UNIVERSITY OF CAMBRIDGE
Faculty of Economics

CAMBRIDGE WORKING PAPERS IN ECONOMICS

The Effect of Affirmative Action on Workers' Outcomes

Noriko Amano-Patiño Julian Aramburu Zara Contractor

Abstract

Fifty-six years after the introduction of affirmative action in employment in the U.S., there is a lack of consensus regarding the effect of this policy on workers' careers (Holzer and Neumark, 2000). This paper contributes to fill this gap by building and analyzing a dataset that allows us to quantify the effects of affirmative action in employment on workers' labor market outcomes. This paper circumvents prior data restrictions by constructing the first administrative database containing worker-level information (from the Longitudinal Employment Household Dynamics) as well as the federal contractor status of workers' employers (from the Equal Employment Opportunity Commission Data and Federal Procurement Data). We estimate the causal effects of affirmative action on workers' outcomes exploiting different features specified by the legal obligations of the regulation in a regression discontinuity setting.

Reference Details

CWPE	2104
Published	12 January 2021
Key Words	J15, J23, J31, J71, J78, K31
JEL Codes	Racial discrimination, affirmative action regulation, unemployment, earnings differentials
Website	www.econ.cam.ac.uk/cwpe

The Effect of Affirmative Action on Workers' Outcomes

Noriko Amano-Patiño* Julian Aramburu† Zara Contractor‡

We are bounded by law to not disclose details of our findings until they are officially approved by the U.S. Census Bureau. This version of our draft may not reflect the current state of the project.

Abstract

Fifty-six years after the introduction of affirmative action in employment in the U.S., there is a lack of consensus regarding the effect of this policy on workers' careers (Holzer and Neumark, 2000). This paper contributes to fill this gap by building and analyzing a dataset that allows us to quantify the effects of affirmative action in employment on workers' labor market outcomes. This paper circumvents prior data restrictions by constructing the first administrative database containing worker-level information (from the Longitudinal Employment Household Dynamics) as well as the federal contractor status of workers' employers (from the Equal Employment Opportunity Commission Data and Federal Procurement Data). We estimate the causal effects of affirmative action on workers' outcomes exploiting different features specified by the legal obligations of the regulation in a regression discontinuity setting.

JEL-codes: J15, J23, J31, J71, J78, K31.
Keywords: Racial discrimination, affirmative action regulation, unemployment, earnings-differentials.

Acknowledgments: We are grateful to Joseph Altonji, Jaime Arellano-Bover, Sebastian Heise, Costas Meghir, Javier Miranda, Giuseppe Moscarini and Asaf Paris Mandoki for helpful discussions. We also thank seminar participants at Yale University for insightful comments and suggestions.
This research uses data from the Census Bureau's Longitudinal Employer Household Dynamics Program, which was partially supported by the following National Science Foundation Grants SES-9978093, SES-0339191 and ITR-0427889; National Institute on Aging Grant AG018854; and grants from the Alfred P. Sloan Foundation.
*University of Cambridge, noriko.amanopatino@econ.cam.ac.uk.
†Yale University, julian.aramburu@yale.edu.
‡Yale University, zara.contractor@yale.edu.

1 Introduction

It has been fifty-five years since the passage of Executive Order 11246, the primary affirmative action policy in the U.S. labor market. To the best of our knowledge, there has been no study quantifying the effects of this regulation on workers' outcomes. This is because, until recently, there was no data connecting workers with the affirmative action regulation status of their employers.

One of the main contributions of this project is to create the first such dataset, by combining the Longitudinal Employer-Household Dynamics (LEHD), matched employer-employee data from unemployment insurance records, with data on firms' regulation status from the Equal Employment Opportunity Commission and the Federal Procurement Data System. The resulting dataset is a large panel that we use to analyze the short- and long-run effect of affirmative action on workers' outcomes; in particular, on employment and earnings.

While there has been prior work on the firm-level effects of affirmative action (notably, Miller (2017)), the literature has focused on the firms' share of minority workers. Our data allow us to analyze the effect of affirmative action regulation on a larger set of firm outcomes, such as relative earnings, diversity of new hires and turnover. Moreover, these data allow us to exploit a novel source of variation to identify the effects of affirmative action.

In addition, we will add valuable information regarding firms' perceptions of affirmative action regulation by implementing a survey. Finally, informed by our results and the survey's information, we will construct a theoretical model of discrimination in the labor market to understand the channels through which affirmative action may be operating to reduce discrimination.

Our contribution to the public policy dialogue is twofold. First, we document the effects that affirmative action has on wages and employment. This will shed light to the debate of whether the policy achieves its intended effects and reduces the institutional failures caused by discrimination. Second, based on our empirical findings, we will analyze the mechanisms through which affirmative action operates. These analyses seek to better understand the implications of public policies targeted at closing racial disparities, and may suggest important dimensions to consider when designing them.

In the remainder of this section, we lay out the institutional setting, then we describe our data and briefly outline our empirical strategy.

Institutional Setting

Executive Order 11246 prohibits firms holding federal contracts or subcontracts from discriminating in employment decisions on the basis of race, color, religion, sex, sexual orientation, gender identity or national origin. Such firms (*federal contractors*, hereafter) must include equal opportunity statements in job ads and posting relevant notices. The Department of Labor estimates that federal contractors employ about 25% of the U.S. labor force (Office of Federal Contract Compliance Programs, 2013).

Federal contractors with over 50 employees and a federal contract of at least 50,000 usd are also subject to affirmative action regulation. Affirmative action mandates regulated federal contractors and subcontractors to, in addition the above, create an affirmative action plan. Such plan identifies under-utilization of minorities and women in any occupation group relative to their availability in the local labor market.[1] As part of the plan, contractors are required to make "good faith" efforts to rectify under-utilization of minorities, including the use of numerical goals with timetables. The affirmative action plan must be submitted within 120 days of the start of the federal contract, and it must be updated annually for as long as the contract is in place. In order to monitor compliance, the Office of Federal Contract Compliance Programs audits the affirmative action plans of a small fraction of contractors each year.

2 Data Description

This project will construct and analyze, to the best of our knowledge, the first dataset containing worker-level information together with the regulation status of their employers. The dataset

[1] This is, "the availability of minorities having requisite skills in an area in which the regulated firm can reasonably recruit" (Office of Federal Contract Compliance Programs, 2013).

consists of a linkage between the following datasets of which we provide details below:

1) Establishment-level information on federal contracts from the Federal Procurement Data System (FPDS),

2) establishment-level data from EEO-1 reports provided by the U.S. Equal Employment Opportunity Commission (EEOC), and

3) the Longitudinal Employer-Household Dynamics (LEHD) covering the years 2001-2014, which defines the sample period of the project.

FPDS

The FPDS is a publicly available dataset containing the universe of federal contracts. For each contract in the years 2001-2014, we observe detailed information such as its dollar value; the duration of the contract; detailed establishment-level information (such as name, DUNS number, location, and a four-digit code describing the product or service); and the applicability of a variety of laws and statutes in addition to affirmative action.

EEOC

EEO-1 records contain the racial, ethnic and gender composition of firms with over 100 employees. The data is disaggregated by 9 broad occupation groups and separately available for establishments with over 100 employees.[2,3] Thus, in contrast with the FPDS, these data include information on all firms with over 100 employees, not only federal contractors.

LEHD

The LEHD is a quarterly panel dataset constructed from unemployment insurance records. The data contains earnings, employment status, state and county, together with demographic char-

[2]Compliance with Title VII of the Civil Rights Act of 1964, requires private sector firms with at least 100 employees to file EEO-1 reports, regardless of whether they hold a contract with the federal government to the EEOC. The employment size threshold above which firms must submit EEO-1 reports goes down to 50 employees if the firm holds a federal contract or sub-contract.

[3]The occupation categories consist of: officials and managers, professionals, technicians, sales workers, administrative support workers, craft workers, operatives, laborers/helpers, and service workers.

acteristics such as age, gender, race, place of birth, and citizenship status.We have access to 19 states[4] from 2000 to 2014. Details of how we process these data can be found in Appendix A. This is a highly restricted database that can only be accessed at a Research Data Center (RDC) of the U.S. Census Bureau.[5]

Finally, we will complement our data by designing and implementing a survey aimed at understanding employers' perceptions of affirmative action. Anecdotal evidence from the survey will complement our quantitative results and shed light on how firms put affirmative action into practice. This will ultimately inform the effects we find. The survey will be administered to a 1% representative sample of establishments (in terms of size and industry) with contracts above and below the threshold for affirmative action regulation in the U.S. Spending Data.

Constructing the dataset

A key challenge in building the novel data for this project is merging our data sources together. We have matched the EEO and LEHD as described below and are currently developing the linkage between the FPDS data and the LEHD. The challenge in merging the data together is twofold. First, the common identifier between the LEHD and the EEO-1 data (the EIN) is either missing or invalid for most of our observations in the EEO-1 sample (69.3% of them). Moreover, there is no common identifier between the LEHD and the FPDS.

Second, the size of the data requires the use of complex matching techniques. Thus, we develop an efficient fuzzy matching program that uses names and addresses to create a bridge between EEO-1 records and their corresponding counterpart in the LEHD. We do this in two steps: first, conditional on the county, we find the best match in terms of the name of the firm. Second, we keep only those cases that satisfy a minimum threshold in the quality of the match in terms of the address. We check the accuracy of our method by using the 30.7% of observations that have a valid common identifier in the LEHD and EEO-1 data. A similar strategy is being developed for the match between FPDS data and the LEHD. More details can be found in

[4] Arizona, Arkansas, Colorado, District of Columbia, Delaware, Hawaii, Indiana, Iowa, Kansas, Maine, Maryland, Montana, Nevada, New Mexico, Oklahoma, Pennsylvania, Tennessee, Washington and Wisconsin.

[5] See https://lehd.ces.census.gov/data/.

Appendix B

3 Empirical Strategy

We aim to identify the effects of affirmative action on outcomes at the worker and firm levels. The outcomes we consider at the worker level are earnings and subsequent employment status. At the firm level, we consider the fraction of minorities — both overall and in new hires, — relative wages and turnover.

Our main specification will employ a regression discontinuity design. The source of the discontinuity is the 50,000 usd threshold in the value of the federal contracts, above which a firm is subject to affirmative action regulation.

We restrict the sample to firms with at least 50 employees, since affirmative action does not apply to smaller firms.[6] The sample will also exclude firms that are perennial contractors, and firms with contracts higher than 10 million usd,[7] which are subject to pre-award compliance evaluations. We analyze the sample of men and women separately.

Identification

The identifying assumption is that observed differences in outcomes at the 50,000 usd threshold are driven only by assignment to treatment (i.e. to being regulated by affirmative action). Under this assumption, the causal effect of affirmative action on outcomes is captured by the jump in the dependent variable at the discontinuity. The identification assumption would be violated if there was another policy impacting outcomes that also takes effect at the 50,000 usd threshold. A thorough revision of the Federal Acquisition Regulation manual (GSA, 2019) indicates that there is no such policy.

Another threat to identification would arise if firms manipulated the size of the contract to avoid the regulation. The contract acquisition process requires firms to make bids in order

[6] Compliance of affirmative action is monitored if the firm has more than 50 employees and has a "sizeable" contract. We do not use the size of the firm as an additional dimension of discontinuity as several other regulations bind when a firm reaches the 50-employees-threshold — e.g. the Family and Medical Leave Act of 1993 (FMLA).

[7] This corresponds to the top 0.3 percent of the distribution of contract values.

to compete for a contract award from the federal government. Therefore, firms could have incentives to select pricing plans that are just below the threshold. This would be problematic if, for example, discriminating firms where minorities have worse outcomes were more likely to be just under the threshold than just over. In such a case, we would overestimate the impact of affirmative action due to selection. However, we find no evidence of firms' contracts bunching below the 50,000 usd threshold, which partially alleviates this concern. In effect, we observe bunching mostly *at* 50,000 usd, which is right above the regulation threshold. Since we also observe bunching of the same order of magnitude at all contract values that are multiples of 10,000, we conclude that this bunching is related to rounding in the value of the contract.

To further ascertain that our results are not driven by selection, we will compare firms just above and below the threshold on observables and past outcomes, and also run a 'donut' specification where we leave out observations sufficiently close to the threshold.

Firm-level outcomes

Firm level outcomes include the proportion of minorities, proportion of minorities among new hires and relative average earnings. We assume outcomes are locally linear in the contract amount and all control variables. The regression discontinuity specification to be estimated close to the cutoff is:

$$y_{jt} = \alpha_0 + \alpha_1 \mathbb{1}_{c_{jt} \geq \bar{c}} + \alpha_2 (c_{jt} - \bar{c}) \mathbb{1}_{c_{jt} < \bar{c}} + \alpha_3 (c_{jt} - \bar{c}) \mathbb{1}_{c_{jt} \geq \bar{c}} + \alpha_4 Z_{jt} + \varepsilon_{jt},$$

where y_{jt} is the outcome of firm j at time t; c_{jt} is the contract amount, $\bar{c}$ is the threshold for the affirmative action policy (i.e. $\bar{c} = 50{,}000$ usd), Z_{jt} is a vector of controls and ε_{jt} captures firm-time unobserved heterogeneity.

The main coefficient of interest is α_1, which represents the impact of being regulated on y_{jt}.

Worker-level outcomes

Worker level outcomes include their earnings, the likelihood that they remain employed at the firm, and the likelihood that they remain employed (at any firm). The latter two outcomes will be analyzed in subsequent periods.

In order to identify the causal effect of affirmative action on worker level outcomes, we first restrict the sample of workers to that of *incumbent* workers; that is, workers who are employed by the firm at the time it gets its first contract in our sample period. The rationale for this is that new hires may select into the firm based on whether it is regulated by affirmative action (Leibbrandt and List, 2018).

As we are interested in measuring the impact of affirmative action on reducing racial disparities, we analyze minorities and non-minorities separately.

As above, the specification is:

$$y_{ijt} = \beta_0^M + \beta_1^M \mathbb{1}_{c_{jt} \geq \bar{c}} + \beta_2^M (c_{jt} - \bar{c}) \mathbb{1}_{c_{jt} < \bar{c}} + \beta_3^M (c_{jt} - \bar{c}) \mathbb{1}_{c_{jt} \geq \bar{c}} + \beta_4^M Z_{ijt} + \varepsilon_{ijt},$$

where coefficients depend on M, the minority status of individual i. All other variables are defined as in the firm-level outcome specification, except that outcomes, controls and the residual are at individual-firm-time level.

Analogous to the specification for firm-level outcomes, the coefficient of interest is β_1^M.

References

GSA (2019) "Federal Acquisition Regulation Manual," General Services Administration Office, available at https://www.acquisition.gov/browse/index/far.

Holzer, Harry J. and David Neumark (2000) "Assessing Affirmative Action," *Journal of Economic Literature*, 38, 483–568.

Hyatt, Henry R., Erika McEntarfer, Kevin L. McKinney, Stephen Tibbets, and Doug Walton (2014) "Job-to-job (J2J) flows: New Labor Market Statistics from Linked Employer-Employee Data," *U.S. Census Bureau Center for Economic Studies Paper*.

Leibbrandt, Andreas and John List (2018) "Do Equal Employment Opportunity Statements Backfire? Evidence from a Natural Field Experiment on Job-Entry Decisions,"Technical Report Working Paper 25035, National Bureau of Economic Research.

Miller, Conrad (2017) "The Persistent Effect of Temporary Affirmative Action," *American Economic Journal: Applied Economics*, 9, 152–90.

Sorkin, Isaac (2018) "Ranking Firms Using Revealed Preference," *The Quarterly Journal of Economics*, 133, 1331–1393.

Appendix A. Details of the LEHD

The LEHD data consist of quarterly worker-level earnings submitted by employers for the administration of state unemployment insurance (UI) benefit programs, linked to establishment-level data collected for the Quarterly Census of Employment and Wages (QCEW) program. The result is incredibly rich administrative dataset which, nevertheless, has some limitations that require us to impose assumptions in order to define unemployment spells and job-to-job transitions. In this section, we describe in detail the assumptions we imposed to deal with these shortcomings.

A.1 Job-to-job Transitions

The data does not contain beginning and end dates of these employment spells. Because of this, we follow Hyatt et al. (2014) to define dominant jobs at the quarterly level.

We say that an individual i is employed at employer j in time t if the worker received positive wages w from that employer in quarter t. We say that an individual i is *beginning-of-quarter employed* at employer j in time t if the worker received positive wages from that employer in both t and $t-1$.

The *dominant beginning-of-quarter* employer in quarter t is the employer from which the worker had the highest earnings summing over quarter t and quarter $t-1$. This job is chosen from among the employers where the worker had positive earnings in both quarter t and quarter $t-1$.

We then consider transitions between employer j and employer k when we see an individual moving from the dominant beginning-of-quarter employer j in t, to the dominant beginning-of-quarter employer k in either $t+1$ or a subsequent quarter. This approach allows us to uniquely link the main job held on the first day of the quarter to the main job held at the start of the subsequent quarter. It does, however, have the obvious disadvantage of dropping job transitions between short duration jobs during the quarter. Thus this categorization restricts each worker to have only one job flow per quarter.

A.2 Unemployment

The LEHD does not contain unemployment indicators. Many papers using the LEHD rely on earnings tests to identify an individual as employed (Sorkin, 2018). However, since unemployment is one of the primary outcomes of interest in this paper, we depart from this approach.

A drawback of the Employment History File (EHF) is that it contains information for the 19 states for which our project was approved.[8] If a worker in one of our approved states was employed in a state we do not have access to, she will appear to have no earnings — and thus would be typically classified as unemployed.

We circumvent this issue by using the EHF National Indicator File from the LEHD, which provides information about the presence of wage records for workers in any state in the U.S.

We say an individual is employed if she has wage records in any state in the U.S. regardless of the level of earnings — the EHF National Indicator File does not contain information about the level of earnings.

Appendix B. Details of the Matching Procedure

A key challenge in building the novel data for this project is merging our data sources together as there is no reliable common variable across them except for establishment name.

In order to overcome this challenge, we develop a fuzzy matching method based on term frequency–inverse document frequency (tf-idf) scores of both names and addresses of each establishment.

The tf-idf score of a word is a numerical statistic that reflects its importance relative to a document within a corpus (a collection of documents). In short, the tf–idf value increases proportionally to the number of times a word appears in the document and is offset by the number of documents in the corpus that contain the word, which helps to adjust for the fact that some words appear more frequently in general.[9]

[8] Arizona, Arkansas, Colorado, District of Columbia, Delaware, Hawaii, Indiana, Iowa, Kansas, Maine, Maryland, Montana, Nevada, New Mexico, Oklahoma, Pennsylvania, Tennessee, Washington and Wisconsin.

[9] An alternative to using tf-idf scores would be to measure the distance between strings of words with measures such as the Levenshtein distance or the Jaro-Winkler distance. These methods however, are inefficient when dealing

We split the U.S. into 3-digit zipcode regions and within each of these regions we proceed as follows. First, we consider the names of establishments and split these into n-grams of size

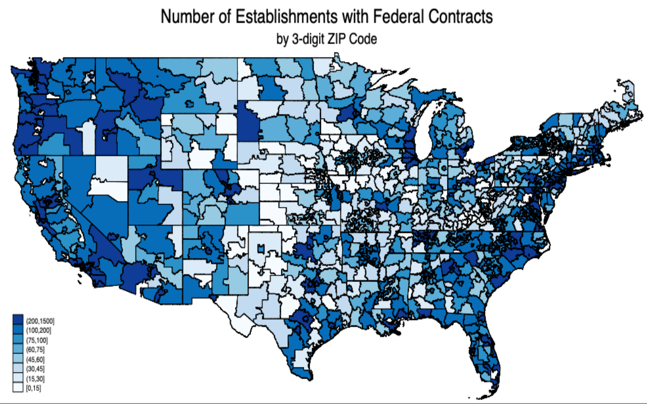

Figure B.1: *Three digit zipcode areas.*

3.[10]

As a reference, we use the establishments' names in the Name and Address File from the SSL at the Census and refer to these data as the *names' training data* in the rest of this section. As a counterpart, we refer to the establishments' names in the FPDS Data as our *names' test or query data*. We map each of the documents (establishments' names) in the names' training data into a vector of the tf-idf scores of the document's n-grams. The tf-idf scores are normalized so that they have norm 1 and they belong to $\mathbb{R}^K$ where K denotes the total number of distinct n-grams in the training corpus.

We then proceed to map the documents in the names' query data into vectors of tf-idf scores using as our reference the set of distinct n-grams in the training corpus. All documents are mapped into a vector in the unit sphere in $\mathbb{R}^K$. If any of the documents in the names' query data have no n-grams in common with the documents in the training data, the document is mapped into a vector of zeros, $0^K \in \mathbb{R}^K$. Otherwise, vectors have norm 1.

In order to match the names' query data with the names' training data, we compute the *cosine similarity scores*[11] between each possible pair of vectors. If the score between a pair is

with large data as the time required to implement them grows quadratically with the number of observations.

[10]That is, we split the name of each establishment into 3-character pieces. As an example, the n-grams of the word 'Department', would be ' De', 'Dep', 'epa', 'par', 'art', 'rtm', 'tme', 'men', 'ent', 'nt '.

[11]The cosine similarity score between vectors $x, y \in \mathbb{R}^K$ is given by $x \cdot y / \|x\| \|y\| = \cos(x, y)$. That is, it measures

"high enough" — meaning the angle between them is "small enough," — we proceed to analyze the establishments' addresses. Otherwise we assume the pair is a bad match.

For the subset of names that are a potential match (i.e. those that have a cosine similarity score above the threshold), we analyze the similarity of establishments' addresses between the Name and Address File from the SSL at the Census and the FPDS data in a similar fashion.

the angle between x and y, where the angle (with range $(0,\pi)$) is mapped into the $(0,1)$ interval. If the angle between x and y exceeds π, then we measure its negative which will be in the $(0,\pi)$ interval.