

3.1.3 Worker Survey

We conducted a telephone questionnaire from May 2013 to September 2013 to survey former CPs about their relationship with each of their managers. We asked each CP about her demographics and experience, but the heart of the questionnaire comprised CPs ranking their managers on a series of 16 questions. Each question provided a manager trait (e.g., the

[11] The order in which these rounds were completed is randomized. We also included practice rounds to mitigate order effects (Nosek et al., 2007).

manager who liked the worker best). Workers were provided with a list of managers they worked with during the CP and then asked to rate (in order) the top three and bottom three of these managers on the given trait.[12] We did not tell workers managers' IAT scores.

At the time of the survey, approximately 20% of workers were still working at the store. Eliminating these workers (fearing they might be reluctant to provide honest evaluations of managers), provides similar results.

3.2 Descriptive Statistics

We start with a sample of 33 stores that provided us with schedule data for both workers and managers, badge data, and worker performance data and at which at least one manager took the IAT. In total, we have 5,099 observations on 218 CP's at 33 stores. There were 154 managers at these stores; stores on average had 4.7 managers each.

We only know two things about CP's universally: whether they are of North or Sub-saharan African origin based on their names (28% are) and whether they are male (7% are). We know a few more things about managers. Very few of the managers are minorities themselves (6%), while a slightly larger fraction are male (10%). On average, they were 41 years old on January 1, 2012. Just over 77% of managers took the IAT and over 67% of managers who took the IAT showed some bias against North Africans (an IAT score above 0.15), while 51% showed moderate to extreme bias (an IAT score above 0.35). Minority CPs work, on average, in stores where managers are just as biased as do majority CPs. In this data, we do not observe a significant correlation between IAT score and manager gender, age, or level (position) in the store. The point estimate indicates that minority managers are less biased against minority workers, but because of the small number of minority managers this coefficient is not statistically significant.

Throughout the analysis, we limit our sample to days where workers were scheduled to

[12] Most workers had six or fewer managers during the CP. In a pilot, we asked workers to rate all of their managers. Subjects found this difficult and thus there was substantial non-response. Subjects also found the time cost of the survey sufficiently expensive that many asked to stop.

work with at least one manager who took the IAT. These leaves us with a sample of 4,374 on 204 CPs.

On average, workers work approximately seven hours per day. The median shift starts at 10:15am and ends at 8:15pm and workers are scheduled for split shifts on just under half (46%) of days. Working days are distributed relatively evenly Monday through Saturday. (We have relatively few observations on Sundays as the firm only opens on Sundays during December in preparation for the Christmas holiday.) Typically, in addition to the training day, workers are scheduled to work four days per week.

4 Exogeneity

Throughout the paper, we want to interpret any change in performance when minority workers work with biased managers as a causal effect of working with those managers. The key assumption is that minority workers were not systematically scheduled to work with biased managers on days or times when their performance would have been particularly high or low for other reasons. Throughout the paper, we use the times that CP workers are scheduled to work, not the times that they actually work, to construct their exposure to bias. This is important because we find that whether workers actually show up to work and how long they stay depends on the managers they are paired with. Here, we first assess whether minority and majority workers were scheduled to work at similar times under similar conditions. We then analyze whether minority and majority workers were scheduled to work with more and less biased managers at similar times under similar conditions.

Panel A of Table 1 compares the times at which minority and majority workers were scheduled to work. Each column in the panel presents a separate regression of a characteristic of a day a worker was scheduled to work on an indicator for the worker being a minority. We control for store fixed effects, as manager assignment is only random within a store, and we cluster standard errors at the store level.[13]

[13] Clustering standard errors at the store level provides us with 33 clusters. This is just over the number of

The first dependent variable is the bias (IAT score) of the manager(s) the CP was scheduled to work with. For workers who were scheduled to work with multiple managers on a given date, this is a weighted average of the managers' IAT scores, where the weights are the amount of time each worker was scheduled to work with each manager. If we do not have a manager's bias score, we simply omit this manager from the calculation. We might have expected that if minority workers had control over their schedules (or were being assigned to schedules non-randomly by managers), they would have been less likely to work with biased managers. Instead, we see that this difference is not significant and the point estimate goes in the opposite direction. The next column investigates whether minority workers are more likely to work with managers who themselves are minorities. Again, we find no effect. Next, we consider minorities' likelihood of working with male managers and "Level 4" managers. (Level 4 managers are higher in the store hierarchy than the remaining, Level 3, managers.) Again, we see no difference between the likelihood that minority and majority workers will work with different types of managers. Nor do we see a difference in the number of managers they work with on a shift (Column 5).

The next column considers the days worked. Workers may simply systematically scan articles faster on some days than others, for example, because stores are busier. To construct a single measure of how productive workers are on a given date, we calculate the average articles scanned per minute from all other stores on that date. (We exclude average articles per minute at the given store as this depends on the CPs and managers working.) We see no evidence that minority workers work on particularly productive or unproductive days. Next, we look at the times the CPs work during the day: we find that minority workers are no more likely to work early in the day, but are slightly more likely to work late in the day. We think this is likely to be due to random chance as only one of the 28 coefficients we test in this section is significantly different from zero. However, we will show that all of our results on the effect of manager bias are robust to controlling for whether workers are

clusters thought to give consistent estimates. In appendix tables, we show the results are robust to clustering our standard errors at other levels of aggregation.

scheduled to work in the evening. The last two columns shows that majority and minority workers work the same number of hours on a given day and are equally likely to have split shifts. Appendix Table 1 shows that minority and majority workers also work under similar conditions when we do not restrict the sample to days in which they are working with at least one manager who took the IAT.

Panel B assesses whether minority workers work with more and less biased managers under the same conditions as do non-minorities. It regresses the dependent variables in Panel A on an indicator for the worker being a minority, the bias of the manager(s) the worker was scheduled to work with, and the interaction of these two terms. This is similar to the specification we will use to analyze outcomes.

The coefficient on the manager bias variable can be significantly different from zero without violating our key assumption. This term measures how the conditions under which biased managers work (with non-minorities) differ from the conditions under which less-biased managers work (with non-minorities). Managers have some control over their schedules and it is reasonable that different types of managers have different preferences. The coefficient on the minority CP dummy measures the conditions minority CPs work under relative to non-minority CPs when they both work with unbiased managers. We see across the board that these coefficients are insignificant. The coefficient on the interaction terms are the coefficients of interest. These coefficients show how the working conditions of minority CPs change relative to non-minority CPs when they are working with more biased managers. Again, all the coefficients are insignificant.

5 Effect of Manager Bias on Performance

We now turn our attention to assessing whether discrimination is a self-fulfilling prophecy: that is, do minority workers perform worse when paired with biased managers? We first consider absence rates and the amount of time spent at work. These metrics are important

to the firm and directly affect workers' pay. Then, we turn to measures of performance while at work: articles per minute scanned (the most important to the firm), inactivity time, and transaction time.

5.1 Time Spent at Work

CPs are rarely absent: they are absent only 1.6% of days.[14] Absences increase throughout the week, starting at a low of 0.7% on Monday and reaching 2.3% on Saturday. Absences are even higher (2.8%) on the rare occasion that workers work on Sunday. However, absences are not significantly different on days with morning or evening shifts.

Panel A of Table 2 investigates how the likelihood of being absent is affected by the bias of the manager the worker is scheduled to work with. The first column displays the results of regressing an indicator for a worker being absent on a given day on the interaction of the manager's bias and an indicator for the worker being a minority, controlling for the manager's bias and worker fixed effects. The coefficient on the interaction term shows how minority workers' performance changes (relative to non-minorities) when working with a biased manager. The subsequent three columns add additional controls: controls for day of the week, shifts that include early mornings or late evenings, and date fixed effects. We cluster standard errors at the store level.

The table shows that minority CPs are more likely to be absent when scheduled to work with a biased manager and that this effect does not change when controls are added. The effect is large in magnitude: the standard deviation of the manager IAT score is 0.36, thus, the estimates indicate that being paired with a manager who is one standard deviation more biased increases a minority worker's absence rate by 1.1 to 1.2 percentage points, 70% of the mean.

While this is a large effect, we expect this estimate to be attenuated due to measurement error. Workers' names do not provide a perfect measure of minority status and we do not

[14] On average, this absence rate would lead to two absences over the six month period.

have IAT scores for all managers.[15] However, the largest source of measurement error is likely to be that managers' IAT scores are not a perfect measure of discrimination. Nosek et al. (2007) summarizes studies measuring the IAT's reliability over time and finds that scores on different IAT administrations have a correlation of approximately 0.56, an effect that doesn't change with the length between testing. This suggests that if the IAT measures some underlying bias plus noise that is uncorrelated across test administrations, the coefficients are attenuated by a factor of approximately 1.8.

The final column of Table 2A includes as controls a dummy for the manager being a minority and an indicator for the worker and the manager both being minorities.[16] Because there are so few minority managers, we do not estimate these terms precisely. However, importantly, including these terms does not change the coefficient of interest (nor does simply eliminating days with minority managers). Thus, the effect of working with a biased manager appears to result from the manager's bias, not manager's own group affiliation. Throughout the panel, the measured effect of working with a biased manager for majority workers is negative, suggesting that non-minorities are less likely to be absent when scheduled to work with biased managers. However, this effect is always insignificant and smaller than the effect for minority workers.

We next investigate the effect of working with a biased manager on the amount of time minority workers spend at the store. Panel B of Table 2 replicates Panel A where the dependent variable is the difference between the actual number of minutes the CP worked and the number of minutes she was scheduled to. Higher values indicate more time working. If workers arrived and left exactly when they were scheduled to, this variable would be negative as workers are allowed to take short breaks during the day which the manager

[15] It is not clear that asking workers whether they self-identify as a minority, about their birth location, or their parents' birth location would have provided a perfect measure of minority status. What matters in part is whether the manager views the worker as being a minority. Managers may view some first-generation workers of Arab descent as non-minorities, while viewing some third-generation workers as minorities.

[16] When the worker is scheduled to work with more than one manager, the dummy for the manager being a minority is a weighted average of the minority status of the managers the worker was scheduled to work with.

schedules. However, on average, CPs work almost exactly the number of minutes they are scheduled to. On 42% of days, workers actually work more than their schedule indicates.

The panel shows that minorities work less when paired with biased managers. When working with a manager who is one standard deviation more biased, they work about 3.5 fewer minutes (relative to non-minorities). This is about one sixth of the interquartile range. As in the prior panel, the result is robust to the addition of controls and is not driven by the minority status of the manager. The point estimates suggest that for non-minority workers, biased managers actually encourage more work. However, as in the prior panel, this estimate is insignificant and smaller than the differential effects for minorities. It is also less robust to the inclusion of date fixed effects and the minority status of the manager. Thus, it could simply be that this coefficient is picking up differences in the conditions under which biased and unbiased managers work.

There are three main ways that a worker could spend less time at work: she could arrive late, leave earlier, or take longer breaks. On average, workers arrive about four minutes before their shifts start and stay nine minutes after their shifts end, but take 13 minutes of breaks during the day. Table 3 uses the same identification strategy to investigate through which channel biased managers affect minorities' time worked. The results show no significant effect of manager bias on arrival time: in fact, the far-from-significant point estimate suggests that minority workers arrived earlier when working with biased managers. There is also no significant effect of manager bias on break time. However, minority workers left substantially earlier when working with biased managers. Unreported results show that they were not more likely to leave before the end of their shift: if anything, they were more likely to stay until the end. But, they were substantially less likely to stay after.[17] This could be because biased managers were less likely to ask minority workers to stay late or because,

[17]For example, 42% of CPs stayed at least 10 minutes after their shift ended. Minority workers working with a manager one standard deviation more biased were (an insignificant) 3.5 percentage points less likely to do so. Just over 5% of CPs stayed 30 minutes or more after their shift; minority CPs facing a one standard deviation more biased manager were only about half as likely to do so. (The difference is statistically significant.)

when working with a biased manager, minority workers were less likely to agree. The final two columns of Table 3 tell a similar story. Minority workers facing biased managers were not more likely to work less than they were scheduled. But they were significantly less likely to work more than they were scheduled to.

5.2 Performance at Work

We now turn our attention to the effect of manager bias on minority performance while at work. We first consider the number of articles per minute workers scan. This is one of the performance metrics over which workers have the most control and, simultaneously, it is one the performance metrics the firm cares most about. In each store, a list of workers' average articles per minute scanned is posted in the breakroom each week. Stores hope that workers will scan 21 articles per minute, but only 24 percent of worker-days in our data hit that threshold. In our data, workers scanned only 18.5 articles per minute on average, but we observe that workers scan faster as they obtain more experience (see Section 6). There are not large day-of-the-week effects in scanning speed, except that workers are exceptionally slow (1.2 articles per minute slower on average), on the few occasions when they have to work Sundays. Workers also scan articles more slowly on shifts that begin in the early morning.

Table 4A replicates the format of Table 2A to show the effect of manager bias on scanning speed. In all specifications, being scheduled to work with a manager who scored one standard deviation higher on the IAT leads the average minority worker to scan 0.4 items per minute less quickly. This is a relatively large effect given that the standard deviation of articles per minute is 2.9. Thus, a one standard deviation increase in manager bias leads to an eighth of a standard deviation decrease in articles per minute scanned for minority workers. The table indicates that biased managers may cause majority workers to scan articles more quickly, though this effect is not robust to adding date effects.

To the extent that cashiers' performance at work is affected by the bias of managers they actually work with (as opposed to the bias of the managers they are scheduled to work

with), the coefficients in this table can be thought of the "reduced form" for an instrumental variables model in which the bias the cashier was scheduled to face instruments for the bias of the manager she actually works with. However, even though there is some strategic and non-strategic switching of schedules, the bias of the manager a cashier was scheduled to work with and the bias of the manager she actually works with has a correlation of 0.93.

Table 5 investigates how manager bias affects the distribution of worker performance. Does manager bias lead minorities to perform very poorly or does it prevent them from achieving excellent performance? To assess this, we consider the effect of manager bias on dummies for achieving different levels of articles per minute. Manager bias does not appear to effect the probability of very poor performance. On 14% of days, workers scanned 16 articles per minute or fewer. Instead of being driven by manager bias, this poor performance appears to be a hallmark of certain workers. While 10% of workers account for half of days with performance this poor, half of workers never have a day where they scan items this slowly. Above this threshold, however, manager bias appears to have a negative effect on minority performance throughout the distribution. The percentage point effect on reaching different targets is relatively constant, but because the standard errors increase with the target, many estimates are statistically insignificant.

Panel B of Table 4 investigates the effect of manager bias on inactivity time, the amount of time that a cashier spends between the end of one transaction and the beginning of another. While this is not an oft-discussed performance metric in the store, managers can observe how quickly cashiers transition between customers and it directly affects line speed. On average, workers spend about thirty seconds between customers.[18] Across all specifications, working with a biased manager leads workers to take more time between customers. The standard deviation in inactivity time is about 13 seconds. Thus, working with a manager who is one

[18] We have eliminated the 25 observations of observations where workers spent more than two minutes on average between customers throughout the day. We think these are likely data errors or they indicate that something else was going on in the store that was outside of the CP's control. (For example, one observation indicates that a worker spent 49 minutes on average between customers.) Spending over two minutes on average between customers is unrelated to manager bias or the interaction of manager bias and minority status.

standard deviation more biased leads to about a tenth of a standard deviation increase in inactivity time.

Finally, Panel C examines the effect of manager bias on minority workers' checkout time: the time between when a worker is done scanning items and the end of the transaction. More so than any of the other metrics. This depends on the speed at which the customer provides payment as well as the cashier's speed of processing it. As with the previous metric, this is not something that is emphasized by the store, but it something that affects the speed of checkout and that the manager could observe. The point estimates all suggest that working with a biased manager increases the amount of time that minority workers take to finish the transaction. However, the point estimates are small and imprecise and double when the controls are added.

Appendix Tables 2, 3, and 4 assess the robustness of the effects of manager bias on time spent at work and on-the-job performance. Appendix Table 2 shows the findings are robust to clustering the standard errors at different levels of aggregation. Appendix Table 3 shows the results are not driven by any other managers characteristics in our data (manager position in the firm, age, or gender). Appendix Table 4 considers different definitions of minority status. Panel A considers as minorities only workers with either a first or last name of North African origin (and eliminates remaining workers with names of Subsaharan African origin), while Panel B does the reverse. The effects of manager bias on workers of North and Subsaharan African origin are similar. Panel C utilizes the original definition of minority workers, but considers as majority workers only workers who have both a first and last name of European origin (eliminating from the majority group workers of indeterminate, mixed, or other origins). Again, the results are virtually unchanged.

6 Evidence on Mechanisms

There are several reasons why minority workers might perform worse than majority workers when paired with biased managers. Perhaps the most straightforward explanation is that minority workers just dislike biased managers. They may believe that biased managers dislike them or have little confidence in their abilities. Alternatively, managers could assign workers to situations or tasks where they are likely to perform worse or get demoralized (e.g., Lehmann, 2013). For example, managers could assign minority workers to unpleasant registers, additional cleaning duties, or undesirable break times. Finally, biased managers may simply put less effort into managing minority workers. For example, throughout the day, managers stop by workers' stations to check on their progress and handle any problems that arise. Biased managers may do this less because (1) they do not enjoy interacting with minority workers, (2) they believe that minority workers are so unproductive that their monitoring effort will be wasted or (3) they want minority workers to perform poorly.[19] If minorities realize that their managers are not monitoring them, it may be rational to exert less effort, since their effort is less likely to be noticed by management. While we find that minorities perform worse under biased managers, theoretically, each of these mechanisms could go in either direction. For example, minorities may have worked harder under biased managers to disprove managers' negative beliefs or biased managers might have put more effort into managing minority CPs because they thought minorities required more managing.

We conducted a worker survey after the conclusion of the CP to shed light on whether any of these three mechanisms was operational.[20] Workers were provided with a list of managers they worked with during the CP and then asked to rate (in order) the top three and bottom three of these managers on a variety of characteristics related to these three hypotheses.

[19] Managers may want minorities to perform poorly to reinforce their beliefs or so minorities do not receive longer-term contracts.

[20] Another general explanation for why minorities may perform worse under biased managers is that biased managers provide them with less training. We think this is unlikely in this context because (1) the managers do not have the primary responsibilities for training workers and (2) even if there was a difference in manager training, a worker could utilize any skills gained under one manager the next day under the next manager.

Using these answers, we gave each manager a ranking on each question from one (the lowest ranked manager) to N (the highest ranked manager), where N is the number of managers the worker worked with. Each manager's ranking represents the number of other managers that manager rated better than or equal to.

We utilize the questions relating to each potential mechanism to create three indices: "affection," "managerial choices," and "managerial effort." Each index is a simple average of the rankings on the related questions.[21] Table 6 shows the results of regressing each index on our key independent variables, controlling for store fixed effects. Appendix Table 5 shows the results for each question separately. The fact that the columns have different numbers of observations reflects the fact that respondents were less likely to answer some questions than others. This likely reflects the fact that subjects found it easiest to identify whom they liked and who liked them and hardest to differentiate managers on the basis of their management choices. Questions relating to the affection and confidence index were asked last, while questions relating to the monitoring and effort index were asked first, so it is not simply survey fatigue that leads to missing data.

The first column of Table 6 indicates that minority workers did not rate biased managers as worse on the affection index. While minority workers indicated that they liked working with biased managers (insignificantly) less, they actually thought that biased managers were (insignificantly) more likely to recommend them for promotion. Similarly, we do not find that biased managers assigned minority workers to worse tasks. While not significant, the relatively large coefficient indicates that minorities thought biased managers assigned them

[21] A principal component analysis suggests that the first principal component of each set of questions has approximately equal weights. Moreover, because the individual questions are rankings on the same scale, they have approximately the same means and standard deviations. The affection index includes rankings of (1) the manager who was most likely to recommend the worker for promotion, (2) who initially made the worker feel the most confident, (3) who liked the worker best, and (4) whom the worker enjoyed working with best. The management choices index includes rankings of (1) whose management of lines and customer flow allowed the worker to perform best, (2) who assigned the worker to cleaning duties least often, (3) who assigned the worker to her preferred register type most often, and (4) who assigned the worker the best breaks. The managerial effort index includes rankings of (1) who the worker had the most interaction with (when both were on duty), (2) who most checked up on the CP's work to make sure it was up to standard, (3) who most noticed how well or poorly the cashier was performing, (4) who demanded the worker perform the hardest, and (5) under whom the cashier exerted the most effort.

to *better* tasks. For example, biased managers were more likely to assign minorities to their preferred register types (significant at the 10% level) and (insignificantly) less likely to assign them to cleaning duties. This is consistent with the finding that minority workers were less likely to stay late when working with biased managers.

However, we do find evidence that biased managers spend less effort managing minority workers. Minorities stated that biased managers were less likely to check up on their work to make sure it was up to standard and that biased managers did not demand they work as hard. Minority workers also said reported they exerted less effort under these managers.

Table 7 investigates how the effect of bias evolves over the contract, focusing on articles per minute. We focus on articles per minute in all but the main analyses because articles per minute has continuous variation, it is mostly determined by workers as opposed to managers or store conditions, and it is important to the store. The first column utilizes both across- and within-person variation in observation timing, interacting the key independent variables with the number of weeks since the beginning of the contract, controlling for the main effect of time worked at the store. The second and third columns utilize only the across-store variation and are limited to stores which provide early week and late week data, respectively. Appendix Table 6 shows that these sets of stores have similar characteristics. In fact, 13 stores are included in both regressions because they have one promotion in the early weeks and one promotion in the later weeks.

Both identification strategies suggest that workers become more productive over time. The first indicates that workers scan an additional 0.21 articles per minute each week, while the second indicates that in the second half of the contract, workers scan approximately two articles per minute faster than they do in the first half. Minority workers improve their scanning speed at the same rate as do non-minority workers. The table shows that the effect of the bias also increases substantially over time. Each week, the effect of having a manager one standard deviation more biased increases by 7% of the mean. In a month, the effect of the bias is expected to increase by 30% of the mean. Similarly, the effect of the bias is

estimated to be four times larger in the second half of the contract than in the first half.

This result is consistent with the fact that minorities report that biased managers spend less effort managing them. At the beginning of their contracts, workers may not know which managers are carefully monitoring them and which managers are not. Over time, minority workers observe that certain (biased) managers stop by their stations less often and require less effort from them. In response, they may put in less effort when these managers are on duty. This result is less consistent with a task assignment explanation in which managers assign workers to registers or breaks which make them unproductive. Under this explanation, biased managers would assign minority cashiers to these unproductive conditions from the beginning of their contracts, so manager bias would lead minority workers to be less productive from the start.

7 Comparison of Minority and Majority Performance

We have seen that minority workers perform worse when working with biased managers. Table 8 examines the consequences of this for the relative performance of minority and majority workers. Panel A simply regresses each performance metric on an indicator for the worker being a minority, conditioning on store fixed effects. Aside from having about 1.6 seconds longer transaction time, minority workers' performance is statistically indistinguishable from that of majority workers. However, Panel B compares minority and majority workers' performance on days when they work with unbiased managers: managers with IAT scores between -0.15 and 0.15.[22] On days when workers spend at least half the day with unbiased managers, minority workers perform substantially better than non-minority workers. They are approximately half as likely to be absent and they scan 0.8 articles per minute faster. Thus, while overall minority and majority workers perform similarly, this analysis suggests that if minority and majority workers both worked in a workplace without bias, minority

[22] Scores with absolute values below 0.15 represent little to no bias, while scores above 0.35 in absolute value represent moderate to strong bias.

workers would perform better.

One caveat to this conclusion is that perhaps minority workers perform particularly well with unbiased managers and particularly poorly with biased managers because they work both with biased and unbiased managers. Minority workers may intertemporally substitute effort over time towards days when their effort is more effective (as in Farber, 2005 and Oettinger, 1999). For example, imagine a worker needs to take a personal day. She would take it regardless of whether she worked full time with a biased manager or full time with an unbiased manager. However, if she works with both types of managers, she can choose to take the personal day on the day she is scheduled to work with the biased manager. If this is the case, the estimates in Table 8 overstate minority performance when working with unbiased managers.[23]

We do not have an experiment where we assign workers randomly to only one manager for long periods of time. However, we can shed light on this hypothesis by looking at workers' response to bias when they work with more and less biased managers for larger durations of time. If, for example, workers plan to be absent once during the week regardless of the manager on duty, absences should not respond to manager bias at the week level. Table 9 considers majority and minority workers' response to bias at longer levels of time aggregation. We first consider blocks of two consecutive working days, then a calendar week (typically four working days), and two calendar weeks. Panel A shows the results for articles per minute and, because the concern is particularly compelling with absences, Panel B examines this outcome. In neither panel, does the measured effect of bias on minority worker performance decrease with the level of aggregation. For absences, the measured effect of discrimination is relatively constant with the level of aggregation, while with articles per minute, the coefficient increases as the level of aggregation increases. This may indicate that

[23] Under other mechanisms for the effect, the estimates in Table 7 could underestimate the performance differences between minority and majority workers working only with unbiased managers. For example, imagine that biased managers train minority workers less. In our context, minority performance under unbiased managers also suffers because of the lack of training they receive under biased managers. Similarly, if minority workers are less eager to continue working at the firm because they have to work with biased managers, this may lead them to work less hard even on the days they work with unbiased managers.

there are some cumulative effects of manager bias on scanning speed.

8 Conclusion

The paper has shown that working with biased managers causes minority workers to perform more poorly. Cashiers working with biased managers were more likely to be absent and spent less time at work. Given that workers are paid based on time worked, this leads to a loss in wages for minority workers. When at work, they scanned items more slowly, and took more time between customers. Yet, it does not appear that biased managers were actively antagonistic towards minority workers. Minority workers do not report disliking biased managers or thinking biased managers dislike them. If anything, biased managers assigned minorities to more pleasant tasks and were less likely to ask them to stay late. Instead, it appears that biased managers simply spent less time with minority workers, leading these workers to exert less effort. Given that we focus on implicit bias, managers may not have even realized they harbored these biases or treated minority workers differently.

The results are striking given that this is a setting where we might have expected little effect of manager bias. Managers and workers have relatively little interaction, managers are not responsible for training workers, and workers can complete their jobs with little input from managers. One might imagine that the effects of manager bias would be even larger in a setting where there is more potential for managers to affect worker performance.

References

Aigner, D. J. and G. G. Cain (1977). Statistical theories of discrimination in labor markets. *Industrial and Labor relations review*, 175–187.

Arrow, K. (1973). The theory of discrimination. *Discrimination in labor markets 3*(10), 3–33.

Bertrand, M., D. Chugh, and S. Mullainathan (2005). Implicit Discrimination. *American Economic Review: Papers and Proceedings 95*(2), 94–98.

Bertrand, M. and S. Mullainathan (2004). Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. *The American Economic Review 94*(4), 991–1013.

Coate, S. and G. C. Loury (1993). Will affirmative-action policies eliminate negative stereotypes? *The American Economic Review*, 1220–1240.

Farber, H. (2005). Is tomorrow another day? the labor supply of new york city cab drivers. *Journal of Political Economy 113*(1), 46–82.

Friese, M., M. Bluemke, and M. Wänke (2007). Predicting voting behavior with implicit attitude measures: the 2002 german parliamentary election. *Experimental Psychology 54*(4), 247.

Green, A. R., D. R. Carney, D. J. Pallin, L. H. Ngo, K. L. Raymond, L. I. Iezzoni, and M. R. Banaji (2007). Implicit bias among physicians and its prediction of thrombolysis decisions for black and white patients. *Journal of General Internal Medicine 22*(9), 1231–1238.

Greenwald, A. G., D. E. McGhee, and J. L. Schwartz (1998). Measuring individual differences in implicit cognition: the implicit association test. *Journal of personality and social psychology 74*(6), 1464.

Kim, D.-Y. (2003). Voluntary controllability of the implicit association test (iat). *Social Psychology Quarterly 66*(1), 83–96.

Lang, K. and J.-Y. K. Lehmann (2012). Racial discrimination in the labor market: Theory and empirics. *Journal of Economic Literature 50*(4), 959–1006.

Lehmann, J.-Y. K. (2013). Job assignment and promotion under statistical discrimination: Evidence from the early careers of lawyers. *Working Paper*.

Nosek, B. A., A. G. Greenwald, and M. R. Banaji (2007). The implicit association test at age 7: A methodological and conceptual review. *Social psychology and the unconscious: The automaticity of higher mental processes*, 265–292.

Oettinger, G. S. An empirical analysis of the daily labor supply of stadium vendors. *Journal of Political Economy 107*, 360–392.

Phelps, E. S. (1972). The statistical theory of racism and sexism. *The american economic review*, 659–661.

Ritter, J. A. and L. J. Taylor (2011). Racial disparity in unemployment. *The Review of Economics and Statistics 93*(1), 30–42.

Rooth, D.-O. (2010). Automatic associations and discrimination in hiring: Real world evidence. *Labour Economics 17*(3), 523–534.

Table 1. Exogeneity of Scheduled Shifts

	Manager Bias (IAT)	Minority Manager	Male Manager	Level 4 Manager	Total Managers	Articles per Min in Other Stores (on that Date)	Morning (9am or earlier)	Evening (8pm or Later)	Total Hours	Split Shift
					A. Minority vs. Majority Shifts					
Minority Worker	0.024	0.000	0.003	-0.003	0.014	0.014	0.015	0.023*	0.016	-0.007
	(0.016)	(0.003)	(0.014)	(0.004)	(0.067)	(0.090)	(0.017)	(0.012)	(0.032)	(0.017)
					B. Majority vs. Minority Shifts when Working with More- and Less-Biased Managers					
Minority Worker ×		0.031	0.029	0.009	-0.023	0.022	-0.018	0.018	0.086	-0.007
Manager Bias		(0.022)	(0.026)	(0.033)	(0.069)	(0.132)	(0.023)	(0.020)	(0.098)	(0.028)
Minority Worker		-0.010	-0.012	-0.009	0.009	0.015	0.022	0.016	-0.014	-0.003
		(0.009)	(0.009)	(0.012)	(0.042)	(0.072)	(0.021)	(0.012)	(0.048)	(0.022)
Manager Bias		-0.068	0.187*	0.084	0.589*	-0.394	0.016	-0.001	-0.120	-0.079*
		(0.048)	(0.107)	0.079	(0.299)	(0.238)	(0.023)	(0.039)	(0.140)	(0.040)
Observations	4,374	4,374	4,374	4,374	4,374	4,241	4,374	4,374	4,371	4,374
Mean of Dep Var	0.40	0.06	0.11	0.17	2.77	18.23	0.14	0.58	7.22	0.46
Store FE.'s	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: Standard errors are clustered at the store level.

*** significant at the 1% level; ** significant at the 5% level; * significant at the 10% level

Table 2. Effect of Manager Bias on Time Spent at Work

	A. Dependent Variable: Absence Indicator				
Minority Worker ×	0.0315***	0.0305***	0.0307***	0.0355***	0.0360***
Manager Bias	(0.0102)	(0.0106)	(0.0106)	(0.0112)	(0.0113)
Manager Bias	-0.0059	-0.0063	-0.0061	-0.0148	-0.0156
	(0.0088)	(0.0090)	(0.0091)	(0.0119)	(0.0123)
Minority Worker ×					0.0082
Minority Manager					(0.0995)
Minority Manager					-0.0061
					(0.0154)
Observations	4,374	4,374	4,374	4,374	4,374
Mean of Dep Var	0.016	0.016	0.016	0.016	0.016
R-squared	0.349	0.350	0.351	0.403	0.403
	B. Dependent Variable: Minutes Worked Relative to Scheduled Minutes				
Minority Worker ×	-9.60**	-9.70**	-9.70**	-9.03**	-8.77**
Manager Bias	(4.067)	(4.057)	(4.108)	(4.052)	(4.031)
Manager Bias	5.273	5.21	5.176	3.125	2.658
	(3.261)	(3.262)	(3.354)	(2.889)	(2.982)
Minority Worker ×					0.421
Minority Manager					(10.776)
Minority Manager					-3.695
					(4.732)
Observations	4,166	4,166	4,166	4,166	4,166
Mean of Dep Var	-0.062	-0.062	-0.062	-0.062	-0.062
R-squared	0.212	0.213	0.218	0.314	0.314
Individual F.E.'s	Yes	Yes	Yes	Yes	Yes
Day of the Week F.E.'s	No	Yes	Yes	No	No
Morning/Evening F.E.'s	No	No	Yes	Yes	Yes
Date F.E.'s	No	No	No	Yes	Yes

Note: Standard errors are clustered at the store level.

*** significant at the 1% level; ** significant at the 5% level; * significant at the 10% level

Table 3. Additional Results on Time Spent at Work

	Minutes Arrived Before Shift Starts	Minutes Stayed After Shift Ends	Break Time	Worked > 10 Minutes Less than Scheduled	Worked > 10 Minutes More than Scheduled
Minority Worker ×	4.95	-9.83*	3.89	0.023	-0.134**
Manager Bias	(5.94)	(5.32)	(4.29)	(0.088)	0.063
Manager Bias	1.14	-1.68	-3.2	-0.011	0.026
	(4.28)	(2.85)	(2.73)	(0.076)	0.040
Minority Worker ×	-6.99	6.67	-0.75	0.089	-0.366**
Minority Manager	(5.54)	(7.30)	5.096	0.101	0.140
Minority Manager	-4.19	-2.89	-3.38	-0.132*	-0.022
	(3.69)	(4.92)	(2.84)	(0.070)	(0.071)
Observations	4,166	4,166	4,166	4,166	4,166
Mean of Dep Var	3.77	8.93	12.76	0.341	0.233
R-squared	0.209	0.191	0.453	0.275	0.321
Individual F.E.'s	Yes	Yes	Yes	Yes	Yes
Day of the Week F.E.'s	No	No	No	No	No
Morning/Evening F.E.'s	Yes	Yes	Yes	Yes	Yes
Date F.E.'s	Yes	Yes	Yes	Yes	Yes

Note: Standard errors are clustered at the store level.

*** significant at the 1% level; ** significant at the 5% level

Table 4. Effect of Manager Bias on Work Performance

	A. Dependent Variable: Articles Scanned per Minute				
Minority Worker ×	-1.159**	-1.158**	-1.158**	-1.017**	-1.060**
Manager Bias	(0.549)	(0.538)	(0.544)	(0.485)	(0.486)
Manager Bias	0.399*	0.396*	0.398*	0.168	0.222
	(0.235)	(0.233)	(0.235)	(0.177)	(0.197)
Observations	3,606	3,606	3,606	3,606	3,606
Mean of Dep Var	18.53	18.53	18.53	18.53	18.53
R-squared	0.597	0.601	0.601	0.675	0.675
	B. Dependent Variable: Inactivity Time (Seconds between Customers)				
Minority Worker ×	3.456**	3.403**	3.358**	4.084**	3.925**
Manager Bias	(1.637)	(1.628)	(1.588)	(1.764)	(1.804)
Manager Bias	-1.841	-1.507	-1.605	-1.820	-1.613
	(1.107)	(1.115)	(1.084)	(1.512)	(1.546)
Observations	3,289	3,289	3,289	3,289	3,289
Mean of Dep Var	30.43	30.43	30.43	30.43	30.43
R-squared	0.402	0.450	0.453	0.541	0.541
	C. Dependent Variable: Checkout Time (Seconds)				
Minority Worker ×	0.467	0.758	0.722	0.894	1.039
Manager Bias	(1.383)	(1.383)	(1.383)	(1.158)	(1.102)
Manager Bias	0.086	0.131	0.095	-1.276	-1.417*
	(0.892)	(0.814)	(0.803)	(0.870)	(0.761)
Observations	3,110	3,110	3,110	3,110	3,110
Mean of Dep Var	50.77	50.77	50.77	50.77	50.77
R-squared	0.352	0.376	0.377	0.455	0.455
Individual F.E.'s	Yes	Yes	Yes	Yes	Yes
Day of the Week F.E.'s	No	Yes	Yes	No	No
Morning/Evening F.E.'s	No	No	Yes	Yes	Yes
Date F.E.'s	No	No	No	Yes	Yes
Manager Minority Status	No	No	No	No	Yes

Note: Standard errors are clustered at the store level. The regressions in the final column in each panel control for manager minority status and its interaction with worker minority status.

*** significant at the 1% level; ** significant at the 5% level; * significant at the 10% level

Table 5. Effect of Manager Bias on Distribution of Articles per Minute
Dependent Variable: Scanning at Least the Indicated Number of Articles per Minute

	15	16	17	18	19	20	21	22	23
Minority Worker × Manager Bias	0.015	-0.008	-0.121***	-0.132**	-0.139**	-0.106	-0.076	-0.106	-0.095
	(0.043)	(0.051)	(0.039)	(0.050)	(0.056)	(0.070)	(0.116)	(0.075)	(0.068)
Manager Bias	0.019	0.006	0.017	0.018	-0.014	-0.043	-0.014	0.005	0.005
	(0.018)	(0.017)	(0.029)	(0.033)	(0.043)	(0.032)	(0.030)	(0.028)	(0.022)
Observations	3,606	3,606	3,606	3,606	3,606	3,606	3,606	3,606	3,606
Mean of Dep Var	0.927	0.864	0.760	0.636	0.494	0.356	0.235	0.136	0.074
R-squared	0.503	0.518	0.585	0.595	0.593	0.596	0.564	0.520	0.476
Individual F.E.'s	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day of the Week F.E.'s	No	No	No	No	No	No	No	No	No
Morning/Evening F.E.'s	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Date F.E.'s	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: Standard errors are clustered at the store level.
*** significant at the 1% level; ** significant at the 5% level; * significant at the 10% level

Table 6. Evidence on Mechanisms

	Affection	Management Choices	Managerial Effort
Minority Worker ×	-0.081	0.525	-0.552**
Manager Bias	(0.366)	(0.358)	(0.248)
Manager Bias	-0.494	-0.870**	-0.301
	(0.453)	(0.390)	(0.349)
Minority Worker	0.078	-0.250	0.212
	(0.198)	(0.239)	(0.178)
Observations	856	397	660
Mean of Dep Var	3.90	3.89	3.96
R-squared	0.355	0.239	0.450
Individual F.E.'s	No	No	No
Store F.E.'s	Yes	Yes	Yes

Note: Standard errors are clustered at the store level. The dependent variable indices are defined in footnote 21 of the text.

*** significant at the 1% level; ** significant at the 5% level; * significant at the 10% level

Table 7. Effect of Bias Early vs. Late in the Contract

	All Observations	Early Weeks	Late Weeks
Minority Worker ×	-0.236	-0.566	-2.022*
Manager Bias	(0.515)	(0.412)	(1.039)
Minority Worker ×	-0.076		
Manager Bias × Week of	(0.062)		
Manager Bias	0.110	0.145	0.844**
	(0.365)	(0.356)	(0.309)
Manager Bias × Week of	0.016		
Contract	(0.026)		
Minority Worker × Week of	0.001		
Contract	(0.070)		
Week of Contract	0.205***		
	(0.035)		
Observations	3,588	2,406	1,200
Mean of Dep Var	18.53	17.88	19.82
R-squared	0.610	0.642	0.452
Individual F.E.'s	Yes	Yes	Yes

Note: Standard errors are clustered at the store level.

*** significant at the 1% level; ** significant at the 5% level; * significant at the 10% level

Table 8. Comparison of Minority and Non-Minority Performance

	Absences	Minutes Worked (Relative to Scheduled)	Articles per Minute	Inactivity Time (Seconds)	Transaction Time (Seconds)
			A. All Days		
Minority Worker	-0.004	0.492	0.239	0.751	1.618**
	(0.007)	(2.306)	(0.295)	(0.767)	(0.651)
Non-Minority Mean	0.0186	-1.189	18.55	28.20	50.06
Observations	4,374	4,166	3,603	3,289	3,110
Store FE.'s	Yes	Yes	Yes	Yes	Yes
			B. Days with Unbiased Managers		
Minority Worker	-0.014*	2.627	0.805*	-1.748	1.769*
	(0.007)	(2.348)	(0.364)	(1.324)	(0.862)
Non-Minority Mean	0.0274	-4.489	18.58	26.81	48.36
Observations	475	437	360	322	301
Store FE.'s	Yes	Yes	Yes	Yes	Yes
			C. Days with Biased Managers		
Minority Worker	-0.009	0.245	-0.152	1.098	1.550*
	(0.010)	(2.996)	(0.360)	(1.075)	(0.790)
Non-Minority Mean	0.0217	-1.020	18.71	28.08	50.00
Observations	2,850	2,727	2,344	2,104	2,039
Store FE.'s	Yes	Yes	Yes	Yes	Yes

Note: Standard errors are clustered at the store level. Days with unbiased managers are days where the worker spent at least 50% of the day with managers with an IAT between -0.15 and 0.15. Days with biased managers are days where the worker spent more than 50% of the day with manager(s) whose IAT score exceeds 0.35.

*** significant at the 1% level; ** significant at the 5% level; * significant at the 10% level

Table 9. Effect of Manager Bias on Minority Performance
Different Levels of Time Aggregation

	Day	Two Working Days	Week	Two Weeks
		A. Articles per Minute		
Minority Worker × Manager Bias	-1.159**	-1.104**	-2.132***	-3.186**
	(0.549)	(0.428)	(0.616)	(1.410)
Manager Bias	0.399*	0.616*	1.117**	2.035*
	(0.235)	(0.356)	(0.469)	(1.148)
Observations	3,606	2,151	1,112	606
Mean of Dep Var	18.53	18.53	18.53	18.53
R-squared	0.597	0.723	0.802	0.866
		B. Absences		
Minority Worker × Manager Bias	0.0315***	0.0095	0.0334	0.0434
	(0.010)	(0.013)	(0.033)	(0.082)
Manager Bias	-0.0059	0.0021	-0.0139	-0.0131
	(0.009)	(0.011)	(0.029)	(0.074)
Observations	4,374	2,386	1,209	651
Mean of Dep Var	0.016	0.016	0.016	0.016
R-squared	0.3492	0.4300	0.4688	0.5491
Individual F.E.'s	Yes	Yes	Yes	Yes

Note: Standard errors are clustered at the store level.
*** significant at the 1% level; ** significant at the 5% level; * significant at the 10% level

Appendix Table 1. Exogeneity of Scheduled Shifts
Including Observations without Manager Bias Scores

	Manager Bias (IAT)	Minority Manager	Male Manager	Level 4 Manager	Total Managers	Articles per Min in Other Stores (on that Date)	Morning (9am or earlier)	Evening (8pm or Later)	Total Hours	Split Shift
					A. Characteristics of Foreign Workers' Shifts					
Minority Worker	0.024	-0.009	-0.008	-0.009	-0.088	0.013	0.010	0.016	0.036	-0.007
	(0.016)	(0.007)	(0.017)	(0.010)	(0.130)	(0.079)	(0.015)	(0.011)	(0.027)	(0.014)
Observations	4,374	5,099	5,099	5,099	5,099	4,945	5,099	5,099	5,094	5,099
Mean of Dep Var	0.40	0.05	0.12	0.16	2.46	18.28	0.14	0.58	7.20	0.46
Store FE.'s	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: Standard errors are clustered at the store level.

Appendix Table 2. Clustering Standard Errors at Different Levels of Aggregation

	Absence Indicator	Minutes Worked (Relative to Schedule)	Articles per Minute	Inactivity Time (Seconds)	Checkout Time (Seconds)
			A. Store - Week Level		
Minority Worker × Manager	0.0357*	-9.03*	-1.017**	4.084**	0.894
Bias	(0.0201)	(5.11)	(0.454)	(1.758)	(1.435)
Manager Bias	-0.0147	3.13	0.168	-1.820*	-1.276*
	(0.0108)	(2.95)	(0.232)	(1.079)	(0.745)
			B. Individual Level		
Minority Worker × Manager	0.0357**	-9.03	-1.017**	4.084**	0.894
Bias	(0.0178)	(5.94)	(0.401)	(1.848)	(1.499)
Manager Bias	-0.0147	3.13	0.168	-1.8195*	-1.276
	(0.0105)	(3.54)	(0.204)	(1.070)	(0.834)
Observations	4,374	4,166	3,606	3,289	3,110
Mean Dependent Variable	0.016	-0.062	18.53	30.43	50.77
R-squared	0.403	0.314	0.675	0.541	0.455
Individual F.E.'s	Yes	Yes	Yes	Yes	Yes
Date F.E.'s	Yes	Yes	Yes	Yes	Yes
Manager Minority Status	No	No	No	No	No

*** significant at the 1% level; ** significant at the 5% level; * significant at the 10% level

Appendix Table 3. Effect of Manager Bias on Performance
Including Controls for Other Manager Characteristics

	Absence Indicator	Minutes Worked (Relative to Scheduled Minutes)	Articles per Minute	Inactivity Time (Seconds)	Checkout Time (Seconds)
			A. Store - Week Level		
Minority Worker ×	0.0386***	-5.96	-1.175**	4.439**	1.398
Manager Bias	(0.0137)	(4.09)	(0.549)	(2.036)	(1.364)
Manager Bias	-0.0190	2.53	0.216	-0.747	-1.069
	(0.0121)	(3.10)	(0.210)	(1.219)	(0.711)
Observations	4,374	4,166	3,606	3,289	3,110
Mean of Dep Var	0.016	-0.062	18.53	30.43	50.77
R-squared	0.406	0.315	0.675	0.545	0.457
Individual F.E.'s	Yes	Yes	Yes	Yes	Yes
Date F.E.'s	Yes	Yes	Yes	Yes	Yes

The regressions all control for manager minority status, manager level in the firm, manager age, and manager gender, and the interaction of these characteristics with the worker's minority status. Standard errors are clustered at the store level.
*** significant at the 1% level; ** significant at the 5% level; * significant at the 10% level

Appendix Table 4. Different Definitions of Minority Status

	Absence Indicator	Minutes Worked (Relative to Scheduled Minutes)	Articles per Minute	Inactivity Time (Seconds)	Checkout Time (Seconds)
		A. Minority = Workers of North African Origin (Workers of Subsaharan African Origin Excluded)			
Minority Worker ×	0.0393**	-5.824	-0.812*	4.124	1.987
Manager Bias	(0.0173)	(4.44)	(0.462)	(2.565)	(2.324)
Manager Bias	-0.0161	3.469	0.180	-1.370	-1.152
	(0.0115)	(2.79)	(0.176)	(1.466)	(0.847)
Observations	3,995	3,796	3,281	3,013	2,832
Mean of Dep Var	0.018	-0.892	18.52	29.00	50.57
R-squared	0.424	0.313	0.682	0.5495	0.464
		B. Minority = Workers of Subsaharan African Origin (Workers of North African Origin Excluded)			
Minority Worker ×	0.0350**	-17.318*	-1.330	3.591	-1.098
Manager Bias	(0.0162)	(9.31)	(0.789)	(2.309)	(1.772)
Manager Bias	-0.0128	2.830	0.124	-2.019	-1.388
	(0.0117)	(3.00)	(0.182)	(1.583)	(0.909)
Observations	3,564	3,383	2,911	2,672	2,526
Mean of Dep Var	0.020	-0.459	18.56	28.97	50.27
R-squared	0.443	0.327	0.678	0.565	0.472
	C. Minority = Workers of North or Subsaharan African Origin (Workers of Indeterminate, Mixed, or Other Origin Excluded)				
Minority Worker ×	0.0356**	-8.654**	-0.859**	3.709**	0.977
Manager Bias	(0.0138)	(4.22)	(0.318)	(1.704)	(1.389)
Manager Bias	-0.0153	2.28	0.049	-1.584	-1.497*
	(0.0134)	(2.93)	(0.199)	(1.361)	(0.809)
Observations	4,003	3,827	3,306	3,035	2,878
Mean of Dep Var	0.019	-0.366	18.61	29.29	50.96
R-squared	0.417	0.329	0.750	0.538	0.459
Individual F.E.'s	Yes	Yes	Yes	Yes	Yes
Date F.E.'s	Yes	Yes	Yes	Yes	Yes

Standard errors are clustered at the store level.

*** significant at the 1% level; ** significant at the 5% level; * significant at the 10% level

Appendix Table 5. Evidence on Mechanisms
Individual Survey Questions

	A. Monitoring & Effort				
	Which Manager You Had the Most Interaction with (On Shifts When Both Working)	Checked Up Most on Your Work	Noticed How Well or Poorly You Were Performing	Demanded You Work the Hardest	Under which Manager You Exerted Most Effort
Minority Worker ×	-0.061	-1.131***	-0.011	-0.527*	-1.029***
Manager Bias	(0.403)	(0.382)	(0.297)	(0.314)	(0.374)
Manager Bias	-0.311	-0.358	-0.517	-0.232	-0.085
	(0.406)	(0.418)	(0.437)	(0.341)	(0.445)
Minority Worker	-0.034	0.487**	0.073	-0.001	0.535**
	(0.230)	(0.226)	(0.212)	(0.180)	(0.238)
Observations	660	660	660	660	660
Mean of Dep Var	3.99	3.95	3.98	3.94	3.93
R-squared	0.312	0.316	0.316	0.322	0.293

	B. Management Choices			
	Best Management of Lines and Customer Flows	Assigned to Cleaning Duties Least Often	Assigned to Preferred Register Type	Assigned Best Breaks
Minority Worker ×	0.308	0.530	0.955*	0.308
Manager Bias	(0.624)	(0.804)	(0.561)	(0.424)
Manager Bias	-0.907*	-0.181	-1.217*	-1.173*
	(0.469)	(0.362)	(0.616)	(0.607)
Minority Worker	-0.071	-0.177	-0.445	-0.306
	(0.355)	(0.369)	(0.379)	(0.364)
Observations	397	397	397	397
Mean of Dep Var	3.88	3.94	3.84	3.89
R-squared	0.279	0.311	0.288	0.313

	C. Affection & Confidence				
	Most Likely to Recommend for Promotion	Manager Initially Made You Feel Most Confident	Manager Liked You Best	You Enjoyed Working with Manager Best	
Minority Worker ×	0.250	-0.190	-0.200	-0.183	
Manager Bias	(0.472)	(0.517)	(0.394)	(0.363)	
Manager Bias	-0.391	-0.429	-0.458	-0.699	
	(0.482)	(0.431)	(0.467)	(0.480)	
Minority Worker	0.014	0.022	0.152	0.123	
	(0.191)	(0.270)	(0.211)	(0.217)	
Observations	856	856	856	856	
Mean of Dep Var	3.92	3.92	3.89	3.88	
R-squared	0.293	0.313	0.295	0.307	
Individual F.E.'s	No	No	No	No	No
Store F.E.'s	Yes	Yes	Yes	Yes	Yes

Standard errors are clustered at the store level.
*** significant at the 1% level; ** significant at the 5% level; * significant at the 10% level

Appendix Table 6. Comparing Observations Early and Late in the Contract

	Early Weeks	Late Weeks	P-value of Difference
Minority Worker	0.273	0.292	0.764
Male Worker	0.056	0.078	0.533
Minority Manager	0.050	0.050	1.000
Male Manager	0.107	0.116	0.839
Level 4 Manager	0.182	0.165	0.736
Age of Manager	40.7	41.8	0.418
Average Manager IAT	0.451	0.495	0.389
Managers per Store			
Number of Stores	24	22	

Note: Cells in the first two columns report the mean of the indicated characteristics for the indicated sample. Standard errors are clustered at the store level.

*** significant at the 1% level; ** significant at the 5% level; * significant at the 10% level