



HAL
open science

Essays in labor economics : discrimination, productivity and matching

Dylan Glover

► **To cite this version:**

Dylan Glover. Essays in labor economics : discrimination, productivity and matching. Economics and Finance. Institut d'études politiques de paris - Sciences Po, 2017. English. NNT : 2017IEPP0025 . tel-03436530

HAL Id: tel-03436530

<https://theses.hal.science/tel-03436530>

Submitted on 19 Nov 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Institut d'études politiques de Paris
ÉCOLE DOCTORALE DE SCIENCES PO
Doctoral program in Economics
Department of Economics
Doctor of Philosophy in Economics

**ESSAYS IN LABOR ECONOMICS:
DISCRIMINATION, PRODUCTIVITY AND
MATCHING**

Dylan Glover

Thesis supervised by Yann Algan

Defended on December 11th, 2017

Jury:

Yann Algan, Professeur des universités, IEP de Paris

Ghazala Azmat, Professeur des universités, IEP de Paris

Pierre Cahuc, Professeur des universités, Ecole Polytechnique (Rapporteur)

Bruno Crépon, HDR, Administrateur INSEE, CREST-INSEE

Esther Duflo, Professor of Economics, MIT (Rapporteur)

Acknowledgments

First and foremost I am indebted to my life companion, Johana Carrier, for her love, encouragement, understanding, flexibility and insight throughout the trials and tribulations of this dissertation. Without her, I would never have been able to complete it.

I thank my brilliant parents, Cathleen and Gregory. They inspired, and though too far away nowadays, still inspire me to be interested in thinking critically about the world and its injustices. The ideas presented in this thesis have their roots in themes they planted in me throughout my childhood and adolescence. And born in the midst this adventure, I am so grateful for my two children, Loé and Mina. They provided perspective and levity in the moments I needed it the most.

I am indebted to my advisors Yann Algan and Bruno Crépon who believed enough in me to give me the opportunity to pursue my Ph.D. and secondly, for imparting their knowledge, experience and wisdom upon me. I have greatly benefited from all my exchanges with them during the last four years.

A deep thank you to Elizabeth Beasley, who took a chance on me way back when I was searching for something worthwhile to do in my life, setting in motion the events that would eventually lead to this body of work. Thank you to Amanda Pallais and William Parienté for their encouragement and apprenticeship bestowed upon me during the 6 years it took for the first chapter to come to fruition. They were unfazed by my pre-economist ramblings and taught me so much, inspiring me to pursue a Ph.D.

I am indebted to the whole team in the DSEE at Pôle emploi for having warmly welcomed me, allowing me to pursue my Ph.D. in excellent conditions. Most notably I would like to thank François Aventur, Anita Bonnet, Stéphane Ducatez and Cyril Nouveau. Without their confidence and support this thesis would never have been realized. I also wholeheartedly recognize the support, both intellectually and financially from the Chaire de Sécurisation

des Parcours Professionnels and its members.

I present my gratitude to the professors, staff and fellow Ph.D. students of the Economics Department at SciencesPo and to J-PAL. Lastly, I have greatly benefited from exchanges with Anil Bhargava, Claire Bernard, Guillaume Blache, Adrien Bouguen, Alexis Bowen, Thomas Brand, David Buchner, Philippe Carrier, Maguelone Cathala, Frédéric Cochinard, Pierre Deschamps, Assia Elgouacem, Alexis Gagnon, Yannick Galliot, Elise Huillery, Olivier Jugand, Jean-Louis Keene, Lucie Lalande, Claudine Lamaze, Murielle Matus, Mawuli Mis-sebukpo, Adrien Pawlik, Elise Pesonel, Régine Petit, Arnaud Philippe, Victor Pouliquen, Alexia Pretari, Marie-José Rabner, Eric Renard, Alexandra Roulet and Juliette Seban along this journey.

Institut d'études politiques de Paris
ÉCOLE DOCTORALE DE SCIENCES PO
Programme doctoral en économie
Département d'économie
Doctorat en sciences économiques

**ESSAIS SUR L'ÉCONOMIE DU TRAVAIL :
DISCRIMINATION, PRODUCTIVITÉ ET
APPARIEMENT**

Dylan GLOVER

Thèse dirigée par Yann ALGAN

Composition du Jury :

Yann ALGAN, Professeur des universités, IEP de Paris

Ghazala AZMAT, Professeur des universités, IEP de Paris

Pierre CAHUC, Professeur des universités, École Polytechnique (Rapporteur)

Bruno CRÉPON, HDR, Administrateur INSEE, CREST-INSEE

Esther DUFLO, Professor of Economics, MIT (Rapporteur)

Soutenue le 11 décembre 2017

Contents

Résumé long (en français)	10
I Discrimination as a Self-Fulfilling Prophecy: Evidence from French Grocery Stores	24
1 Introduction	25
2 Setting	29
3 Data and Descriptive Statistics	31
3.1 Data Sources	31
3.1.1 Administrative Store Data	31
3.1.2 Manager Survey and IAT	33
3.1.3 Worker Survey	35
3.2 Descriptive Statistics	36
3.3 Exogeneity	37
4 Effect of Manager Bias on Performance	39
4.1 Time Spent at Work	39
4.2 Performance while at Work	41
5 Evidence on Mechanisms	45
5.1 Effects of Manager Bias on Minority Performance	45
5.1.1 Theories	45
5.1.2 Evidence	46
5.2 Hiring Decisions	48
5.2.1 Theories	48
5.2.2 Evidence	49
6 Conclusion	50
Tables and figures	52

Appendix tables and figures	62
II The Value of a Vacancy: Evidence from a Randomized Evaluation with Local Employment Agencies in France	76
1 Introduction	77
2 Description of intervention and heterogeneity dimension	79
2.1 Description of intervention	79
2.2 Heterogeneity	80
3 Data, sampling and randomization	81
3.1 Data	81
3.2 Sampling and Randomization	82
3.3 Empirical Specification	83
4 Balance, sample description and compliance	84
5 Compliance and treatment intensity	85
6 Impacts	86
6.1 Vacancy and hiring flows	86
6.2 Employment creation	87
6.3 Distribution of impacts	89
6.4 Simply intertemporal substitution?	90
6.5 Displacement?	90
7 Potential Mechanisms	92
7.1 Outline	94
7.2 Setup	94
7.3 Predictions	96
7.3.1 Number and types of vacancies posted	97
7.3.2 Number of applications	99
7.3.3 Number of hires	100
8 Conclusion	102
Figures and Tables	103

Appendix Figures and Tables	112
Variance Computation Appendix	123
Proof Appendix	123
III Job Search and Intermediation under Discrimination: Evidence from Terrorist Attacks in France	126
1 Introduction	127
2 Data	130
2.1 Distribution statistics	133
3 Identification and empirical specification	134
4 Impacts	138
4.1 Potential matches	138
4.2 Impacts on hires	142
4.3 Robustness	143
5 Heterogeneity in impacts	147
5.1 Theoretical motivation	147
5.2 Empirical results	149
5.3 Compensatory effects	154
5.4 Discussion	157
6 Conclusion	158
Figures and Tables	160
Appendix	178
Weighted Least Squares regression	178
A synthetic placebo year	178
Appendix Figures and Tables	182
Bibliography	190

Résumé long

Cette thèse examine la manière dont la discrimination influence la recherche d'emploi et le rendement au travail des minorités. Elle explore également comment les politiques actives du marché du travail visant à aider les entreprises dans leur processus de recrutement affectent leur offre d'emploi.

Dans le premier chapitre, co-écrit avec Amanda Pallais et William Parienté, nous étudions la façon dont le rendement au travail des minorités change selon qu'elles travaillent avec des managers qui ont plus ou moins de préjugés à l'encontre de leurs origines. Ce biais est mesuré par un test d'association implicite (TAI). Nous créons une mesure de la probabilité de faire partie de la minorité ou de la majorité d'après leurs noms étymologiquement arabes ou français. En exploitant l'assignation quasi-aléatoire des caissiers d'une grande chaîne de supermarchés en France, nous montrons que lorsque les minorités travaillent avec des managers plus biaisés, leurs performances sont nettement moins bonnes que celles des employés issus de la majorité, et ce sur une série d'indicateurs de performance : leur absentéisme est plus élevé et ils sont moins enclins à faire du temps de travail additionnel à la fin de leur journée de travail. Au travail, ils scannent les articles plus lentement et mettent plus de temps à passer d'un client à un autre. Nous présentons des résultats qui soutiennent le fait que les managers biaisés ont tendance à moins interagir avec les employés issus des minorités et que ces employés font alors moins d'effort. Nous constatons qu'en moyenne, les performances des employés issus des minorités et de la majorité sont équivalentes lorsqu'ils travaillent avec des managers non biaisés, ce qui suggère que les minorités peuvent être mieux qualifiées pour un même salaire. Nous affirmons que c'est un fait empirique qui révèle une prophétie auto-réalisatrice selon laquelle les managers biaisés rendent les minorités moins productives et cela génère une discrimination statistique dans la politique d'embauche de l'entreprise.

Le deuxième chapitre explore comment changer la valeur d'un poste vacant en offrant des services de recrutement gratuits aux entreprises a une incidence sur son offre d'emploi. Les co-auteurs, Yann Algan, Bruno Crépon et moi-même, constatons que lorsque le service public de l'emploi (SPE), Pôle emploi, sélectionne au hasard des entreprises pour déterminer leurs besoins de recrutement et offrir des services de recrutement gratuits, cela augmente fortement le nombre d'offres d'emploi postées sur Pôle emploi. En outre, cela se traduit par une augmentation significative des embauches en contrats permanents. Nous constatons l'effet sur l'affichage des postes vacants en utilisant l'ensemble de notre échantillon, mais l'effet sur les recrutements est exclusivement centré sur les entreprises qui étaient en contact avec le SPE au cours de l'année précédant l'intervention. Pour approfondir la compréhension

de ces effets hétérogènes, nous développons un modèle dans lequel les entreprises ont la possibilité de recruter par différents canaux, tels que les candidats qui passent par le SPE, par l'intermédiaire d'un conseiller SPE ou qui sont autonomes dans leur recherche. Le modèle prédit que lorsque le SPE fournit des services de présélection supplémentaires, les entreprises utilisent les canaux SPE, ce qui augmente sans ambiguïté les affichages des postes vacants, mais a des effets ambigus sur l'embauche. Cela est dû au fait que les employeurs réduisent leurs propres efforts, mais deviennent également plus exigeants sur les personnes qu'ils recrutent ; lorsque les canaux SPE sont ouverts, l'entreprise est exposée à une nouvelle distribution des compétences des candidats qui peuvent être plus ou moins adaptées aux besoins des entreprises. En testant le modèle à l'aide de données sur les appariements potentiels avec les offres d'emploi, nous trouvons des preuves à l'appui des prédictions du modèle. Un nombre significativement plus élevé de services de présélection ont été fournis aux entreprises du groupe de traitement par rapport aux entreprises du groupe de contrôle, et ces entreprises consacrent moins d'efforts à leurs propres recherches. Les résultats de cette étude suggèrent que les politiques actives du marché du travail visant à générer la demande de main-d'œuvre des entreprises peuvent avoir une valeur ajoutée substantielle sur le marché du travail.

Dans le dernier chapitre, j'explore les effets sur les démarches de recherche d'emploi des minorités musulmanes en France suite aux attentats terroristes de janvier 2015 contre Charlie Hebdo, la police et le supermarché HyperCacher. Je le fais en examinant les efforts de trois acteurs clés du marché du travail : les demandeurs d'emploi, leurs conseillers SPE et les employeurs. Comme dans le premier chapitre, les demandeurs d'emploi issus des minorités et de la majorité sont catégorisés en utilisant leurs prénoms. Les résultats montrent que les demandeurs d'emploi issus des minorités réduisent nettement leurs efforts de recherche d'emploi par rapport à ceux issus de la majorité dans les semaines qui suivent les attaques. Les employeurs réduisent également leurs efforts de recherche, mais seulement pour les types de contrats les plus prisés, c'est-à-dire les contrats à durée indéterminée, tout en augmentant leur recherche de candidats issus des minorités pour des postes en contrat à durée déterminée et en intérim. Cette baisse est partiellement compensée par une augmentation des démarches des conseillers après le choc, mais seulement dans les régions où la discrimination latente, mesurée grâce à la part des votes locaux pour le Front national, est relativement faible. De plus, cet effet de conseiller est beaucoup plus fort pour les conseillers qui sont eux-mêmes issus des minorités et pour les conseillers issus de la majorité qui se spécialisent dans le retour au travail des demandeurs d'emploi les plus marginalisés. Je trouve également des impacts sur les types de contrats proposés aux minorités, mais il est difficile d'en faire une interprétation précise en raison des effets d'équilibre général potentiellement importants provoqués par les

changements d'intensité de recherche. Il est donc difficile de démêler l'effet d'un changement de comportement de recherche des demandeurs d'emploi issus des minorités par rapport à un changement dans la discrimination des employeurs. En examinant la création totale d'emploi, je trouve peu de résultats soutenant que les demandeurs d'emploi issus des minorités ont trouvé moins de travail que ceux issus de la majorité dans les semaines qui ont suivi les attaques. Cela suggère que les intermédiaires du marché du travail peuvent jouer un rôle clé dans l'atténuation des effets négatifs de la discrimination.

Chapitre 1. La discrimination comme prophétie auto-réalisatrice : des résultats empiriques provenant des supermarchés français

Ce chapitre présente des résultats qui mettent en valeur le fait que les croyances discriminatoires peuvent affecter directement les performances professionnelles des employés issus des minorités dans un environnement de travail réel. Une vaste littérature économique teste l'existence de la discrimination dans le marché du travail. La discrimination étant présente quand une personne issue d'une minorité et une autre issue de la majorité faisant preuve de la même productivité sont traitées différemment en termes d'embauche, de rémunération ou de promotion. Le travail pionnier de Becker, *The Economics of Discrimination* (1957), introduit la notion de discrimination de goût : les employeurs éprouvent de la désutilité lorsqu'ils emploient une personne issue d'une minorité et ils compensent en versant des salaires moins élevés ou en les obligeant à être plus productifs pour le même salaire. Des recherches postérieures, en commençant par Phelps (1972), puis Arrow (1973), ont envisagé la discrimination non pas comme de l'animosité, mais plutôt comme de l'information imparfaite. Les a priori défavorables sur la productivité des minorités ou de la présélection imparfaite amènent les employeurs à traiter de manière inégale les employés issus des minorités et de la majorité bien qu'ils soient également qualifiés. Lundberg et Startz (1983) et Coate et Loury (1993) ont élaboré à partir de ces idées en montrant comment la discrimination statistique pourrait potentiellement réduire les investissements des minorités dans les compétences en les menant à croire, à raison, que ces investissements ne seraient pas pleinement récompensés. En conséquence, la discrimination statistique peut conduire à une prophétie auto-réalisatrice selon laquelle les croyances antérieures défavorables des employeurs portant sur les niveaux de compétences des minorités s'auto-confirment en équilibre. L'hypothèse implicite que les goûts et les croyances des employeurs n'ont pas d'incidence directe sur la productivité des employés est partagée par cette littérature. Bien que la discrimination statistique puisse inhiber l'investissement dans les compétences, elle n'affecte pas directement la performance

des employés. Cependant, Steele et Aronson (1995) documentent que les stéréotypes négatifs sur les groupes issus des minorités peuvent affecter directement les performances. Cet ensemble de travaux démontre qu'aux États-Unis, lorsque les stéréotypes sont mis en valeur avant l'exécution (par exemple, les candidats sont priés de déclarer leur race ou leur sexe), les noirs, les hispanophones et les femmes font preuve de performances nettement moins bonnes que dans les cas où l'appartenance à un groupe n'est pas connue. Cette ligne de recherche qui a eu, jusqu'à présent, une influence limitée sur la littérature économique, implique que les croyances négatives des employeurs à l'égard des minorités – qu'elles soient issues d'une discrimination de goût ou bien statistique – pourrait être auto-réalisatrices. Non pas parce qu'ils inhibent l'investissement dans les compétences, mais parce qu'ils induisent une performance dégradée.

Nous étudions 34 hypermarchés d'une chaîne de grande distribution en France. Ces magasins emploient une proportion considérable de salariés que nous identifions, d'après leurs prénoms, comme ayant une origine maghrébine ou africaine subsaharienne¹. Nous étudions les caissiers en contrat de professionnalisation embauchés pour une durée de six mois, car ces employés sont affectés quasi-aléatoirement aux managers. Ces caissiers, comme tous les caissiers des magasins, travaillent avec des managers différents à des jours différents. Ils ne sont pas autorisés à soumettre des préférences de planification. Leurs plannings sont déterminés par un programme informatique qui attribue des horaires selon la demande clientèle prévue, en tenant compte des préférences des employés seniors. Ainsi, les employés issus des minorités et de la majorité de notre échantillon ne choisissent pas les managers avec lesquels ils travaillent d'une part, et, d'autre part, travaillent avec les mêmes managers dans des conditions similaires ; enfin les employés savent au préalable avec quels managers ils vont travailler.

Nous avons mesuré les a priori des managers envers les minorités avec un test d'association implicite, qui est fréquemment utilisé pour mesurer les biais, en particulier dans le domaine de la psychologie (Voir Nosek, Greenwald et Banaji 2007 ; Lane et al. 2007 ; et Greenwald et al. 2009 pour les résumés de la littérature). Les scores TAI sont corrélés avec de nombreuses décisions du monde réel et sont difficiles à manipuler².

Le test utilise la vitesse à laquelle les sujets catégorisent les items pour déterminer leur association implicite entre deux concepts : ici (1), les prénoms qui sonnent comme ayant une origine française, maghrébine ou africaine subsaharienne et (2), des mots liés aux compétences ou incompétences au travail. Le score d'a priori mesure ainsi l'étendue

¹Les employés sont classés en fonction de leur noms. ISM CORUM, expert dans le testing de discrimination en France, a fait la catégorisation.

²Voir, par exemple, Kim (2003) , Friese, Bluemke et Wänke (2007) , Green et al. (2007) , Greenwald et al. (2009) , et Rooth (2010) .

de l'association que font les managers entre noms maghrébins ou africains et performance moindre. Les magasins nous ont fourni des données sur les absences et le temps travaillé (déterminés par les données de pointage), la vitesse à scanner les articles et le temps passé entre chaque client. Nous évaluons si les performances des caissiers issus des minorités sont moins élevées les jours où ils travaillent avec les managers qui ont des préjugés contre leur origine. Nous utilisons la méthode des doubles différences, en comparant l'évolution de la performance des employés issus des minorités travaillant sous des managers plus ou moins biaisés à la performance des employés issus de la majorité. Nous constatons que les minorités sont plus susceptibles d'être absents lorsqu'il est prévu que ces salariés travaillent avec des managers ayant plus d'a priori contre leur type. Quand ils se trouvent sur leur lieu de travail, ils passent moins de temps au magasin : spécifiquement, ils sont beaucoup moins susceptibles de rester après la fin de leur temps de travail. Parce que les employés sont payés en fonction du temps travaillé, nous estimons que les minorités gagnent 2,5 % de moins en raison des a priori des managers. De plus, les minorités scannent les articles plus lentement et passent plus de temps entre les clients lorsqu'ils travaillent avec des managers biaisés.

Nous combinons les données d'une enquête auprès des employés réalisée après la fin du contrat avec nos données administratives pour distinguer quelles théories de discrimination expliquent le mieux nos résultats. Tout d'abord, nous trouvons peu de preuves que les managers traitent mal les minorités. Les employés issus des minorités ne signalent pas que ces managers biaisés les détestent ou qu'ils détestent eux-même les managers biaisés. Ils signalent que les managers biaisés sont *moins* susceptibles de les affecter à des tâches désagréables (nettoyage) et ne sont pas plus enclins à les attribuer aux caisses désagréables ou à leurs donner des temps de pauses moins pratiques. Nos résultats sont plus cohérents avec une théorie dans laquelle les managers biaisés interagissent moins avec les employés issus des minorités. La recherche en psychologie sur le « racisme aversif » a montré que les individus avec des préjugés implicites envers les minorités sont moins susceptibles de parler, plus hésitants à parler et moins amicaux envers les membres de ces groupes³. Ils peuvent se sentir moins à l'aise d'interagir avec les personnes issues des minorités ou ils peuvent être préoccupés par le fait de laisser transparaître leurs a priori. En effet, nous trouvons que les managers biaisés interagissent moins avec les minorités est ceci est compatible avec le fait qu'il sont moins susceptibles de demander aux minorités de faire le ménage ni de faire du temps de travail additionnel à la fin de leur journée de travail.

En faisant baisser la productivité des minorités, la discrimination peut conduire l'entreprise à agir comme si, bien qu'ayant les mêmes caractéristiques, les salariés issus des minorités

³Voir, par exemple, McConnell et Leibold (2001), Dovidio, Kawakami et Gaertner (2002) et Hebl et al. (2002). Dovidio et Gaertner (2008) font un résumé de cette littérature.

étaient moins productifs que ceux issus de la majorité. Parce que l'entreprise doit offrir le même salaire à ses employés, indépendamment de leurs performances, elle fixerait un seuil d'embauche plus élevé pour les minorités. Ainsi, en l'absence de discrimination de la part des managers, les personnes issues des minorités embauchées devraient être plus productives que celles issues des majorités embauchées. Dans l'ensemble, nous trouvons que les minorités et la majorité se comportent de manière équivalente. Cependant, conformément à la discrimination statistique, lorsqu'ils travaillent avec des managers sans a priori, les employés issus des minorités sont nettement plus productifs que ceux issus de la majorité. Ils sont deux fois moins susceptibles d'être absents et ils scannent plus rapidement les articles ; quand ils travaillent avec ces managers impartiaux, les minorités servent 9 % de clients en plus. Alors que la moyenne d'une personne issue d'une minorité se situe au 53e centile de la moyenne de performance en général, sa performance est au 79e centile les jours travaillés avec des managers sans préjugé,. Cela suggère que la discrimination a des conséquences importantes sur la performance des employés et donc, en théorie, sur leurs résultats sur le marché du travail.

Chapitre 2. La valeur d'une offre : une évaluation randomisée avec les agences locales de Pôle emploi

Le modèle d'appariement du marché du travail de Mortensen and Pissarides (1994) définit explicitement les coûts de recrutement comme un paramètre clé qui déterminent l'offre d'emploi. Pourtant, les politiques actives du marché du travail et les études qui examinent leurs effets se sont concentrées presque exclusivement sur la relation entre l'offre d'emploi d'une entreprise, la productivité et les coûts attendus après l'embauche. En effet, de nombreuses études ont exploré les impacts de l'aide aux demandeurs d'emploi à travers des programmes de formation ou dans leur recherche d'emploi, et il existe des preuves solides que ces programmes peuvent être efficaces (Card et al., 2015). Dans cet article, nous explorons l'idée qu'une intervention symétrique qui aide les entreprises dans leur recrutement pourrait également apporter une valeur ajoutée sur le marché du travail.

Les grands thèmes de l'appariement et de la recherche coûteuse face à l'information asymétrique de l'employeur et de l'employé forment la base de la littérature économique, mais beaucoup moins d'attention a été accordée à la façon dont les employeurs abordent le processus de recrutement. Par exemple, Manning (2011) étudie la littérature et montre que les estimations des coûts de recrutement, par rapport à la masse salariale, varient de 1,5 %

à 11 % selon le contexte et le type d'emploi⁴. Dans le processus de recrutement, les coûts monétaires de poster une offre d'emploi par le biais du service public de l'emploi peuvent être très faibles, comme Dickens et al. (2001) le soulignent, bien que Manning (2011) souligne que la majeure partie du coût ne vient pas de la génération de candidats, mais plutôt de la sélection, puis de la formation⁵. Ayant consacré presque exclusivement des ressources à l'aide aux demandeurs d'emploi depuis 2008, le service public de l'emploi (SPE) en France a réorganisé son programme de services aux entreprises en 2015⁶. Ce nouveau service aux entreprises reposait sur un traitement plus intensif et dynamique des offres d'emploi, notamment par l'introduction de services de présélection, dans le but d'aider les entreprises de manière plus efficace à trouver le bon candidat pour le poste. La prestation de ces nouveaux services devait être complétée par la prospection : les conseillers des agences devaient étudier activement les besoins des entreprises et les contacter de manière proactive pour proposer les nouveaux services de recrutement afin d'améliorer le placement des demandeurs d'emploi⁷. Nous étudions les impacts de ces nouveaux services de recrutement en tirant au sort les entreprises prospectées. 8 232 entreprises ont participé à l'étude et la moitié d'entre elles ont été prospectées intensivement au cours d'une fenêtre de 10 semaines à la fin 2014. Les conseillers ont ensuite été invités à incuber et maintenir des relations avec les entreprises prospectées jusqu'en avril 2015. Pendant ce temps, aucune action proactive n'a été prise envers le groupe de contrôle. Au cours de cette période de sanctuarisation, les entreprises du groupe de contrôle ne se sont pas vu refuser le service si elles l'ont demandé.

En imposant de la variation aléatoire dans les coûts de recrutement attendus des entreprises, nous essayons donc de « choquer » la valeur d'une offre pour l'entreprise afin d'examiner comment ce changement de sa valeur affecte l'offre d'emploi. Cet éloignement d'une focalisation exclusive sur les demandeurs d'emploi est concomitant du récent travail sur le potentiel d'effets de déplacement qui peut nier l'efficacité des politiques actives du marché du travail visant à améliorer la qualité des demandeurs d'emploi (Crépon et al., 2013). La

⁴Le tableau 2 de Manning (2011) résume les pourcentages tirés de huit études distinctes dans plusieurs pays différents et souligne que nous ne pouvons pas déterminer si ces coûts de recrutement sont marginaux ou moyens. Il suggère que la majeure partie du coût est liée à la formation de nouveaux employés, mais que le coût du processus d'embauche lui-même demeure une question importante.

⁵En mai 2017, poster une offre d'emploi pendant 30 jours sur Monster.fr coûte 855 euros. Ce coût s'élève à 1 390 euros pour un affichage complet de 60 jours. Il n'est donc pas évident que les coûts de flux soient négligeables dans ce contexte.

⁶Pôle emploi compte plus d'un milliard d'agences en France métropolitaine et en territoires d'outre-mer. Il a été créé en 2008 à la suite d'une fusion entre l'agence nationale pour l'emploi (ANPE), l'agence gouvernementale de conseil en emploi et de recrutement, et l'assedic (Association pour l'emploi dans l'industrie et le commerce).

⁷Dans des entretiens qualitatifs menés dans la phase de faisabilité de l'étude en début 2014, les conseillers qui avaient précédemment travaillé pour l'ANPE ont souligné de manière anecdotique que l'absence de relations avec les entreprises depuis la fusion était une perte pour les opportunités de placement.

motivation de cette étude est donc aussi de tester l'efficacité d'une politique innovante qui pourrait éviter ces effets de déplacement entre les demandeurs d'emploi.

Nous constatons que ce programme de prospection a entraîné une hausse de 30 % des offres d'emploi postées auprès du SPE et, plus important encore, une augmentation de 9 % des embauches en CDI, ce qui se traduit par 48 jours de travail supplémentaires créés par les entreprises du groupe de traitement, en moyenne. Nous trouvons des impacts sur les marges intensives et extensives, les entreprises prospectées étant beaucoup plus susceptibles d'embaucher un demandeur d'emploi inscrit au SPE. Nous sommes en mesure d'exclure que les effets positifs importants sur la création d'emplois et les embauches sont motivés par la substitution intertemporelle, c'est-à-dire une accélération du processus de recrutement. Nous affirmons également qu'il est peu probable que ces effets soient atténués par des « effets de déplacement d'entreprise », un cas dans lequel l'intervention ayant pu amener les employés des entreprises du groupe témoin à passer de façon disproportionnée dans les emplois créés dans le groupe de traitement. Nous ne croyons pas non plus que les entreprises de traitement ont probablement déplacé les processus de recrutement des entreprises de contrôle ou des entreprises hors de notre échantillon. Nous affirmons que parce que l'intervention a eu lieu dans des marchés du travail à faible tension, le bassin de demandeurs d'emploi disponibles est important et donc la concurrence entre les candidats est faible. De plus, les entreprises de l'échantillon ne représentent qu'une très faible proportion des entreprises locales, de sorte que les effets d'équilibre général liés à la mise en œuvre du programme sont susceptibles d'être faibles.

Enfin, alors que nous trouvons des impacts forts sur la création d'offres d'emploi avec l'ensemble des entreprises de l'échantillon, nous constatons que les impacts sur les embauches sont exclusivement centrés sur les entreprises ayant été en contact avec le SPE au cours des neuf mois précédant l'intervention. Enfin, pour structurer notre compréhension des résultats, nous présentons un modèle de recrutement multicanal dans lequel la valorisation d'une offre par une entreprise repose non seulement sur les marges de rentabilité par les coûts salariaux et les flux des coûts de poster une offre, mais aussi sur le processus de sélection coûteux disponible dans chaque canal. Bien que suggestif, nous croyons qu'il donne un aperçu de la façon dont la demande de main-d'œuvre d'une entreprise peut changer en raison d'un choc sur les marges qui ne s'adresse pas à la profitabilité de l'emploi créé. Bien que non expérimental, nous trouvons des preuves empiriques pour étayer les prédictions du modèle. Nous ne trouvons aucune différence dans la marge de la rentabilité salariale entre les offres des groupes de traitement et de contrôle, mais nous constatons que des services gratuits de réduction des coûts ont été beaucoup mieux fournis aux offres du groupe de traitement. Ces services étaient axés sur la mise en œuvre de la présélection par les conseillers, ce qui

a entraîné une réduction des efforts de recherche de la part des entreprises et une baisse du nombre moyen de candidats envoyés à l'entreprise pour être examinés. Ces résultats sont robustes quand on contrôle les effets de sélection corrélés avec la distribution des caractéristiques des offres. Cela suggère que l'augmentation de la création d'emplois et des embauches pourrait résulter d'une réduction des frictions d'appariement dans le processus de recrutement lui-même et, à notre connaissance, il s'agit de la première étude à démontrer ce lien étroit entre les coûts de recrutement et l'offre d'emploi.

Chapitre 3. La recherche d'emploi et l'intermédiation face à la discrimination : des résultats empiriques suite aux attentats terroristes en France

Fondés sur les idées de Becker et Phelps et Arrow, les études qui ont suivi se sont concentrées principalement sur l'interaction des travailleurs avec l'employeur, soit pendant le processus d'embauche (commencer par Lang and Lehmann (2012) pour la revue de littérature théorique et Bertrand and Dufflo (2017) pour les expériences appliquées), plus récemment, sur la performance des employés (Hjort (2014) et le premier chapitre de cette thèse.). Un sujet qui a reçu moins d'attention est la façon dont la discrimination, ou sa perception, pourrait affecter la recherche d'emploi elle-même.

Dans cette étude, nous examinons les effets sur les résultats du marché du travail d'un choc qui aurait pu considérablement augmenter les attitudes discriminatoires envers un groupe minoritaire spécifique. Bien qu'il existe de nombreuses preuves que la discrimination contre les minorités musulmanes est réelle en France (Adida, Laitin, and Valfort, 2014), nous exploitons les attaques terroristes à Paris et en région parisienne contre le journal satirique Charlie Hebdo, la police et le supermarché Hypercacher entre le 7 et le 9 janvier 2015 comme un choc exogène qui peut avoir sensiblement accru les préjugés contre les minorités musulmanes en France. Pour illustrer le lien entre ces attaques terroristes et une potentielle augmentation de la discrimination, nous montrons une figure qui montre le grand nombre de recherches du mot « islamophobie » sur Google, en France. L'intérêt pour cette recherche est proche de zéro jusqu'à la semaine des attaques de janvier 2015, lors de laquelle il fait un bond pour atteindre 100 dans la semaine de l'attaque, date à laquelle tous les autres points sont normalisés. Et même si nous ne sommes pas en mesure de déterminer exactement pourquoi ou qui recherche ce terme, le chiffre indique que les préjugés contre les minorités musulmanes sont apparus dans la conscience publique dans les jours et les semaines qui ont suivi les attaques.

Nous concentrons notre analyse sur les effets de ce choc sur les efforts de recherche que

montrent trois acteurs du marché du travail : les demandeurs d'emploi inscrits auprès du service public de l'emploi (SPE) Pôle emploi, leurs conseillers d'emploi et les employeurs. En utilisant la stratégie empirique des doubles différences en contrôlant pour la tendance de l'année précédente, nous constatons que le choc a conduit à une baisse transitoire importante des efforts de recherche des personnes issues des minorités, par rapport à celles issues de la majorité, mesurés par le nombre moyen de candidatures par demandeur d'emploi aux offres postées sur le SPE. En moyenne, les demandeurs d'emploi issus des minorités, définis comme dans le premier chapitre en utilisant les prénoms d'origines maghrébines, réduisent leurs démarches de recherche de 11 % par rapport aux demandeurs d'emploi issus de la majorité – ceux dont le prénom est d'origine française – dans les 10 semaines qui suivent les attaques terroristes. En revanche, nous constatons que les conseillers Pôle emploi augmentent de 13 % le nombre d'offres d'emplois qu'ils envoient à leurs demandeurs d'emploi issus des minorités par rapport à ceux issus de la majorité dans les semaines qui suivent les attaques. Fait important, cet effet de conseiller est plus fort pour les conseillers qui sont eux-mêmes issus des minorités et pour les conseillers issus de la majorité dont les fonctions impliquent de travailler avec les demandeurs d'emploi les plus marginalisés. Nous constatons également que le choc incite les employeurs à moins ouvrir leurs offres d'emploi aux minorités, mais uniquement pour les postes en contrat à durée indéterminée, le type de contrat le plus recherché. Nous constatons qu'ils augmentent leurs ouvertures de postes en direction des minorités pour leurs contrats à durée déterminée et en intérim. Pour tous les types de contrats, nous ne voyons aucun impact sur les appariements initiés par les employeurs en raison du choc. Compte tenu de ces impacts importants sur les principaux canaux de recherche d'emploi et d'appariement, nous affirmons qu'il peut être imprudent d'attribuer les impacts mesurés sur les résultats d'embauche exclusivement à la discrimination accrue des employeurs. Ainsi, tout au long du document, nous faisons une distinction conceptuelle entre les effets d'une augmentation de la discrimination sur la recherche d'emploi et ses effets sur les résultats de l'emploi. Nous faisons la distinction entre ces résultats de « premier » et de « deuxième ordre », parce que notre stratégie d'identification ne nous permet pas de distinguer de changement dans les résultats du niveau d'emploi des minorités en raison d'un changement du comportement de recherche des demandeurs d'emploi eux-mêmes (ou bien de leurs conseillers) et les résultats qui passent par un changement du niveau de la discrimination des employeurs. Cet article contribue ainsi à la littérature qui s'intéresse à la manière de mesurer la discrimination sur le marché du travail, et à son existence même. Les études par correspondance (voir Riach and Rich (2002) pour une introduction) trouvent généralement de grandes différences dans les taux de rappel des demandeurs d'emploi issus des minorités ayant des curriculum vitae équivalents à ceux des candidats issus de la majorité (voir Bertrand and Mullainathan (2004)

aux États-Unis et Behaghel et al. (2015), Petit et al. (2011), Berson (2012) et Adida et al. (2016) pour le contexte français). Pourtant, Heckman (1998) a fait valoir que les résultats des études par correspondance sont, à tout le moins, difficiles à interpréter comme indiquant la présence de discrimination sur le marché du travail parce qu'ils sont basés sur le principe que les travailleurs candidatent au hasard. Les résultats présentés dans cet article sur ces résultats de premier ordre montrent que l'impact de la discrimination sur le comportement de recherche peut être important et qu'il est donc difficile de démêler l'effet direct de la discrimination sur les résultats de l'emploi. Une contribution secondaire de ce chapitre est donc basée sur le fait que nous observons le comportement réel de recherche des employeurs pour les entreprises qui recrutent via le SPE.

Dans notre analyse, nous illustrons cette difficulté de séparer l'effet des changements dans l'intensité de la recherche de la discrimination accrue des employeurs en utilisant l'effet des attaques sur les embauches. Nous constatons qu'après le choc, les demandeurs d'emploi issus des minorités signent des contrats à très court terme de manière disproportionnée par rapport demandeurs d'emploi issus de la majorité, mais cet effet est entraîné par une baisse relative de l'utilisation de ces contrats par la majorité. Par conséquent, une explication possible pourrait être qu'en raison d'une baisse effective de la tension du marché du travail associée à la baisse de la recherche de CDI de la part des demandeurs des minorités, les candidats de la majorité ont un taux de recherche d'emploi plus élevé ; ils sont ainsi donc moins dépendants de contrats courts. Une autre explication pourrait être simplement que, parce que les employeurs sont moins enclins à embaucher des membres des minorités dans des contrats permanents, ils se substituent aux membres de la majorité et que cela représente le mécanisme moteur de l'écart dans les embauches à court terme. Malheureusement, nous ne pouvons pas faire la distinction entre les deux, et il existe certainement beaucoup d'autres scénarios qui déterminent la dynamique de l'emploi. Par conséquent, tout au long du document, nous conserverons un positionnement agnostique quant aux raisons de changement de la dynamique d'embauche après le choc. La seule conclusion pour laquelle nous pouvons fournir des preuves à l'appui est que, lorsque nous additionnons les jours de travail sur tous les types de contrat, nous ne trouvons aucune preuve que le choc a entraîné moins de créations d'emplois pour les minorités que pour la majorité.

Cette étude se joint également à la littérature sur les effets des attentats ethno-religieux sur les résultats du marché pour les minorités musulmanes. Les résultats antérieurs sont mitigés en ce qui concerne les effets sur les résultats d'emploi. Åslund and Rooth (2005) constatent que les attitudes à l'égard des minorités musulmanes en Suède ont changé après les attentats terroristes du 11 septembre, mais ne trouvent aucun effet sur l'emploi de ces minorités. De même, Kaushal et al. (2007) ne trouve aucun effet sur l'emploi des minorités aux

États-Unis après le 11 septembre, mais constate un impact sur les salaires. En outre, Gautier et al. (2009) trouve que les prix du logement diminuent dans les quartiers minoritaires et que la ségrégation résidentielle augmente à la suite d'une attaque terroriste ethno-religieuse, reflétant des résultats similaires trouvés par Ratcliffe and von Hinke Kessler Scholder (2015) qui montrent que des attaques peuvent avoir eu un effet négatif sur les niveaux d'emploi dans les quartiers comptant un pourcentage plus élevé de résidents issus des minorités. Cet article contribuera ainsi à une meilleure compréhension de la façon d'interpréter les résultats de second ordre analysés dans ces études.

Nous affinons l'analyse dans ce document en utilisant une mesure de discrimination existante au niveau local. Nous utilisons le pourcentage des voix au niveau de la municipalité pour le parti d'extrême droite en France, le Front national (FN), lors des élections présidentielles françaises de 2012. Ces données servent à capturer la discrimination existante que les demandeurs d'emploi peuvent rencontrer dans leurs recherches d'emploi. Cette analyse révèle des effets fortement hétérogènes du choc de la discrimination pour les demandeurs d'emploi et les conseillers. Alors que nous observons un effet négatif sur les démarches de recherche des demandeurs d'emploi issus des minorités dans les zones à fort et faible vote FN, l'impact est 3,5 fois plus important dans les zones à faible vote FN. Cet effet est contrebalancé par les efforts mis en œuvre par les conseillers dans les zones à faible vote FN. Nous motivons l'interprétation de ces résultats, peut-être frappants, en modélisant l'effet marginal de la discrimination comme décroissant sur la discrimination latente. Le coût de recherche des demandeurs d'emploi est endogène à la probabilité de trouver un emploi, ce qui est affecté par la discrimination réelle et/ou perçue sur le marché. Étant donné que l'effet de la discrimination est décroissant, l'effet marginal d'un choc important sur la discrimination devrait être plus grand dans les zones où les niveaux initiaux de discrimination sont faibles. Pour les conseillers, nous interprétons leurs efforts comme un intrant dans le paramètre d'efficacité d'une fonction d'appariement du marché du travail. Si leurs efforts représentent une tension entre les efforts faits auprès des demandeurs d'emploi qui ont du mal à trouver un emploi et le coût de leur propre discrimination, de tels résultats empiriques hétérogènes sont tout à fait plausibles.

Pour soutenir l'utilisation de la part de vote du FN en tant que proxy pour la discrimination existante, nous revenons à l'utilisation des données de recherche de Google. Nous regardons d'abord les tendances de recherches dans l'année précédant le choc, ventilées par régions françaises. Nous constatons que la part de vote FN est en effet fortement corrélée avec les recherches Google qui connotent la prévalence de la discrimination et de l'animosité envers notre groupe minoritaire. Nous mesurons ensuite le changement dans les tendances de recherche autour de la date des attaques de janvier 2015. Nous montrons que le volume

de ces termes de recherche négatifs augmente brusquement en raison du choc et continue d'être très positivement corrélé avec la part de vote FN. Nous examinons ensuite les termes qui évoquent la cohésion sociale et constatons un intérêt élevé suite au choc, mais, à l'opposé, ces termes sont négativement corrélés avec la part des votes pour le FN. Bien que spéculatifs, ils constituent un scénario dans lequel le service public de l'emploi a un impact qui atténue les effets d'un tel choc par une internalisation des difficultés croissantes auxquelles sont confrontés ces demandeurs d'emploi. Et que cet « effet compensatoire » des conseillers d'emploi peut être lié aux préférences pour l'intégration et la cohésion sociale. Il est cohérent donc que l'effet soit beaucoup plus grand pour les conseillers qui sont eux-mêmes issus des minorités. Nous constatons également des effets positifs importants pour les conseillers issus de la majorité, mais seulement parmi ceux qui se spécialisent dans le retour au travail des demandeurs d'emploi les plus marginalisés, et nous n'observons aucune augmentation de l'effort des conseillers issus de la majorité à qui sont attribués des demandeurs d'emplois « normaux ». Enfin, l'existence d'un effet compensatoire peut ne pas être isolée de l'intermédiation par des conseillers d'emploi professionnels, car de nombreux emplois sont trouvés par des canaux informels et par des réseaux personnels. Des travaux futurs sur ce sujet sont nécessaires pour mieux comprendre comment la discrimination se filtre à travers les acteurs du marché du travail pour élaborer des politiques efficaces pour lutter contre ses effets néfastes.

Part I

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

joint with Amanda Pallais and William Parienté¹

Abstract

Examining the performance of cashiers in a French grocery store chain, we find that manager bias negatively affects minority job performance. In the stores studied, cashiers work with different managers on different days and their schedules are determined quasi-randomly. When minority cashiers, but not majority cashiers, are scheduled to work with managers who are biased (as determined by an Implicit Association Test), they are absent more often, spend less time at work, scan items more slowly, and take more time between customers. This appears to be because biased managers interact less with minorities, leading minorities to exert less effort. Manager bias has consequences for the average performance of minority workers: while on average minority and majority workers perform equivalently, on days where managers are unbiased, minorities perform significantly better than do majority workers. These facts are consistent with a model of statistical discrimination in which, because manager bias depresses minority performance, the firm sets a higher hiring standard for minorities.

¹We would like to thank Yann Algan, David Autor, Thomas Le Barbanchon, Eric Cediey, Raj Chetty, Bruno Crépon, Muriel Dejemppe, Eleanor Dillon, Esther Duflo, Erin Fletcher, Roland Fryer, Ed Glaeser, Lisa Kahn, Lawrence Katz, Danielle Li, Florian Mayneris, four anonymous referees, and seminar participants at Berkeley Haas, London School of Economics, MIT, NBER Labor Studies, Nova Business School of Economics, Paris School of Economics, Sciences Po, University of California Santa Barbara, University of Illinois at Urbana-Champaign, University of Namur, and UC Louvain for their helpful comments. We would also like to thank Lisa Abraham, Jenna Anders, Frédéric Cochinard, Elizabeth Mishkin, and Adrien Pawlik for superb research assistance. We gratefully acknowledge funding from the Fonds d'Expérimentation pour la Jeunesse (France).

1 Introduction

A vast economic literature tests for the presence of labor market discrimination, a setting in which equally-productive minority and non-minority workers are treated differently in terms of hiring, pay, or promotion. Becker’s pioneering work, *The Economics of Discrimination* (1957), introduced the notion of taste-based discrimination: employers experience disutility when employing minority workers and compensate by paying minorities less or requiring them to be more productive for the same wage. A subsequent body of work, starting with Phelps (1972) and Arrow (1973), conceived of discrimination not as a matter of animus but one of imperfect information. Unfavorable priors about minority workers’ productivity or imperfect screening precision causes employers to treat equally-skilled minority and majority workers unequally. Building on these insights, Lundberg and Startz (1983) and Coate and Loury (1993) showed how statistical discrimination could potentially depress minorities’ skill investments by leading minorities to correctly believe that these investments would not be fully rewarded. As a result, statistical discrimination may lead to a self-fulfilling prophecy whereby employers’ adverse prior beliefs about minorities’ skill levels are self-confirming in equilibrium.

Something that unites these strands of literature is the implicit assumption that employers’ tastes and beliefs do not directly impact worker productivity. Although statistical discrimination might inhibit skill investment, it does not directly affect the performance of workers with given skill levels. However, a strand of literature beginning with Steele and Aronson (1995) documents that adverse stereotypes about minority groups’ abilities can directly impair group members’ performance. This body of work demonstrates that when stereotypes are made salient prior to performance (e.g., test-takers are asked to report their race or gender), blacks, Hispanics, and women tend to perform significantly worse than in settings where group membership is not made salient. This line of research, which has so far had limited influence on the economic literature, implies that adverse employer beliefs about minorities – whether stemming from animus or statistical discrimination – could be self-fulfilling, not because they inhibit minority skill investment but because they induce poorer performance. Related research shows that individuals’ own stereotypes can negatively impact their performance (Coffman 2014) and that these stereotypes need not be fully accurate (Bordalo et al. 2016).

This paper presents a novel test of whether discriminatory beliefs directly affect minority workers’ job performance in a real-world workplace.² We study 34 outlets of a French grocery

²This paper is related to the literature showing that workers and students benefit from interacting with co-ethnics. See, for example, Dee (2004, 2005), Stoll, Raphael, and Holzer (2004), Stauffer and Buckley (2005), Giuliano, Levine, and Leonard (2009, 2011), Price and Wolfers (2010), and Hjort (2014). It is also

store chain. These stores employ a sizable proportion of minority workers that, based on their names, we identify as having a North African or Sub-Saharan African origin.³ We study new cashiers hired on six-month contracts since these workers are assigned quasi-randomly to managers. These cashiers, like all cashiers in the stores, work with different managers on different days. Unlike more senior workers, however, they are not allowed to submit schedule preferences. Their schedules are determined by a computer program which assigns shifts to meet predicted demand, taking into account the preferences of more senior workers. Thus, the minority and majority workers in our sample do not choose the managers they work with and they work with the same managers under similar conditions. Workers know which managers they will be working with beforehand as both worker and manager schedules are publicly posted several weeks in advance.

We measured managers' bias towards minorities with an Implicit Association Test (IAT), which is widely used to measure bias, particularly in psychology (see Nosek, Greenwald, and Banaji 2007; Lane et al. 2007; and Greenwald et al. 2009 for summaries of the literature). IAT scores have been correlated with many real-world decisions and are difficult for subjects to manipulate.⁴ The test uses the speed with which subjects categorize prompts to determine their implicit association between two concepts: here (1) traditionally French or North African sounding names and (2) words indicating worker competence or incompetence. Our manager bias score thus measures the extent to which managers associate North African names with poor worker performance. This concept is correlated with, but distinct from, managers' distaste for minorities (Agerström, Carlsson, and Rooth 2007).

Each of the stores in our sample tracks individual performance at a daily level. The stores provided us with data on absences and time worked (determined by time clock data), scanning speed, and time taken between customers. The firm considers absences particularly important: being absent three times is one of the few ways a worker can be fired during her initial six-month contract. The firm also prioritizes scanning speed, posting a list of workers' articles scanned per minute in the break room each week. The firm uses these performance metrics along with the managers' observations about workers' performance and customer relations to determine whether workers will be offered a longer contract at the end of their six-month contract. Approximately 30-40% of workers are offered a longer contract.

related to the literature started by Rosenthal and Jacobson (1968) showing that teachers' expectations about student performance can directly affect student outcomes.

³Workers are categorized into minority and non-minority status based on their names because in France it is illegal to ask workers their ethnicity. 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 it would not be able to identify any individual in the study.

⁴See, for example, Kim (2003), Friese, Bluemke, and Wänke (2007), Green et al. (2007), Greenwald et al. (2009), and Rooth (2010).

We assess whether minority cashiers perform worse on the days they work with managers who are biased against their minority group. Because there may be other differences between more- and less-biased managers – biased managers may simply be less skilled, for example – we do not want to simply attribute any change in minority performance when working with more-biased managers to manager bias. Instead, we utilize a difference-in-difference methodology, comparing the change in minority workers’ performance under more- and less-biased managers to the change in non-minority performance.

We find that manager bias leads minorities to perform worse. Minorities are more likely to be absent when scheduled to work with more-biased managers. When they do come to work, they spend less time at the store: specifically, they are much less likely to stay after their scheduled shift ends. While workers are allowed to leave when their shift ends, managers can ask them to work late. Because workers are paid based on time worked, we estimate that minorities earn 2.5% less as a result of manager bias.

Minorities also scan items more slowly and take more time between customers when working with biased managers. Throughout our analyses, none of the differential effects of working with more-biased managers are explained by the other manager characteristics we have, including the managers’ own minority status. The effect of manager bias is concentrated in stores with fewer minority workers and appears to grow during the contract (though this latter difference is not statistically significant).

We combine data from a worker survey conducted after contract expiration with our administrative data to distinguish between theories of discrimination that can explain our results.⁵ First, we find little evidence that animus – or biased managers treating minorities poorly – can explain our results. Minority workers do not report that biased managers disliked them or that they disliked biased managers. They report that biased managers were *less* likely to assign them to unpleasant tasks (cleaning) and no more likely to assign them unpleasant registers or breaks.

Our evidence is most consistent with a theory in which biased managers interact less with minority workers. Research in psychology on “aversive racism” has found that individuals with implicit biases towards minority groups are less likely to speak to, more hesitant in speaking to, and less friendly towards members of those groups.⁶ They may feel less comfortable interacting with minorities or they may be concerned about appearing biased. Using

⁵One explanation that cannot drive the day-to-day differences in performance we find is that biased managers depress minority human capital accumulation. Minorities may accumulate fewer skills under biased managers, but we would not detect this since minorities would have any skills they learned working with unbiased managers on the days when they work with biased managers.

⁶See, for example, McConnell and Leibold (2001), Dovidio, Kawakami, and Gaertner (2002) and Hebl et al. (2002). Dovidio and Gaertner (2008) summarize this literature.

whether a worker remembered each manager as an indicator for worker-manager interaction, we find minorities were less likely to remember biased managers. Worker-manager interaction appears to be a key determinant of performance: workers performed substantially better when working with managers they later remembered.⁷ Biased managers interacting less with minorities is consistent with their being less likely to ask minorities to do cleaning duties. It may also explain why minorities are less likely to stay after the end of their shifts when working with biased managers: the managers might simply not ask them to.

An alternative explanation concerns self-stereotyping or stereotype threat. Minorities may hold negative stereotypes about their suitability for the job or be aware of existing stereotypes about their group and biased managers may activate these negative stereotypes. To test whether this occurred, we asked workers which managers gave them the most confidence in their abilities. Minorities do not report that biased managers gave them less confidence in their abilities. Nevertheless, this does not rule out an explanation whereby biased managers subconsciously activated minorities' negative stereotypes.

Finally, we find that the negative impact of manager bias on minority performance may lead to statistical discrimination in hiring. Under statistical discrimination, the firm infers worker productivity from workers' observable characteristics and minority status. By depressing minority productivity, manager bias can lead the firm to act as if minority workers are less productive than majority workers with the same characteristics. Because the firm has to pay workers the same wage independent of their performance, it would set a higher hiring threshold for minorities. Thus, in the absence of manager bias, hired minorities would perform better than hired majority workers. Overall, we find that minority and majority workers perform equivalently. There is no difference in their average absence rates, time spent at work, articles scanned per minute, or time taken between customers.⁸ However, consistent with statistical discrimination, when working with unbiased managers, minority workers perform substantially better than non-minorities. They are half as likely to be absent and scan significantly faster.

With some assumptions, we can combine our performance metrics to estimate the number of customers each worker serves per day. On average, minorities serve an insignificant 2% more customers than do majorities. However, when they work with unbiased managers, minorities serve 9% more customers than majorities. While the average minority is at the 53rd

⁷This is consistent with Mas and Moretti (2009) which finds that monitoring improves cashiers' performance.

⁸The similar performance of minority and majority workers is (weakly) inconsistent with a model of taste-based discrimination in which the firm faces a utility cost of employing minorities. In this model, the firm requires higher average productivity from minorities to hire them at the same wage. While minorities may perform slightly better, we can reject that hired minorities perform more than 4% better on average than hired majority workers.

percentile of average worker performance, on days with unbiased managers she is at the 79th percentile. This suggests there are substantive consequences of manager bias on minority workers’ performance and thus, in theory, workers’ subsequent labor market outcomes.⁹

2 Setting

We study entry-level cashiers in a large French grocery store chain. These cashiers are hired on a specific contract called *Contrat de Professionnalisation* (CP): a six-month contract subsidized by the government. In return for the subsidy, the firm trains CP workers (or CPs) to be cashiers and on 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 contract start date in which workers are trained without pay. During this week, either party can walk away from the contract without penalty.

CP cashiers perform 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, CPs 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.¹⁰ Schedules are determined three weeks at a time and, once determined, publicly posted. Manager schedules are also publicly posted in advance, so workers know ahead of time which managers they will be working with. The chief cashier (the managers’ boss) can, in theory, revise the schedules assigned by the computer system. However, this happens very rarely.

The stores typically have around 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 at once. The manager on duty sits in a special station in the middle of the registers. When a cashier arrives for her shift, she “badges” (clocks) in near the manager station. She typically has a brief conversation with the manager, who gives her the day’s

⁹Unfortunately, we do not have data on which workers were offered a second contract. Our identification strategy also does not lend itself well to determining the effect of manager bias on workers’ subsequent labor market outcomes since over the six-month contract, there is little variation in the average bias workers are exposed to.

¹⁰A split shift occurs when a worker is scheduled to work for two separate periods in the same day (for example, from 9 am to 12 pm and from 3 pm to 6 pm).

news and assigns her to a register. Some workers are assigned to special cash registers, such as the 10-items-or-less line or the self-checkout, though this is rare for CPs. The worker then gets her till (cash box) from the safe, sets it up at her station, and starts receiving customers. There are no baggers in these stores; customers bag the items themselves.

The manager roams the store, talking with cashiers and monitoring them at their stations. She manages the lines, opening and closing new ones and directing customers to short ones. Cashiers whose lines are closed are assigned to other tasks such as aisle arrangement, the welcome desk, or assisting managers. The manager also decides when workers can go on break, though the *amount* of break time is specified in workers' contracts. Workers are allowed to leave at the end of their shifts, but the manager can ask them to work late. Before leaving, cashiers confer with the manager, return their tills to the safe, and badge out near the manager station.

CP workers are hired in waves: approximately twice a year each store has a "promotion," in which new CPs are hired. The managers we study are rarely involved in the hiring process, which is conducted by the chain's central office and the store's chief cashier.

The most important performance metrics for workers are showing up to work, arriving on time, and having the correct amount of money in the till. During their initial contract, workers can be fired only for misconduct, which includes having more than three absences, being late more than three times, having more than three warnings for misbehavior, or having even one report of violent conduct or one large till deviation. If misconduct occurs, the chief cashier decides whether to fire the CP, relying on the advice of the managers. Aside from misconduct, 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 is posted in the employee break room.

CPs are not paid based on performance; they are paid solely based on time worked. In particular, CPs are not paid for days they are absent, though after three sick days and a doctor's authorization, the government pays 70% of workers' pay during their sick leave.¹¹ CPs' paychecks are also adjusted if they work more or less than scheduled.

CPs' main incentive to perform well is the opportunity to receive another contract. After their initial six-month contract, about 30% to 40% of workers are offered another contract. 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.

Managers are on indefinite-term contracts; their pay is fixed, not dependent on their

¹¹CPs earn vacation time, but cannot use it for days off: they are paid for their vacation days after the end of the contract.

performance. Managers' performance is assessed annually based on customers' checkout experience, which is determined by how quickly the lines move, and to a lesser extent register cleanliness, the stocking of the small shelves at each register, and effective handling of customer problems. While managers are graded on their support of inexperienced cashiers, they do not have the primary responsibility for cashier training.

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 worker performance. The manager survey data provide our measure of manager bias. We use the worker survey data to learn about the mechanism for the effects of manager bias.

3.1.1 Administrative Store Data

We collected daily data for each CP in a given promotion over a six-week period between July 2011 and August 2012. We have schedule data: the precise times at which workers and managers were supposed to begin and end their shifts, allowing us to determine which manager(s) a worker was scheduled to work with on a given day.

We also have badge data: the precise times that workers and managers badged in and out of the stores. Both managers and CPs must badge in and out at the beginning and end of their shifts and for breaks, so we have actual working times to the minute. Combining these data with the schedule data provides our first two metrics of worker performance: absence and the number of minutes worked relative to the number of minutes the worker was scheduled to work.¹² Time spent at the store can differ from the schedule for three reasons: (1) workers arrive earlier or later than scheduled, (2) workers leave earlier or later than scheduled, and (3) workers take breaks. While workers are entitled to breaks, breaks are not scheduled by the computer program.

We also have daily worker performance data, most importantly, articles scanned per minute. The time over which articles per minute is calculated starts when a worker scans a customer's first item and stops when a worker scans a customer's last item, so it is not affected by the time between customers. We also have two other determinants of line speed. The first is inter-customer time: the time between finishing one customer's transaction and

¹²We analyze time worked relative to time scheduled instead of simply time worked to gain precision.

starting to scan the next customer’s items. The second is payment time: the time between the scanning of a customer’s last item and the completion of the customer’s transaction, during which time she is paying. While the firm tracks both of these metrics, it does not emphasize them as key performance measures.

Our final sample has 34 stores, 204 workers, and 4,371 worker-day observations.¹³ While we asked for data from all of the chain’s stores in France, we received the necessary administrative data from 45 of them. From these 45 stores, we eliminate 11 in which managers did not take the IAT. (The process of getting managers to take the IAT is explained more below.) Because most stores had multiple promotions during the year, we have data on 51 promotions from the 34 stores in our sample. Two stores did not provide data on inter-customer time, while four did not provide data on payment time, so we have slightly smaller samples for these outcomes.¹⁴

Because we wanted variation in the timing of the observations during the contract, we asked for data on weeks three through eight of the contract for some promotions and weeks 18 to 23 for others. We have data on weeks 18 to 23 for promotions that occurred chronologically earlier and data on weeks three to eight for promotions that occurred chronologically later because stores kept data for only one year.

In addition to these data, the stores provided a few other worker and manager characteristics, most importantly, their names. In France, it is illegal to ask people about their ethnicity. Thus, we use workers’ names as an indicator of their minority status. 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. We provided ISM CORUM with separate lists of first and last names, so that it 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) North African, (3) Sub-Saharan African, (4) Mixed or undetermined, and (5) Other (including names of Turkish and Asian origins). We consider workers with a North African and Sub-Saharan sounding first or last name as the minorities in this context. In the appendix we show results are robust to using other definitions of minority status.

We also classified workers’ and managers’ genders using their names. The chain also

¹³Throughout the paper, we cluster standard errors at the store level to allow for correlation in performance both within and across days in a store. While we have more than 30 clusters, we show that p-values are similar when we use a wild cluster bootstrap procedure that is robust to having a small number of clusters (Cameron, Gelbach, and Miller 2008; Garthwaite, Gross, and Notowidigdo 2014).

¹⁴One store did not provide data on either of these outcomes. The remaining stores provided data on the total amount of inter-customer time or payment time during the worker’s shift, not scaled by the number of customers served.

provided managers' ranks (positions) within the store and managers' dates of birth.

3.1.2 Manager Survey and IAT

We measure managers' bias towards minority workers using an Implicit Association Test (Greenwald, McGhee, and Schwartz 1998; Nosek, Greenwald, and Banaji 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 sides of a computer screen. In our case, subjects were presented with (1) names typically indicating a French origin (e.g., Jean) or names traditionally indicating a North African origin (e.g., Ahmed) and (2) adjectives that describe good employees (e.g., reliable) or bad employees (e.g., incompetent).

In all rounds, one word at a time (either a name or adjective) comes onto the screen and subjects are told how to categorize it (for example, adjectives describing good employees to the left, adjectives describing bad employees 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 told to categorize French sounding names and negative adjectives to the same side of the screen, while in the other, they are tasked with categorizing North African sounding names and negative adjectives to the same side. The idea behind the test is that if a subject has an implicit association between two concepts (e.g., workers of North African origin and bad employees), it should be easier and quicker to do the categorization 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 North African sounding names and negative adjectives are categorized on the same side of the screen, relative to when French sounding names and negative adjectives are categorized on the same side.¹⁵

IAT scores have been found to be correlated with judgments, choices, and psychological responses (Bertrand, Chugh, and Mullainathan 2005). For example, IAT scores are correlated with voting behavior (Friese, Bluemke, and Wänke 2007), callback rates of minority job applicants (Rooth 2010), and doctors' provision of differential medical treatments by race (Green et al. 2007). Moreover, research suggests that it is very difficult to fake an IAT score.¹⁶

¹⁵We randomized the order in which subjects completed these rounds. We also included practice rounds to mitigate order effects (Nosek, Greenwald, and Banaji 2007). We used the computer software Inquisit to administer the IAT.

¹⁶See Banse, Seise, and Zerbes (2001), Egloff and Schmukle (2002), Kim (2003), Greenwald et al. (2009), and Hu, Rosenfeld, and Bodenhausen (2012). Faking a score on an IAT requires a specific strategy of slightly speeding up or slowing down in certain blocks, a strategy that few participants spontaneously discover (Greenwald et al., 2009).

The chain's human resource office contacted the chief cashier in each store, asking her to get the managers to take the IAT as part of a study. While managers could likely tell from the IAT that the study concerned their beliefs about minorities, they did not know the exact purpose of the study. Managers were allowed to take the test during work hours, but did not receive any payment for doing so. Initially, managers received an email with a link to the IAT so that they could take the IAT at their convenience. We sent email reminders and periodically called the chief cashiers to induce more managers to take the test. We also visited stores that had technical difficulty accessing the IAT website, administering the IAT in person to these managers.

The managers took the IAT on average 17 months after the administrative data in our sample. Thus, neither taking the IAT nor knowledge of our study could have affected managers' treatment of minority workers in our data. Managers' experience in the store could have affected their implicit beliefs, but it seems very unlikely that interaction with the CPs in our study would have led to variation in those beliefs. In particular, our identification strategy ensures that the more- and less-biased managers we compare worked with the same CPs. The vast majority of our managers (85%) had been at the store for over 10 years, so would have seen at least 20 different CP promotions, several more recent than the ones we study. These managers have managed 100 to 250 workers at a time for many years (relative to an average of six CPs per manager in our study), most of whom they work with much longer than with CPs. Moreover, the effects do not change with the length of time between the administrative data and when managers took the IAT. Finally, we use male names in the IAT (over 90% of our CPs are female) so managers are not prompted by the names of specific workers.

While it is unlikely that interacting with the CPs in our study affected managers' IAT scores, interacting with minorities in general might have. For example, if minorities disliked some managers and, as a result, performed badly for them, these managers might have developed negative beliefs about minority workers' performance. We think this is unlikely: minorities do not report disliking biased managers and the negative impacts of manager bias on minority performance appear to be driven by manager actions, not solely worker actions. Nevertheless, we cannot fully rule out this alternative explanation.

We have IAT scores for 77% of the managers in the 34 stores. On most dimensions we have, managers who did and did not take the IAT look similar. They were the same average age and were equally likely to be a minority and to have a high position in the store. Calculating manager fixed effects for all of our performance outcomes produces no significant differences between managers who did and did not take the IAT. (These differences are also inconsistently signed.) There is no correlation between the number of days it took managers

to take the IAT after we requested it and their IAT scores, so it does not necessarily appear that more-biased managers were more reluctant to take the test. Male managers were less likely than were female managers to take the IAT. We show, however, that our results are robust to including controls for manager gender and manager gender interacted with worker minority status.

3.1.3 Worker Survey

We conducted a telephone questionnaire from May 2013 to September 2013, surveying former CPs about their relationship with each of their managers. The heart of the questionnaire comprised CPs ranking their managers on a variety of dimensions. Respondents rated the extent to which they remembered each manager, which we use as a measure of worker-manager interaction. We also described manager traits or actions (e.g., the manager who liked the worker best) and asked workers to rate in order the top and bottom three managers on each trait.¹⁷ We provided workers with a list of managers, but did not tell workers managers' IAT scores, nor did we ask whether they thought the managers were biased.

Half of surveyed workers responded. The main cause of non-response was that CPs no longer had the same contact information and their phone numbers had been disconnected. (Only 2% of workers answered the phone but refused to answer the survey.) We have survey responses for 94 workers in our main sample. Because we did not know which stores would provide performance data when we conducted the survey, we surveyed a larger sample. We also have survey data for 74 workers for whom we have manager IAT scores but not performance data and 10 workers for whom we have performance data but not manager IAT scores.

Controlling for store fixed effects, minorities were 7.4 percentage points less likely to respond to the survey (off a base of 52.6% for majorities). While this is not a statistically significant difference, it is not a small one. Appendix Table A.1 compares the characteristics and performance of workers who responded and workers who did not, for the whole sample and separately for minority and majority workers. There are few differences between respondents and non-respondents. Consistent with chance differences, of the 30 comparisons in the table, one is significant at the 5% level and two are significant at the 10% level.¹⁸

¹⁷Most workers had six or fewer managers. In a pilot, we asked workers to rate all of their managers, however, workers found this difficult. There was substantial non-response and a few workers asked to stop the survey.

¹⁸Respondents worked more minutes per day than did non-respondents both among majority workers and in the overall sample. Respondents in the overall sample were less likely to be from the Paris region.

3.2 Descriptive Statistics

Table 1 reports descriptive statistics. We know only two things about all CPs: their minority status and gender, both based on their names. While 28% of workers are minorities, only 7% are male. The worker survey paints a slightly richer picture of workers in these stores. Despite this being an entry-level job, the average worker is 30 years old and has had four previous jobs. Only 11% of the sample has had no prior employment. Most workers (58%) do not have a high school degree, while relatively few (7%) have more than a high school education.

Managers tend to be older, averaging 41 years of age. Relatively few of the managers are minorities themselves (6%) and few are male (10%). Managers' IAT scores suggest that most are biased against minorities. For ease of interpretation, throughout the paper, we divide managers' raw IAT scores by the standard deviation in our sample (0.36). Positive scores indicate a preference for majorities while negative scores indicate a preference for minorities. The average (scaled) manager IAT score is thus 1.35, which means that the average manager is 1.35 standard deviations away from being completely unbiased. Using the typical thresholds in the literature,¹⁹ 9% of managers show little to no bias against minorities, 20% show a slight bias against minorities, and 66% of our sample shows moderate to severe bias against minorities. Only 4% of our sample shows a preference for minorities. The managers seem approximately as biased as US undergraduates are against African-Americans though more biased than Americans who choose to take an IAT online (Amodio and Devine 2006; Smith-McLallen et al. 2006; Mooney 2014).

Appendix Table A.2 shows the results of regressing manager IAT score on manager characteristics. The point estimates suggest that older managers tend to be more biased, while minority managers are less biased against their own group. However, none of these coefficients are significant, partially because we have so few minority managers.²⁰ Minority and majority CPs work in stores where managers are equally biased.

Workers are scheduled to work just over four days per week on average (in addition to the training day). Working days are distributed relatively evenly Monday through Saturday. We have relatively few observations on Sundays as the stores open on Sundays only during December. Workers are scheduled to work just over seven hours per day on average. The

¹⁹See, for example, Greenwald, Nosek, and Banaji (2003), Rooth (2010), Haider et al. (2011), and Hahn et al. (2014). Raw IAT scores below -0.15 indicate some preference for minorities; scores between -0.15 and 0.15 indicate little to no bias; scores between 0.15 and 0.35 indicate a slight bias against minorities; and scores above 0.35 show moderate to severe bias against minorities.

²⁰The coefficients suggest that, on average, minority managers are 0.44 standard deviations less biased against minorities and a manager 10 years older is 0.08 standard deviations more biased. Controlling for other manager characteristics and store fixed effects, these effects decrease to 0.14 and 0.03 standard deviations, respectively.

median shift starts at 10:15 am and ends at 8:15 pm.

Table 1 also provides the means of the dependent variables. First, CPs are absent less than 2% of working days, an absence rate that leads to an average of two absences over the six-month contract. Second, CPs work almost exactly the number of minutes they are scheduled. Workers badge out of the store during breaks, but they tend to arrive earlier and stay later than scheduled. On average, CPs scan approximately 18.5 articles per minute, take just under 30 seconds between finishing one customer’s transaction and starting the next²¹ and spend approximately 50 seconds per customer in payment time.

3.3 Exogeneity

Throughout the paper we want to interpret any change in performance when minorities worked with biased managers – relative to when majorities worked 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. 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. Throughout the paper, we use the CPs’ and managers’ schedules to construct CPs’ exposure to bias since CPs’ actual working times respond to the managers they are paired with.

Panel A of Table 2 compares the shifts minority and majority workers were scheduled to work. Each column in the panel presents a separate regression of a characteristic of a scheduled working day on an indicator for the worker’s minority status. We control for store fixed effects, as shift assignment is only quasi-random within a store. We cluster standard errors at the store level.

The first dependent variable is the bias (IAT score) of the CP’s scheduled manager. For workers scheduled to work with multiple managers on a given day, this is a weighted average of the managers’ IAT scores, where the weights are based on 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 their schedules were assigned non-randomly by managers, they would have been less likely than majority workers to work with

²¹We eliminate 25 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 outside 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 worker minority status.

biased managers. Instead, we see that the difference is not significant and the point estimate goes in the other 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 (who are higher in the store hierarchy than Level 3 managers). We see no difference in the likelihood that minority and majority workers are scheduled to work with different types of managers. Nor do we see a difference in the number of managers they work with on a shift.

Workers may 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 in all other stores (excluding the store itself) on that date. We see no evidence that minority workers work on particularly productive or unproductive days, nor that minority workers are any more or less likely to work in the early morning or late evening. Minority and majority workers work the same number of hours per day and are equally likely to have split shifts. Panel A of Appendix Table A.3 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 of Table 2 assesses whether minority workers work with more- and less-biased managers under the same conditions as do non-minorities. It presents the results of estimating the equation

$$y_{ist} = \alpha + \beta_1(\text{minority}_i \times \text{bias}_{ist}) + \beta_2\text{bias}_{ist} + \beta_3\text{minority}_i + \delta_s + \varepsilon_{ist}. \quad (1)$$

Here, y_{ist} is a characteristic of the shift worker i in store s who was scheduled to work on day t . Minority_i is an indicator for worker i being a minority and bias_{ist} is the scaled IAT score of the manager the worker was scheduled to work with on day t in store s . Store fixed effects, δ_s , are included. The coefficient β_2 can be significantly different from zero without violating our key assumption, though it never is. This term measures how the conditions under which more- and less-biased managers work with non-minorities differ. The coefficient β_3 measures how the working conditions of minority and majority CPs differ when working with unbiased managers. We see across the board that these coefficients are insignificant. The primary coefficient of interest, β_1 , shows how the working conditions of minority CPs change relative to those of majority CPs when both work with a manager one standard deviation more biased. Again, all the estimated coefficients are insignificant. Panel B of Appendix Table A.3 shows that the all the coefficients are also insignificant when we include worker instead of store fixed effects. Appendix Table A.4 shows that minority CPs are not

differentially likely to be scheduled to work at the same times as other minority CPs.

4 Effect of Manager Bias on Performance

We now turn our attention to assessing whether minority workers perform worse when paired with biased managers. We first consider absence rates and the amount of time spent at work, which are important to the firm and directly affect workers' pay. Then, we consider measures of performance while at work, the most important of which is articles scanned per minute.

To determine the effect of manager bias on worker performance, we estimate the equation

$$y_{ist} = \alpha + \beta_1(\text{minority}_i \times \text{bias}_{ist}) + \beta_2\text{bias}_{ist} + \delta_i + X_{ist}\beta_3 + \varepsilon_{ist}. \quad (2)$$

Here, y_{ist} is a performance metric for worker i in store s on day t . Minority_i and Bias_{ist} are defined as in the previous section. The regression controls for worker fixed effects, δ_i , and shift characteristics X_{ist} . Standard errors are clustered at the store level. The coefficient of interest, β_1 , measures how minorities' performance changes (relative to the change in non-minority performance) when working with a manager one standard deviation more biased.

We expect the estimate of β_1 to be attenuated due to measurement error. Workers' names do not provide a perfect measure of minority status and we do not have IAT scores for all managers. However, the largest source of measurement error is likely to be that managers' IAT scores are not a perfect measure of bias. Nosek, Greenwald, and Banaji (2007) summarizes studies measuring the IAT's reliability over time and finds that individuals' scores on different IAT administrations have a correlation of approximately 0.56, an effect that doesn't change with the length of time between testing. If the IAT is a combination of managers' true implicit bias and noise that is uncorrelated across test administrations, the coefficients of interest will be attenuated by a factor of approximately 1.8 due to measurement error in the IAT score.

4.1 Time Spent at Work

CPs are absent less than 2% of days. Absence rates increase throughout the week, starting at a low of 0.7% on Monday and reaching 2.3% on Saturday; absence rates 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 3 shows the effect of manager bias on absence rates. The first column adds no controls, X_{ist} . It shows that working with a manager with an IAT score one standard deviation higher leads minorities to have an absence rate one percentage point higher. The

effect is large (70% of the mean) even before measurement error corrections. The subsequent two columns add (1) day of the week fixed effects and controls for the shift starting in the early morning and ending in the late evening and (2) date fixed effects, respectively. Adding these controls does not change the coefficient of interest.

The final column includes as controls a dummy for the manager being a minority and an indicator for the worker and the manager both being minorities. Because there are so few minority managers, we do not estimate these terms precisely. However, 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 the manager's 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 biased managers on the amount of time spent at work. Panel B of Table 3 replicates Panel A where the dependent variable is the number of minutes the CP worked in excess of the number of minutes she was scheduled to work. These regressions exclude days the worker was absent.²²

The panel shows that even when not absent, minorities work less when paired with a biased manager. When working with a manager one standard deviation more biased, they work about 3.3 fewer minutes (one twelfth of a standard deviation before correcting for measurement error). As in the prior panel, the result is robust to the addition of controls and is not driven by the manager's minority status. However, unlike in the previous panel, the point estimates do not suggest that majority workers spend more time at work when scheduled to work with biased managers.

Aside from being absent, there are three main ways that a worker could spend less time at work: she could arrive later, leave earlier, or take longer breaks. We find that working with a biased manager primarily leads minorities to leave work earlier. Appendix Table A.5 presents regressions similar to those in Table 3, where the dependent variables are different aspects of time spent at work. There is no significant effect of manager bias on arrival time or break time. However, a minority worker paired with a manager one standard deviation

²²Because we eliminate days workers were absent, these regressions could be biased. For this bias to be driving our results, it would have to be that minority workers would have chosen to work *more* than average (relative to their schedules) on the days that being scheduled to work with a biased manager led them to be absent. Instead, we believe that any bias likely attenuates our results. It seems reasonable that the days that minority workers were absent as a result of being paired with a biased manager are days that they would have worked relatively less had they arrived.

more biased left the store 3.7 minutes earlier on average. Minorities were not more likely to leave before the end of their shift when working with biased managers. But, they were substantially less likely to stay after.²³ Biased managers may be less likely to ask minority workers to stay late or minorities may be less likely to agree when working with biased managers.

While we do not have information on workers' pay, we can use estimates of time spent at work to estimate how much more minorities would earn (relative to majorities) if they worked only with unbiased managers. Re-estimating the minutes worked regression including days that workers were absent suggests that, for each standard deviation increase in manager bias minorities spend 8 fewer minutes per day at work. Eliminating manager bias would thus increase the time minorities spent at work and their pay by 2.5% (before correcting for measurement error).

4.2 Performance while at Work

We now examine the effect of manager bias on minority performance while at work, first considering the number of articles scanned per minute. This is one of the performance metrics over which workers have the most control and the firm cares most about. In each store, a list of workers' average articles scanned per minute is posted in the break room each week.

On average, CPs scan 18.5 articles per minute. 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 the few occasions when they have to work Sundays. Workers also scan articles more slowly on shifts that begin in the early morning.

Before we show our causal estimates, Figure I plots the relationship between manager bias and articles scanned per minute for minorities (Panel A) and majorities (Panel B). Each point plots the average articles scanned per minute on days that workers faced the indicated level of bias. Manager IAT scores are aggregated into bins of 0.2 and the size of the point indicates the number of observations in each bin. These graphs do not remove any individual

²³Consistent with this, when we include in our regressions the bias of the manager on duty at shift end instead of the day's average manager bias, the effect of manager bias on minutes worked more than doubles to -6.71 (without controls) and -7.18 (with our full controls). However, neither of these estimates is significant at conventional levels, with p-values of 0.132 and 0.114, respectively. Appendix Figure A.1 shows the measured effect of manager bias on minorities staying different lengths of time after the shift ends. Manager bias appears to decrease the number of minutes minorities stay after their shifts throughout the distribution. For example, it decreases the probability that a minority will stay at least an hour and a half after her shift ends by 1.5 percentage points. If these minorities who would have stayed at least 90 minutes after the end of their shifts instead leave at shift end, a decrease in these longer shifts would account for approximately half of the overtime effect.

or store fixed effects, so differences in scanning speed may result from cross-store differences in the types of workers hired or items purchased, instead of manager bias. Nonetheless, these graphs tell the same story as the regressions: minorities tend to scan slower when working with more-biased managers. Majority workers appear to scan more quickly when working with more-biased managers, but the effect for majorities is smaller than the effect for minorities.

Panel A of Table 4 replicates the format of Table 3's Panel A, showing the effect of manager bias on scanning speed. 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 were scheduled to work with), these regressions can be thought of as the reduced form for instrumental variables regressions in which the bias of the scheduled manager instruments for the bias of the manager on duty. These two bias measures have a high correlation (0.93).

Being scheduled to work with a manager one standard deviation more biased leads the average minority worker to scan 0.28 fewer items per minute (Table 4, Panel A). (The standard deviation of articles per minute is 2.9.) Unreported regressions show that manager bias does not appear to induce minorities to perform extremely poorly (in the bottom 15%), but otherwise it affects 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. The coefficients indicate that biased managers may cause majority workers to scan more quickly, though this effect is only significant at the 10% level and in only one of the specifications.

Panel B of Table 4 investigates the effect of manager bias on inter-customer time: the amount of time that a cashier spends between finishing one customer's transaction and starting to scan the next customer's items. While this is not an oft-discussed performance metric in the store, it directly affects the speed at which the lines move. On average, workers spend just under thirty seconds between customers. Working with a manager one standard deviation more biased leads minority workers to spend about 1.2 more seconds (one tenth of a standard deviation or 4% longer) between customers.

In contrast, there is no effect of manager bias on payment time: the time between the scanning of a customer's final item and the end of the transaction, during which the customer is paying (Appendix Table A.6). Payment time largely may largely depend on the customer.

Our results on the effect of manager bias are not driven by the other manager characteristics in our data: manager position in the firm, age, or gender (Appendix Table A.7). Appendix Table A.8 shows the results are robust to using 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 Sub-Saharan African ori-

gin), while Panel B does the reverse. The effects of manager bias on workers of North and Sub-Saharan African origins 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 workers of indeterminate, mixed, or other origins). The results are virtually unchanged. The results are also robust to eliminating managers who are unbiased or biased in favor of minorities and to using a wild cluster bootstrap, a method suggested for small numbers of clusters (Cameron, Gelbach, and Miller 2008).²⁴

An alternative explanation for our results is that CPs' performance does not respond to managers, but responds instead to senior cashiers who tend to work disproportionately with certain managers. More-biased workers may sort towards more-biased managers or senior minority workers may sort away from more-biased managers.²⁵ It is difficult to test this theory directly since we do not have data on senior cashiers (just managers and CPs). However, we think this explanation is unlikely. While senior cashiers have some control over their schedules, they can only submit preferences over the times they work, not the people they work with. A worker could attempt to work with a particular manager by requesting certain shifts that the manager prefers. For example, if a biased manager tended to work Monday mornings, but not Tuesday afternoons, biased senior cashiers who wanted to work with this manager could request these times. There is a limit to how much workers can control their schedules: because the firm values everyone doing their "fair share" of different kinds of shifts, it has included this as a constraint in the assignment algorithm.

Appendix Table A.9 shows that the measured effects of manager bias on minority performance are virtually identical if we control for shift (day of the week \times morning or evening) within store and thus the likely sorting of senior cashiers. While senior workers could sort towards managers in a more sophisticated way, the fact that this most likely method explains none of the effect of manager bias suggests that senior cashier sorting does not drive our results.

Finally, in Appendix Table A.10, we examine heterogeneity in the impact of biased managers on minority workers' performance. We show how the impact of bias evolves over the contract (Panel A) and how it differs in more- and less-diverse stores (Panel B). We focus on articles scanned per minute because it is important to the store and has continuous variation.

²⁴Using a wild cluster bootstrap, the p-values for the coefficients on the minority worker \times manager bias term in the first columns of Tables 3 and 4 are 0.016 (absences), 0.009 (minutes worked), 0.049 (articles scanned per minute), and 0.093 (inter-customer time).

²⁵Even though we do not think this describes what is happening in the stores, this alternative explanation might have similar implications to manager bias directly depressing minority performance. If minority performance is negatively affected by more-biased senior workers, then the bias of more-senior store personnel would still be harming minority workers. If senior minority workers prefer not to work with biased managers, then this explanation would still include manager bias affecting minorities' work patterns.

The two columns in Panel A show the results of estimating Equation (2), where observations are limited to promotions for which we have early- and late-week data, respectively. Appendix Table A.11 shows that promotions for which we have early- and late-week data have similar worker and manager characteristics. In fact, many stores are included in both regressions because they have one promotion with early-week and one promotion with late-week data.

CPs become more productive over time. In the latter part of the contract, workers scan two more articles per minute than they do in the early part. The negative effect of manager bias on minority performance also appears to increase over time. The effect of the bias is estimated to be twice as large in the latter weeks of the contract as it is in the early weeks, though the coefficients are not significantly different.

Panel B shows the effect of manager bias by workforce diversity. While we do not have data on any non-CP workers, our survey asked managers about the diversity of their stores. Since it is illegal to directly solicit data on ethnicity, we asked managers to report the fraction of workers they had managed from “sensitive urban zones” (ZUS). ZUS have high concentrations of immigrants and first generation citizens (ONZUS 2011) and ZUS residence has been used as a signal of minority status in a study of anonymous resumes (Behaghel, Crépon, and Barbanchon 2015). We average manager responses within a store to create a proxy for store diversity. The reported fraction of workers who come from ZUS is positively correlated with the fraction of minority CPs in our data. Panel B shows that the negative effect of manager bias is driven by stores with relatively little diversity. Appendix Table A.11 shows that there are other differences between more- and less-diverse stores, though none of these differences seem to drive the effect of manager bias.²⁶

Why do firms employ biased managers given that biased managers negatively impact minority performance? One explanation is that biased managers do not decrease average performance. Regressing worker performance on manager IAT scores indicates that biased managers don’t generate worse *average* worker performance. While biased managers depress the performance of minority workers, minorities are a small share (28%) of the entire workforce. For three out of the four main outcomes, point estimates suggest that biased managers (insignificantly) improve majority worker performance. Because managers can choose when they work, the estimates of manager bias on overall worker performance may include the effects of different store conditions and do not necessarily isolate managers’ effects on worker performance. Nevertheless, adding controls for shift characteristics does not change the estimates.

²⁶More-diverse stores have more male workers and male managers (both significant at the 10% level) and younger managers (significant at the 1% level) than do less-diverse stores.

5 Evidence on Mechanisms

In this section, we discuss several theories of discrimination and their predictions for our context. We then use worker survey and administrative data to test these predictions. We break these theories into two types: (1) theories that explain why minorities perform worse when working with biased managers and (2) theories that explain the firm’s hiring decisions, and discuss them in turn.

5.1 Effects of Manager Bias on Minority Performance

We consider only theories that can explain why minorities are absent more often, leave work earlier, and have worse on-the-job performance on days when they work with more-biased managers. For example, while manager bias may impede minority skill development (see, for example, Lundberg and Startz 1983 and Coate and Loury 1993), our study cannot assess whether this occurs. Even if it did occur, minorities would have similar skills on days when they worked with both more- and less-biased managers.

5.1.1 Theories

Animus

Biased managers may simply dislike minorities. Animus could lead biased managers to treat minorities badly and give them unpleasant tasks, thereby causing minorities to dislike coming to work. Minorities would be absent more often and leave work earlier. Animus could also impede minority on-the-job performance. To test whether manager animus drives our results, we assess whether, in the worker survey, minorities report that biased managers liked them less and assigned them to unpleasant tasks.

Less Interaction

Individuals with higher implicit biases towards a minority group have been found to spend less time talking to, have more hesitation in speaking to, and act less friendly towards minority group members (McConnell and Leibold 2001; Dovidio, Kawakami, and Gaertner 2002; and Hebl et al. 2002). Biased individuals may be uncomfortable interacting with minorities or their actions can be driven by a desire to seem unprejudiced (see Dovidio and Gaertner 2008 for a summary of the literature on aversive racism). Alternatively, biased managers may believe minorities are so unproductive that there are low returns to expending managerial effort on them. As a result, biased managers may spend less time at minority workers’ stations. When they need a task accomplished – even an unpleasant one – they may be less likely to ask a minority. Minorities may realize which managers are paying

less attention to them and exert less effort when these managers are on duty. We use the worker survey linked to administrative data to test whether (1) biased managers interact less with minority workers, (2) in our context, more worker-manager interaction correlates with better worker performance and (3) biased managers are less likely to assign minorities to tasks whose assignment requires interaction with workers.

Self-Stereotyping or Stereotype Threat

Under self-stereotyping (Coffman 2014), workers' expectations about their group's suitability for a given task affect their performance. Here, minorities might think that workplace environments (even relatively low-skill ones) are not environments where minorities thrive. Biased managers may activate these negative stereotypes. Relatedly, under stereotype threat (Steele and Aronson 1995), the risk of confirming negative group stereotypes leads minorities to become anxious and perform worse. To assess the extent to which managers differentially activate self-stereotypes or trigger stereotype threat, we asked workers the extent to which different managers made them feel confident in their abilities. While this tests for conscious activation of stereotypes, it would not capture stereotypes activated subconsciously.

5.1.2 Evidence

The evidence seems inconsistent with manager animus. Minorities do not perceive biased managers as disliking them. In the worker survey, we asked workers to rank their managers on the extent to which the manager liked the worker and the manager was likely to recommend the worker for promotion. We use workers' answers to order managers from the 1 (the lowest ranked) to N (the highest ranked).²⁷ The first two columns of Table 5's Panel A show the results of estimating Equation (2) with workers' rankings of their managers on these dimensions as the dependent variables. While neither coefficient is significant, both are positive, suggesting that, if anything, minorities perceived biased managers as liking them better. We also asked workers to rate how much they enjoyed working with each manager. Minorities again rated biased managers insignificantly more positively.

To assess whether biased managers activated minorities' negative self-stereotypes or triggered stereotype threat, we asked workers which managers initially made the worker feel most confident in their abilities. There is no evidence that biased managers made minorities anxious about confirming stereotypes or activated self-stereotypes of poor performance: minorities rated biased managers as making them feel insignificantly more confident about their abilities (Table 5, Panel A). Though, as discussed above, this does not rule out subconscious stereotype activation.

²⁷We eliminate from this ranking managers workers indicated they did not remember since workers almost never ranked these managers. We analyze whether workers remembered their managers separately.

Panel B analyzes task assignment. Animus would lead biased managers to assign minorities to unpleasant tasks more often, while if biased managers avoid interacting with minorities, they would assign minorities to additional tasks less often, regardless of task pleasantness. The first two columns of Table 5’s Panel B examine workers’ register assignments and the desirability of their break times.²⁸ Since all cashiers need to be assigned to a register and given breaks, these two assignments test for animus, but not whether biased managers avoid interacting with minorities. We find that biased managers do not assign minorities unpleasant registers or break times.

The final column considers assignment to cleaning duties. Telling a worker to shut down her register and start cleaning requires interaction. Moreover, cleaning is typically considered to be cashiers’ least pleasant task. Consistent with biased managers avoiding minorities but inconsistent with animus, biased managers are significantly *less* likely to assign minorities to cleaning duties.²⁹ The interaction theory can also explain why biased managers may be less likely to ask minorities to stay late as well as why the effect of manager bias is larger in stores with less diversity: it may be more difficult for managers to avoid minority workers when a larger share of the workforce is minority.

To further test the interaction theory, we asked CPs to rate the extent to which they remembered each manager. We utilize remembering a manager (ranking the manager at least two out of 10) as an indicator of the amount of interaction the worker and manager had. Workers performed much better when working with managers with whom they interacted more. The first column of Table 6’s Panel A shows that workers scanned 1.5 more articles per minute when working with a manager they later remembered. However, workers did not perform better when working with managers they had been scheduled to work with more often (Column 2), nor does the effect of remembering the manager on worker performance decrease when we control for the amount of time spent working together (Column 3) or manager fixed effects (Column 4). Taken together, this suggests that worker-manager interaction within a shift leads workers to perform substantially better. Mas and Moretti (2009) similarly finds that cashiers exert more effort when their performance is being noticed by coworkers they value.

Minorities report interacting less with biased managers: they were about 1.5 percentage points less likely to remember a manager one standard deviation more biased (Table 6, Panel B).³⁰ The final columns in Panel B suggest how worker-manager interaction affects the

²⁸Managers choose when workers can take their breaks, but not how much break time workers receive, which is stipulated in workers’ contracts.

²⁹Managers worried about appearing biased might be particularly concerned with the optics of assigning minorities to clean.

³⁰Unsurprisingly, workers were more likely to remember managers they were scheduled to work with more.

measured impact of manager bias on minority performance. While the results are imprecise, when we limit the sample to days the worker remembered the manager, the effect of manager bias on minority performance falls by 25%.

Summarizing our results, we find the strongest evidence for the theory that biased managers interact less with minority workers and assign them to new tasks – even unpleasant ones – less often. This may be because they feel less comfortable around minorities, they are concerned with appearing biased, or they believe there is a low return to expending effort managing minorities. We find no evidence of animus: minorities do not report that biased managers disliked them or assigned them to unpleasant tasks. We have less clear evidence on whether self-stereotyping or stereotype threat plays a role: we do not find that biased managers made minorities consciously anxious about their abilities, but this does not rule out that they activated minority stereotypes on a subconscious level.

5.2 Hiring Decisions

Comparing the performance of minority and majority workers allows us to shed light on whether the firm engages in statistical or Beckerian (taste-based) discrimination in the hiring process.

5.2.1 Theories

Statistical Discrimination

By making minorities less productive, manager bias may lead to statistical discrimination in hiring (Phelps 1972; Arrow 1973). Under statistical discrimination, the firm uses workers' observable characteristics and minority status to infer worker productivity. Wages are fixed and cannot depend on performance, so the firm hires the workers with the highest expected productivities.³¹ Even if minorities and majorities with the same observable characteristics are equally productive in the absence of manager bias, because manager bias depresses minority productivity, the firm infers that minority workers are less productive than majority workers with the same characteristics. To be hired, minorities would need better qualifications than hired majority workers and hired minorities would be more productive than hired majorities when not exposed to manager bias.

Beckerian Discrimination

Under taste-based discrimination, the firm doesn't necessarily have uncertainty about worker productivity. Instead, the firm (or its employees) faces a utility cost of employing minority workers. Since all hired workers must be paid the same wage, hired minorities need

³¹We assume the firm is risk neutral.

to have higher average productivity than hired majorities to compensate the firm for hiring them.

5.2.2 Evidence

Our evidence is most consistent with statistical discrimination, though we cannot fully rule out that the firm engages in taste-based discrimination.

We compare minority and majority workers on a summary measure of worker performance, the number of customers served, as well as on the individual performance metrics. To construct this summary measure, we combine the time spent at work with a worker’s average articles scanned per minute, inter-customer time, and payment time. We assume that the average customer has 25 items, though our results are not very sensitive to this assumption. We also assume that cashiers spend all day at their registers receiving customers. To the extent that this is not true, but relative performance differences are similar on other tasks, this can still be thought of as a summary measure of performance.

Table 7 compares minority and majority performance by presenting estimates from the regression

$$y_{ist} = \alpha + \beta \text{minority}_i + \delta_s + \varepsilon_{ist} \quad (3)$$

where y_{ist} is a metric of performance of worker i in store s on date t . As before, δ_s are store fixed effects. The coefficient of interest, β , shows how minority workers’ performance compares to the performance of majority workers in the same store. Panel A includes all days. It shows that minority workers’ average performance is statistically indistinguishable from that of majority workers. We estimate that the average majority worker serves 162 customers per day. The average minority serves an additional 2.8 customers, a difference which is far from significant, and which places the average minority worker at the 53rd percentile of average worker performance. The similar average performance of minority and majority workers is (weak) evidence against taste-based discrimination. We do not find that hired minorities perform substantially better than hired majority workers and can rule out that they perform more than 4% better on average.

Panel B compares minority and majority workers’ performance on days when they work with unbiased managers. (Because most managers are biased, minorities only work with unbiased managers on a small fraction of days.) On days when workers spend at least half their time with unbiased managers, minority workers perform substantially better than non-minority workers. They are approximately half as likely to be absent, they scan 0.75 more articles per minute, and they take two fewer seconds between customers. On days with unbiased managers, the average minority cashier serves 14 customers more per day than does

the average majority. This 9% better performance places the average minority working with an unbiased manager at the 79th percentile of worker performance. That overall minority and majority workers perform similarly, but minorities perform substantially better when not exposed to manager bias is consistent with statistical discrimination.³²

These facts would not be evidence of statistical discrimination if minority workers were simply intertemporally substituting effort towards days with unbiased managers. If they were fully intertemporally substituting, manager bias would not affect average minority performance, it would just lead minorities to perform worse on some days and better on others. That minorities performed better on days without bias would not indicate their higher general productivity. However, we provide two tests that suggest that minorities are not simply intertemporally substituting effort. Table 8 shows the effect of working with a more-biased manager on one day on the worker’s performance in the rest of the week.³³ Under intertemporal substitution, working with a more-biased manager on one day should lead minorities to perform better in the rest of the week. We find no evidence that this is the case. In fact, minorities scan items significantly slower in the rest of the week when they spend one day with a more-biased manager.

We can also look at workers’ response to manager bias aggregated over longer periods (Appendix Table A.12). If a worker is intertemporally substituting her effort within a given period (e.g., at the week or two-week level), performance should be uncorrelated with manager bias at that level of aggregation. However, we find no evidence that the impact of manager bias is attenuated when performance is aggregated over longer periods. For absences, the measured effect of manager bias is relatively constant with the level of aggregation, though it is no longer statistically significant once the data is aggregated. For scanning speed, the measured effect increases with the level of aggregation. Consistent with the results in Table 8, this may indicate that there are some cumulative effects of manager bias on scanning speed.

6 Conclusion

Working with biased managers leads minorities to perform more poorly. When scheduled to work with more-biased managers, minority cashiers are absent more often and they leave

³²Statistical and taste-based discrimination predict that hired minorities should have better observable characteristics than hired majorities. We have only one pre-employment characteristic that should have a clear relationship with productivity: educational attainment. Minorities are more than three times as likely to have above high school education as majority workers (16% vs. 5%) and this difference remains nearly as large (9.1 percentage points) when store fixed effects are added.

³³We control for manager bias in the rest of the week and its interaction with the worker’s minority status to eliminate the effects of intra-week correlation in manager bias.

work earlier. This depresses minority wages since workers are paid based on time worked. Minorities also scan items more slowly and take more time between customers when working with biased managers. Biased managers do not appear to treat minorities poorly. Instead, they seem to simply interact less with minorities, leading these workers to exert less effort. By making minorities less productive, manager bias appears to generate statistical discrimination in hiring.

These results come from one setting: entry-level cashiers in a large French grocery store chain. However, they may be applicable to many other workplace settings. In our setting, biased managers' discomfort with minorities can lead them to monitor minorities less, assign minorities to new tasks less frequently, and not ask minorities to stay late. In other settings, interacting less with minority employees may have larger consequences if it leads biased managers to train, mentor, advise, or challenge minorities less.

Our results raise the question of the type of policy responses that could be used to ameliorate the impacts of manager bias on minority workers. One set of potential policies would aim to directly reduce implicit bias. Beaman et al. (2009) finds that having female leaders reduces implicit bias against women. Outside of the workplace, Rao (2014) and Boisjoly et al. (2006) find that exposure to a group can reduce bias against it. Another set of potential policies would attempt to mitigate the effect of manager bias by directly targeting manager actions. For example, these interventions could encourage managers to interact with all workers equally or provide more specific guidelines about how to manage workers. Investigating the effects of such policies is an interesting question for future research.

Tables and figures

Table 1. Descriptive Statistics

	Administrative Data Sample		Worker Survey Sample	
	All Observations from Included Stores	Regression Sample (Observations with Manager IAT Scores)	All Survey Respondents	Survey Respondent Regression Sample
A. Worker Characteristics				
Minority	28%	28%	29%	25%
Male	6.9%	7.4%	7.7%	7.3%
Age			29.9	30.1
Number of Previous Jobs			3.9	4.0
Less than High School Education			58%	61%
High School Degree			35%	32%
More than High School Education			7%	7%
Number of Workers	218	204	310	178
B. Manager Characteristics				
Minority	6%	8%		
Male	10%	7%		
Level 4 (High Position)	18%	18%		
Age	41.1	41.1		
Average IAT Score (in Standard Devs)		1.35		
Moderate to Severe Bias		66%		
Slight Bias		20%		
Little to No Bias		9%		
Preference for Minorities		4%		
Number of Managers	154	119		
C. Shift Characteristics				
Scheduled Days per Week	4.2	4.2		
Scheduled Hours per Day	7.2	7.2		
Absent	1.8%	1.6%		
Minutes Worked in Excess of Schedule	-0.31	-0.06		
Articles Scanned per Minute	18.5	18.5		
Inter-Customer Time (Seconds)	29.2	28.7		
Payment Time (Seconds)	50.7	50.8		
Number of Shifts	5 099	4 371		
Number of Stores	34	34	70	51

Note: The first two columns of data provide descriptive statistics for the sample for whom we have administrative data. The first includes all observations from the 34 included stores, while the second includes only observations for which we have the manager's IAT score. The final two columns provide descriptive statistics for the worker survey sample. The first includes all survey respondents, while the second includes only those workers for whom we either have managers' IAT scores or performance data and are thus included in the analysis. Level 4 managers have a higher position in the store than the remaining managers. Manager age is reported as of January 1, 2012. *Moderate to Severe Bias* is defined as having a raw IAT score above 0.35, *Slight Bias* is defined as having a raw IAT score between 0.15 and 0.35, *Little to No Bias* is defined as having a raw IAT score between -0.15 and 0.15, and *Preference for Minorities* is defined as having a raw IAT score below -0.15.

Table 2. Exogeneity of Scheduled Shifts

	Manager Bias	Minority Manager	Male Manager	Level 4 Manager	Total Managers	Articles per Minute in Other Stores on Date	Shift Includes Early Morning	Shift Includes Late Evening	Total Hours	Split Shift
<u>A. Minority Workers</u>										
Minority Worker	0.005 (0.022)	0.000 (0.003)	-0.011 (0.009)	-0.001 (0.004)	-0.025 (0.043)	0.045 (0.086)	0.014 (0.015)	0.021 (0.013)	0.017 (0.034)	-0.000 (0.011)
<u>B. Minority Workers and Manager Bias</u>										
Minority Worker × Manager Bias		0.009 (0.007)	-0.002 (0.008)	0.008 (0.010)	-0.013 (0.024)	0.049 (0.044)	-0.004 (0.008)	0.007 (0.007)	0.038 (0.037)	0.002 (0.012)
Minority Worker		-0.009 (0.009)	-0.009 (0.007)	-0.010 (0.012)	-0.011 (0.034)	-0.007 (0.074)	0.019 (0.017)	0.013 (0.013)	-0.024 (0.049)	-0.002 (0.020)
Manager Bias		-0.032 (0.021)	0.015 (0.025)	0.049 (0.033)	0.066 (0.048)	-0.043 (0.093)	-0.002 (0.011)	-0.003 (0.019)	-0.058 (0.062)	-0.017 (0.022)
Observations	4,371	4,371	4,371	4,371	4,371	4,238	4,371	4,371	4,368	4,371
Dependent Variable										
Mean	1.13	0.060	0.114	0.171	2.74	18.23	0.141	0.580	7.22	0.465
Store FE.'s	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: Each column in Panel A shows the result of regressing the dependent variable indicated by the column heading on an indicator for the worker being a minority. Each column in Panel B shows the results of regressing the same dependent variable on a dummy for the worker's minority status, the manager's IAT score (in standard deviation terms), and the interaction of the worker's minority status and the manager's IAT score. Both the dependent variables and the manager's IAT score are based on the store's schedule, not actual realizations. For example, *Shift Includes Early Morning* is a dummy variable for the shift being scheduled to start at 9 am or earlier, regardless of whether the worker arrived by that time. *Shift Includes Late Evening* is an indicator for the shift being scheduled to end at 8 pm or later. *Manager Bias* is the manager's IAT score, while *Minority Manager*, *Male Manager*, and *Level 4 Manager* are indicators for a manager being a minority, being male, and having a high-level management position, respectively. When workers are scheduled to work with more than one manager, manager variables are averages, weighted by the amount of time workers were scheduled to work with each manager. Observations are worker-days and standard errors are clustered at the store level.

Table 3. Effect of Manager Bias on Time Spent at Work

A. Dependent Variable: Absence Indicator				
Minority Worker × Manager Bias	0.0098** (0.0039)	0.0095** (0.0040)	0.0117*** (0.0042)	0.0118*** (0.0043)
Manager Bias	-0.0021 (0.0031)	-0.0021 (0.0032)	-0.0050 (0.0040)	-0.0052 (0.0042)
Minority Worker × Minority Manager				0.0081 (0.0972)
Minority Manager				-0.0057 (0.0153)
Observations	4,371	4,371	4,371	4,371
Dependent Variable Mean	0.0162	0.0162	0.0162	0.0162
R-squared	0.0005	0.0031	0.0835	0.0835
B. Dependent Variable: Minutes Worked in Excess of Schedule				
Minority Worker × Manager Bias	-3.295** (1.550)	-3.279** (1.588)	-3.327* (1.687)	-3.237* (1.678)
Manager Bias	-0.002 (1.141)	-0.002 (1.167)	-0.005 (0.969)	-0.005 (1.009)
Minority Worker × Minority Manager				0.349 (10.501)
Minority Manager				-3.712 (4.592)
Observations	4,163	4,163	4,163	4,163
Dependent Variable Mean	-0.068	-0.068	-0.068	-0.068
R-squared	0.001	0.008	0.129	0.129
Individual F.E.'s	Yes	Yes	Yes	Yes
Day of the Week F.E.'s	No	Yes	No	No
Morning/Evening F.E.'s	No	Yes	Yes	Yes
Date F.E.'s	No	No	Yes	Yes

Note: Each column in each panel shows the result of regressing the dependent variable on the interaction of the worker's minority status and the manager's IAT score (in standard deviation terms), controlling for the manager's IAT score and worker fixed effects. The dependent variable in Panel A is an indicator for the worker being absent. The dependent variable in Panel B is the number of minutes worked in excess of the number of minutes the worker was scheduled to work. The first column includes no additional controls. The second column adds day of the week fixed effects, an indicator for the shift starting at 9 am or earlier, and an indicator for the shift ending at 8 pm or later. The third column includes date fixed effects and drops the day of the week fixed effects. The last column adds a dummy for the manager being a minority and the interaction of the worker's and the manager's minority status. Observations are worker-days and standard errors are clustered at the store level. *, **, *** denote significance at the 10%, 5% and 1% levels,

Table 4. Effect of Manager Bias on Performance at Work

<u>A. Dependent Variable: Articles Scanned per Minute</u>				
Minority Worker × Manager Bias	-0.276** (0.109)	-0.279** (0.111)	-0.233** (0.108)	-0.249** (0.111)
Manager Bias	0.140* (0.083)	0.140 (0.083)	0.080 (0.065)	0.102 (0.073)
Observations	3,601	3,601	3,601	3,601
Dependent Variable Mean	18.53	18.53	18.53	18.53
R-squared	0.001	0.013	0.195	0.195
<u>B. Dependent Variable: Inter-Customer Time (Seconds)</u>				
Minority Worker × Manager Bias	1.213** (0.590)	1.228** (0.553)	1.417** (0.649)	1.360** (0.665)
Manager Bias	-0.648 (0.386)	-0.571 (0.376)	-0.656 (0.521)	-0.580 (0.534)
Observations	3,287	3,287	3,287	3,287
Dependent Variable Mean	28.70	28.70	28.70	28.70
R-squared	0.001	0.013	0.195	0.195
Individual F.E.'s	Yes	Yes	Yes	Yes
Day of the Week F.E.'s	No	Yes	No	No
Morning/Evening F.E.'s	No	Yes	Yes	Yes
Date F.E.'s	No	No	Yes	Yes
Manager Minority Variables	No	No	No	Yes

Note: Each regression shows the result of regressing the dependent variable on the interaction of the worker's minority status and the manager's IAT score (in standard deviation terms), controlling for the manager's IAT score and worker fixed effects. The dependent variables are the number of articles per minute scanned (Panel A) and the average number of seconds between finishing one customer's transaction and starting to scan the next customer's items (Panel B). The first column includes no controls. The second column adds day of the week fixed effects, an indicator for the shift starting at 9 am or earlier, and an indicator for the shift ending at 8 pm or later. The third column includes date fixed effects and drops the day of the week fixed effects. The last column adds a dummy for the manager being a minority and the interaction of the worker's and the manager's minority status. Observations are worker-days and standard errors are clustered at the store level. *, ** denote significance at the 10% and 5% levels, respectively.

Table 5. Worker-Manager Affection and Task Assignment

	<u>A. Worker-Manager Affection</u>			
	Manager Liked You Best	Manager Most Likely to Recommend You for Promotion	You Enjoyed Working with Manager Best	Manager Initially Made You Feel Most Confident
Minority Worker × Manager Bias	0.019 (0.246)	0.078 (0.212)	0.243 (0.234)	0.194 (0.196)
Manager Bias	0.152 (0.131)	0.251* (0.148)	-0.061 (0.162)	0.134 (0.127)
Observations	3,036	2,862	3,209	3,189
Dependent Variable Mean	3.991	4.053	4.062	4.073
R-squared	0.015	0.042	0.010	0.026
	<u>B. Task Assignment</u>			
	Manager Assigned to Preferred Register Type	Manager Assigned Best Breaks	Management of Lines and Customer Flows Encouraged Performance	Manager Assigned to Fewest Cleaning Duties
Minority Worker × Manager Bias	-0.035 (0.391)	0.146 (0.469)	-0.153 (0.308)	0.673*** (0.189)
Manager Bias	0.021 (0.157)	-0.083 (0.146)	0.129 (0.137)	-0.276 (0.182)
Observations	2,288	2,553	2,864	2,235
Dependent Variable Mean	4.010	3.922	4.215	3.373
R-squared	0.002	0.008	0.018	0.045

Note: Each column in each panel shows the result of regressing the dependent variable on the interaction of the worker's minority status and the manager's IAT score (in standard deviation terms), controlling for the manager's IAT score and worker fixed effects. The dependent variable is the worker's ranking of the manager on the question indicated by the column heading. This ranking ranges from 1 (the lowest ranked manager) to N (the highest ranked manager), where N is the number of managers the worker had. Observations are worker-days and standard errors are clustered at the store level. *, *** denote significance at the 10% and 1% levels, respectively.

Table 6. Worker-Manager Interaction

Panel A. Effect of Worker-Manager Interaction on Performance				
Dependent Variable: Articles Scanned per Minute				
Remembers Manager (Indicator)	1.510** (0.635)		1.587** (0.630)	2.053*** (0.744)
Fraction of Time Scheduled with Manager		-1.172 (1.352)	1.724 (3.638)	4.021 (3.828)
(Fraction of Time Scheduled with Manager) ²			-4.454 (4.886)	-6.603 (5.407)
Individual F.E.'s	Yes	Yes	Yes	Yes
Manager F.E.'s	No	No	No	Yes
Observations	1,885	1,885	1,885	1,885
Dependent Variable Mean	18.42	18.42	18.42	18.42
R-squared	0.008	0.001	0.010	0.095
Panel B. Minority Status, Manager Bias, and Worker-Manager Interaction				
	Dependent Variable: Remembers Manager (Indicator)		Dependent Variable: Articles Scanned per Minute	
	All Days		All Days	Days where Worker Remembers Manager
Minority Worker × Manager Bias	-0.0152* (0.0086)		-0.415* (0.209)	-0.311 (0.314)
Manager Bias	0.0190* (0.0097)		0.271** (0.114)	0.203** (0.095)
Fraction of Time Scheduled with Manager	0.6362* (0.3351)		-1.932 (4.159)	-5.115 (4.161)
(Fraction of Time Scheduled with Manager) ²	-0.5605 (0.3981)		1.017 (6.270)	4.114 (5.757)
Individual F.E.'s	Yes		Yes	Yes
Observations	3,958		1,584	1,317
Dependent Variable Mean	0.932		18.52	18.66
R-squared	0.017		0.006	0.005

Note: Each regression in Panel A shows the results of regressing articles scanned per minute on the variables listed in the left-most column, controlling for worker fixed effects. *Remembers Manager* is an indicator for the worker reporting in the worker survey that she remembered the manager she was scheduled to work with that day. *Fraction of Time Scheduled with Manager* is a number between 0 and 1. It is the fraction of the worker's time in the administrative data that she was scheduled to work with the given manager, averaged over all working days. The first column of Panel B regresses an indicator for whether the worker remembered the manager on the interaction of the worker's minority status and the manager's IAT score (in standard deviation terms), controlling for the manager's IAT score, the fraction of time in the administrative data the worker spent with the manager, this fraction squared, and worker fixed effects. The next column in the panel regresses articles per minute scanned on these same variables. The final column in the panel replicates the previous column, but eliminates days where the worker did not remember the manager. Throughout, observations are worker-days and standard errors are clustered at the store level. *, **, *** denote significance at the 10%, 5%, and 1% levels respectively.

Table 7. Comparison of Minority and Non-Minority Performance

	Absence Indicator	Minutes Worked in Excess of Schedule	Articles Scanned per Minute	Inter-Customer Time (Seconds)	Estimated Customers Served
			<u>A. All Days</u>		
Minority Worker	-0.0041 (0.0072)	0.522 (2.213)	0.282 (0.329)	0.504 (0.719)	2.80 (2.02)
Non-Minority Mean	0.0187	-1.186	18.55	28.21	162
Observations	4,371	4,163	3,601	3,287	3,086
Store F.E.'s	Yes	Yes	Yes	Yes	Yes
			<u>B. Days with Unbiased Managers</u>		
Minority Worker	-0.0127* (0.0067)	2.572 (2.331)	0.745** (0.323)	-2.075* (1.113)	13.94** (4.84)
Non-Minority Mean	0.0267	-4.268	18.65	26.59	162
Observations	482	444	367	330	301
Store F.E.'s	Yes	Yes	Yes	Yes	Yes
			<u>C. Days with Biased Managers</u>		
Minority Worker	-0.0047 (0.0094)	0.271 (2.872)	0.006 (0.383)	0.936 (0.935)	2.21 (2.68)
Non-Minority Mean	0.0194	-1.106	18.65	27.94	162
Observations	3,474	3,319	2,832	2,555	2,395
Store F.E.'s	Yes	Yes	Yes	Yes	Yes
			<u>D. All Other Days</u>		
Minority Worker	0.0026 (0.0027)	0.379 (1.625)	1.292** (0.325)	0.178 (0.186)	-0.25 (2.05)
Non-Minority Mean	0.0036	0.928	17.69	31.88	156
Observations	445	429	422	421	410
Store F.E.'s	Yes	Yes	Yes	Yes	Yes

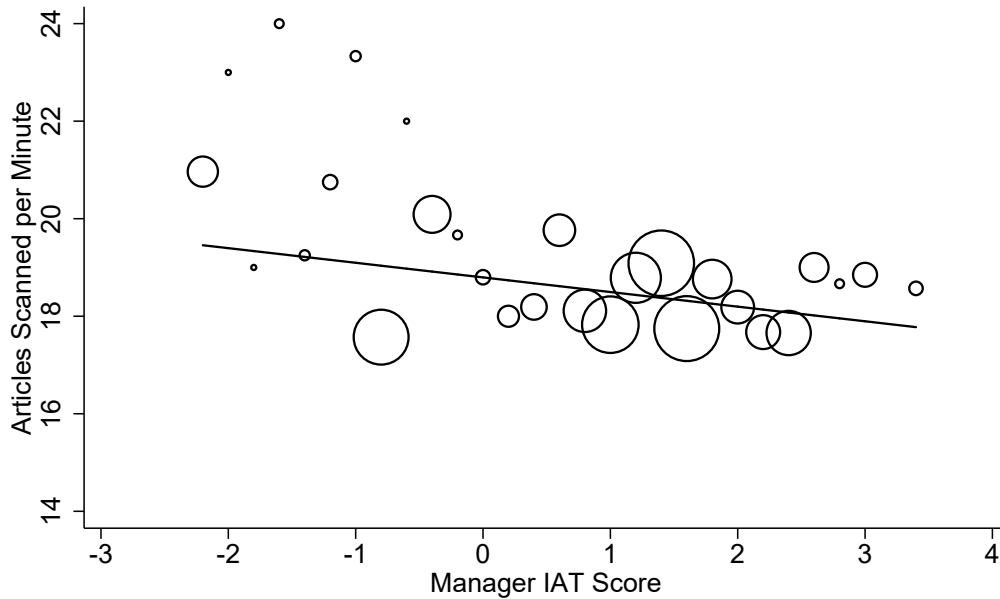
Note: Each column in each panel shows the result of a separate regression of the dependent variable indicated by the column on an indicator for the worker being a minority, controlling for store fixed effects. *Estimated Customers Served* is calculated under the assumptions that customers average 25 items. 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 a raw (unscaled) IAT score between -0.15 and 0.15. *Days with Biased Managers* are days where the worker spent more than 50% of the day with managers whose raw IAT score exceeds 0.15. Days where a worker spent more than 50% of her time with managers biased in favor of minorities (managers with raw IAT scores below -0.15) and days where a worker did not spend more than 50% of her time with managers in any bias category are included in Panel D as *All Other Days*. *, ** denote significance at the 10% and 5% levels, respectively.

Table 8. Effect of Manager Bias on Performance in Rest of the Week
 Dependent Variable: Average Performance in Rest of the Week

	Absence Indicator	Minutes Worked in Excess of Schedule	Articles Scanned per Minute	Inter-Customer Time (Seconds)
Minority Worker × Manager Bias	0.0023 (0.0020)	-0.628 (0.940)	-0.099** (0.041)	-0.336 (0.274)
Manager Bias	-0.0029 (0.0017)	1.185* (0.638)	0.063 (0.046)	-0.030 (0.203)
Observations	4 271	4 174	3 935	3 610
Dependent Variable Mean	0.0153	-0.053	18.43	29.04
R-squared	0.0010	0.0058	0.0061	0.0026


Notes: Each column shows the result of regressing the workers' average performance in the rest of the week on the interaction of the worker's minority status and the manager's IAT score (in standard deviation terms) on a given day. Regressions control for the manager's IAT score on that day and worker fixed effects. They also control for the average manager IAT score in the rest of the week and this score interacted with the worker's minority status. Standard errors are clustered at the store level. ** denotes significance at the 5% level.

Figure I.A. Manager Bias and Worker Performance
Minority Workers



Note: The size of each marker indicates the number of observations in the bin.

Figure I.B. Manager Bias and Worker Performance
Majority Workers



Note: The size of each marker indicates the number of observations in the bin.

Appendix tables and figures

Appendix Table A.1. Characteristics of Respondents and Non-Respondents to the Worker Survey

	All Workers			Minority Workers			Majority Workers		
	Respondents	Non-Respondents	p-value of Difference	Respondents	Non-Respondents	p-value of Difference	Respondents	Non-Respondents	p-value of Difference
Male	7.7%	8.1%	0.855	6.7%	6.8%	0.994	8.1%	9.0%	0.743
Absence Rate	0.018	0.028	0.423	0.016	0.006	0.364	0.016	0.037	0.235
Minutes Worked in Excess of Schedule	2.59	-2.51	0.064	2.92	1.29	0.765	2.51	-4.58	0.031
Articles Scanned per Minute	18.57	18.43	0.699	18.84	18.20	0.298	18.51	18.56	0.919
Inter-Customer Time (Seconds)	28.51	29.50	0.420	27.94	31.24	0.130	28.65	28.52	0.933
Payment Time (Seconds)	50.51	50.73	0.768	51.83	52.13	0.823	50.17	50.02	0.868
Paris Region	42%	49%	0.095	63%	59%	0.585	33%	41%	0.103
Municipality has Less than 25,000 Residents	33%	29%	0.330	28%	24%	0.544	35%	33%	0.639
Municipality has 25,000 to 75,000 Residents	49%	52%	0.435	54%	54%	0.984	47%	50%	0.457
Municipality has More than 75,000 Residents	18%	19%	0.877	18%	21%	0.524	19%	17%	0.696

Notes: The *Paris Region* or "Ile-de-France" is one of the 13 administrative regions in France. Municipality population data comes from 2013 Census data. P-values are calculated from t-tests.

Appendix Table A.2. Correlates of Manager IAT Scores
 Dependent Variable: Standardized Manager IAT Score

Minority	-0.441 (0.345)						-0.364 (0.388)	-0.135 (0.445)
Male		-0.078 (0.448)					0.109 (0.617)	0.092 (0.675)
Age			0.008 (0.010)				0.008 (0.012)	0.003 (0.018)
Level 4 (High Position)				0.019 (0.266)			-0.093 (0.290)	0.329 (0.366)
Fraction ZUS in Store					-0.005 (0.004)		-0.005 (0.007)	
Far Right Vote Share						0.005 (0.016)	-0.006 (0.016)	
Observations	119	119	119	119	110	119	110	119
R-squared	0.014	0.000	0.005	0.000	0.003	0.001	0.021	0.269
Dependent Variable Mean	1.346	1.346	1.346	1.346	1.410	1.346	1.410	1.346
Store F.E.'s	No	No	No	No	No	No	No	Yes

Note: Each column shows the results of regressing a manager's (standardized) IAT score on her characteristics. Level 4 managers have a higher position in the store than the remaining managers. Manager age is reported as of January 1, 2012. *Fraction ZUS in Store* is the fraction of workers that come from "sensitive urban zones," zones with high concentrations of immigrants and first generation citizens, and is measured on a scale of 1 to 100. *Far Right Vote Share* is the share of votes received by the Front National Party in the first round of the 2012 presidential election. This data comes from the French Ministry of the Interior. Standard errors are clustered at the store level.

Appendix Table A.3. Exogeneity of Scheduled Shifts
Robustness to Alternative Specifications

	Manager Bias	Minority Manager	Male Manager	Level 4 Manager	Total Managers	Articles per Minute in Other Stores on Date	Shift Includes Early Morning	Shift Includes Late Evening	Total Hours	Split Shift
<u>A. Minority Workers, Including Observations with No Manager IAT</u>										
Minority Worker	0.005 (0.022)	-0.009 (0.007)	-0.020 (0.014)	-0.008 (0.010)	-0.127 (0.115)	0.043 (0.072)	0.008 (0.013)	0.014 (0.012)	0.034 (0.029)	-0.001 (0.010)
Observations	4,371	5,099	5,099	5,099	5,099	4,945	5,099	5,099	5,094	5,099
Dependent Variable Mean	1.13	0.051	0.116	0.156	2.44	18.28	0.141	0.581	7.20	0.461
Store F.E.'s	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<u>B. Minority Workers and Manager Bias, with Worker Fixed Effects</u>										
Minority Worker × Manager Bias		0.034 (0.024)	0.024 (0.024)	0.031 (0.037)	-0.059 (0.050)	-0.057 (0.064)	-0.013 (0.021)	0.016 (0.025)	-0.077 (0.110)	-0.027 (0.031)
Manager Bias		-0.040 (0.026)	0.007 (0.023)	0.030 (0.039)	0.019 (0.052)	0.065 (0.055)	0.006 (0.013)	-0.015 (0.023)	-0.045 (0.067)	-0.013 (0.024)
Observations		4,371	4,371	4,371	4,371	4,238	4,371	4,371	4,368	4,371
Dependent Variable Mean		0.060	0.114	0.171	2.74	18.23	0.141	0.580	7.22	0.465
Individual F.E.'s		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: Panel A replicates Panel A of Table II, including observations where we do not have the manager's IAT score. Panel B replicates Panel B of Table II, replacing the store fixed effects and minority worker indicator with worker fixed effects. That is, Panel B shows the results of regressing the dependent variable indicated by the column on the manager's IAT score (in standard deviation terms) and the interaction of manager's IAT score and the minority worker indicator, controlling for worker fixed effects.

Appendix Table A.4. Exogeneity of CP Coworkers

	Scheduled to Work at Same Time	Minutes Scheduled to Work Together
Minority-Minority Pair	-0.015 (0.036)	-6.10 (11.66)
Minority-Majority Pair	-0.014 (0.022)	-5.73 (7.26)
Store F.E.'s	Yes	Yes
Observations	15 791	15 791
Mean Dep. Var.	0.594	186.7

Note: The table shows the result of regressing an indicator for whether two CPs in the same store were scheduled to work together on a given day (Column 1) and the number of minutes they were scheduled to work together on a given day (Column 2) on an indicator for both CPs being minorities and an indicator for one worker being a minority and the other being a majority. The omitted category is both workers being majorities. Observations are CP pair-days. Store fixed effects are included and standard errors are clustered at the store level.

Appendix Table A.5. Additional Results on Time Spent at Work

	Minutes Arrived Before Shift Start	Break Time (Minutes)	Minutes Stayed After Shift End	Stayed Until Shift End	Stayed at Least 10 Minutes After Shift End
Minority Worker × Manager Bias	1.617 (1.858)	1.081 (1.381)	-3.773** (1.674)	0.014 (0.024)	-0.041* (0.023)
Manager Bias	0.633 (1.358)	-0.698 (0.616)	-0.402 (1.027)	-0.005 (0.017)	0.003 (0.017)
Observations	4,163	4,163	4,163	4,163	4,163
Dependent Variable Mean	4.63	15.55	10.84	0.844	0.437
R-squared	0.121	0.136	0.101	0.129	0.133

Note: Each regression shows the result of regressing the dependent variable indicated by the column on the interaction of the worker's minority status and the manager's IAT score (in standard deviation terms), controlling for the manager's IAT score and worker fixed effects. The regressions additionally control for the manager's minority status, the interaction of the worker's and the manager's minority status, date fixed effects, and dummies for the shift starting at 9 am or earlier and ending at 8 pm or later. Observations are worker-days and standard errors are clustered at the store level. *, ** denote significance at the 10% and 5% levels, respectively.

Appendix Table A.6. Effect of Manager Bias on Payment Time
 Dependent Variable: Payment Time (Seconds)

Minority Worker × Manager Bias	-0.188 (0.457)	-0.091 (0.413)	-0.064 (0.361)	-0.011 (0.341)
Manager Bias	0.046 (0.314)	0.049 (0.281)	-0.453 (0.325)	-0.506* (0.290)
Observations	3,108	3,108	3,108	3,108
Dependent Variable Mean	50.77	50.77	50.77	50.77
R-squared	0.000	0.039	0.159	0.159
Individual F.E.'s	Yes	Yes	Yes	Yes
Day of the Week F.E.'s	No	Yes	No	No
Morning/Evening F.E.'s	No	Yes	Yes	Yes
Date F.E.'s	No	No	Yes	Yes
Manager Minority Variables	No	No	No	Yes

Note: Each regression shows the result of regressing the dependent variable on the interaction of the worker's minority status and the manager's IAT score (in standard deviation terms), controlling for the manager's IAT score and worker fixed effects. The first column includes no controls. The second column adds day of the week fixed effects, an indicator for the shift starting at 9 am or earlier, and an indicator for the shift ending at 8 pm or later. The third column includes date fixed effects and drops the day of the week fixed effects. The last column adds a dummy for the manager being a minority and the interaction of the worker's and the manager's minority status. Observations are worker-days and standard errors are clustered at the store level. * denotes significance at the 10% level.

Appendix Table A.7. Effect of Manager Bias on Time at Work and Work Performance
Including Controls for Other Manager Characteristics

	Absence Indicator	Minutes Worked in Excess of Schedule	Articles Scanned per Minute	Inter-Customer Time (Seconds)
Minority Worker × Manager Bias	0.0139** (0.0058)	-2.101 (1.639)	-0.277** (0.123)	1.523** (0.743)
Manager Bias	-0.0064 (0.0043)	1.007 (1.082)	0.1011 (0.078)	-0.269 (0.428)
Observations	4,371	4,163	3,601	3,287
Dependent Variable Mean	0.0162	-0.068	18.53	28.70
R-squared	0.088	0.131	0.196	0.241

Note: The regressions in this table replicate the regressions in the final columns of Table IIIA, IIIB, IVA, and IVB, respectively, adding additional control variables. The regressions include individual and date fixed effects, dummies for early morning and late evening shifts, an indicator for manager minority status, and the interaction of the worker's and manager's minority status. The regressions also include a dummy for the manager having a Level 4 position, a dummy for the manager being male, manager age as of January 1, 2012, and the interaction of each of these variables with the worker's minority status. ** denotes significance at the 5% level.

Appendix Table A.8. Effect of Manager Bias on Time at Work and Work Performance
Different Definitions of Minority Status

	Absence Indicator	Minutes Worked in Excess of Schedule	Articles Scanned per Minute	Inter-Customer Time (Seconds)
<u>A. Minorities as Workers with North African Names</u>				
Minority Worker × Manager Bias	0.0150** (0.0060)	-2.437 (1.791)	-0.228 (0.153)	1.305 (0.890)
Manager Bias	-0.0053 (0.0039)	1.238 (0.934)	0.075 (0.065)	-0.516 (0.503)
Observations	3,994	3,795	3,277	3,012
Dependent Variable Mean	0.0168	-0.725	18.51	28.46
R-squared	0.091	0.141	0.206	0.237
<u>B. Minorities as Workers with Sub-Saharan African Names</u>				
Minority Worker × Manager Bias	0.0082* (0.0044)	-6.010 (3.786)	-0.225 (0.138)	1.506* (0.854)
Manager Bias	-0.0043 (0.0038)	0.982 (1.001)	0.067 (0.062)	-0.701 (0.545)
Observations	3,561	3,380	2,906	2,670
Dependent Variable Mean	0.0177	-0.346	18.59	28.47
R-squared	0.100	0.131	0.217	0.249
<u>C. Workers of Indeterminate, Mixed, or Other Origin Excluded</u>				
Minority Worker × Manager Bias	0.0116** (0.0053)	-3.166* (1.755)	-0.181** (0.076)	1.282* (0.627)
Manager Bias	-0.0052 (0.0045)	0.812 (0.984)	0.030 (0.064)	-0.565 (0.465)
Observations	4,000	3,824	3,301	3,033
Dependent Variable Mean	0.0173	-0.020	18.61	28.69
R-squared	0.093	0.139	0.231	0.235

Note: The regressions in this table replicate the regressions in the final columns of Table IIIA, IIIB, IVA, and IVB, respectively. Regressions in Panel A define as minorities workers with a first or last name of North African origin. Remaining workers with a Sub-Saharan African first or last name are eliminated. Regressions in Panel B define as minorities workers with a first or last name of Sub-Saharan African origin, eliminating other workers with North African names. Regressions in Panel C use the original definition of minority (a first or last name of North or Sub-Saharan African origin), but eliminate workers with names of indeterminate, mixed, or other origin. *, ** denote significance at the 10% and 5% levels, respectively.

Appendix Table 9. Effect of Manager Bias on Time at Work and Work Performance
Controlling for Within-Store Shift Effects

	Absence Indicator	Minutes Worked in Excess of Schedule	Articles Scanned per Minute	Inter-Customer Time (Seconds)
Minority Worker × Manager Bias	0.0139*** (0.0046)	-3.366* (1.871)	-0.276** (0.128)	1.326* (0.755)
Manager Bias	-0.0086 (0.0057)	1.062 (1.568)	0.022 (0.098)	-0.699 (0.547)
Observations	4,371	3,221	3,601	3,287
Dependent Variable Mean	0.016	0.509	18.530	28.700
R-squared	0.227	0.423	0.382	0.393

Note: The regressions in this table replicate the regressions in the second-to-last columns of Table IIIA, IIIB, IVA, and IVB, respectively. These regressions add controls for the shift (day of the week × morning or evening) separately within each store. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Appendix Table A.10. Heterogeneity in the Effect of Manager Bias on Work Performance
 Dependent Variable: Articles Scanned per Minute

	A. Time During Contract		B. Fraction ZUS in the Store	
	Early Weeks	Late Weeks	Below Median	Above Median
Minority Worker × Manager Bias	-0.200 (0.141)	-0.422** (0.160)	-0.477** (0.203)	-0.002 (0.112)
Manager Bias	0.051 (0.123)	0.305** (0.115)	0.267** (0.099)	-0.102 (0.168)
Observations	2,404	1,197	1,864	1,340
Dependent Variable Mean	17.88	19.82	18.64	18.39
p-value: Coefficients are equal	0.249	0.249	0.037	0.037
R-squared	0.001	0.004	0.003	0.001

Note: Each regression shows the result of regressing the dependent variable on the interaction of the worker's minority status and the manager's IAT score (in standard deviation terms), controlling for the manager's IAT score and worker fixed effects. The dependent variable is the number of articles per minute scanned. Observations in the first column are limited to the early weeks (weeks 3 to 8) of the contract, while observations in the second column are limited the late weeks (weeks 18 to 23) of the contract. Observations in the third column are limited to stores in which managers reported relatively few workers from ZUS (below the median in our sample), while the final column includes the remaining stores. Observations are worker-days and standard errors are clustered at the store level. ** denotes significance at the 5% level.

Appendix Table A.11. Comparing Observations by Time During the Contract and Store Diversity

	A. Time During Contract			B. Fraction ZUS in the Store		
	Early Weeks (1)	Late Weeks (2)	p-value of Difference (3)	Below Median (4)	Above Median (5)	p-value of Difference (6)
Minority Worker	0.273	0.292	0.764	0.179	0.410	0.000
Male Worker	0.078	0.056	0.533	0.047	0.115	0.085
Minority Manager	0.055	0.059	0.906	0.040	0.089	0.242
Male Manager	0.110	0.118	0.864	0.071	0.178	0.052
Level 4 Manager	0.174	0.167	0.883	0.192	0.133	0.393
Manager Age	40.9	41.4	0.742	42.5	37.2	0.001
Manager IAT Score	1.28	1.44	0.315	1.44	1.35	0.640
Fraction ZUS in Store	0.166	0.176	0.458	0.108	0.255	0.000

Note: Cells in Column 1 report the mean of the indicated characteristic for the promotions for which we have data on the early weeks (weeks 3 to 8) of the contract. Cells in Column 2 report the mean of the indicated characteristic for the promotions for which we have data on the late weeks (weeks 18 to 23) of the contract. Column 3 presents the p-values from a test of the hypothesis that the means of both samples are the same. Cells in Column 4 report the mean of the indicated characteristic for stores in which managers report they managed relatively few workers from ZUS (below the median in our sample), while cells in Column 5 report means for the remaining stores. Column 6 reports p-values from a test of the hypothesis that the means of both samples are the same.

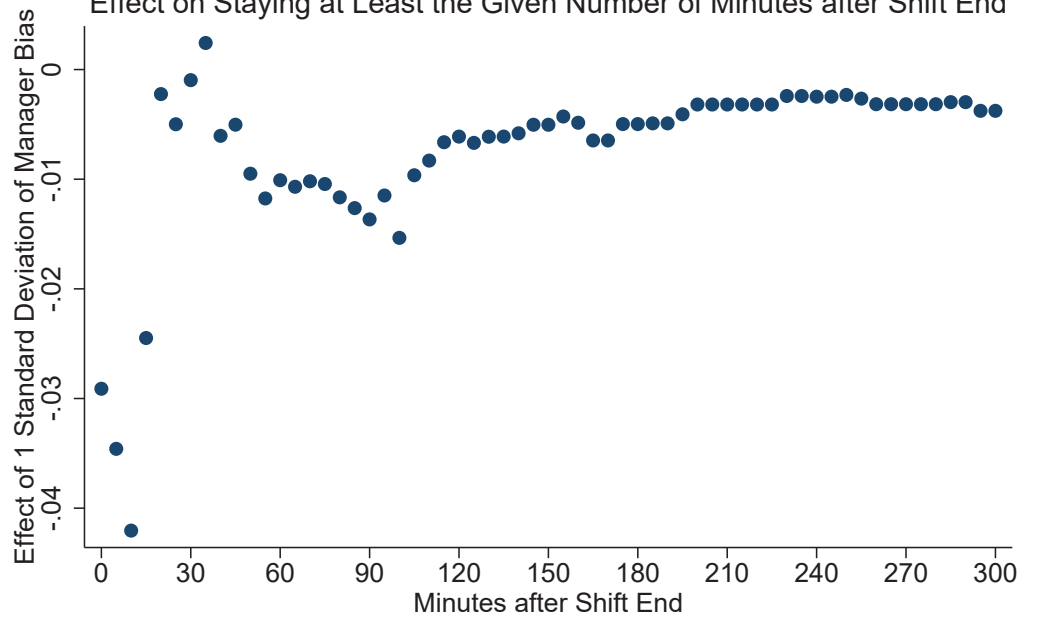
Appendix Table A.12. Effect of Manager Bias Over Different Levels of Time Aggregation

	<u>Level of Time Aggregation</u>			
	One Day	Two Working Days	One Calendar Week	Two Calendar Weeks
	<u>A. Absences</u>			
Minority Worker × Manager Bias	0.0098** (0.0039)	0.0021 (0.0043)	0.0091 (0.0106)	0.0115 (0.0245)
Manager Bias	-0.0021 (0.0031)	0.0008 (0.0037)	-0.0051 (0.0098)	-0.0047 (0.0223)
Observations	4,371	2,386	1,209	651
Dependent Variable Mean	0.0162	0.0176	0.0226	0.0252
R-squared	0.0005	0.0001	0.0004	0.0004
	<u>B. Articles Scanned per Minute</u>			
Minority Worker × Manager Bias	-0.276** (0.109)	-0.305*** (0.092)	-0.508*** (0.180)	-0.604 (0.361)
Manager Bias	0.140* (0.083)	0.221* (0.124)	0.400** (0.159)	0.716** (0.341)
Observations	3,601	2,149	1,111	605
Dependent Variable Mean	18.53	18.46	18.59	18.60
R-squared	0.0015	0.0034	0.0130	0.0316
Individual F.E.'s	Yes	Yes	Yes	Yes

Note: Each column in each panel shows the results of regressing the dependent variable on the interaction of the worker's minority status and the manager's IAT score (in standard deviation terms), controlling for the manager's IAT score and worker fixed effects. No other controls are included. The first column reproduces results from Tables IIIA and IVA, respectively. The dependent variable is an indicator for the worker being absent (Panel A) and the number of articles per minute scanned (Panel B). The remaining columns aggregate observations over longer time periods. In these regressions, both manager bias and the dependent variables are averaged (by worker) over the relevant time frame, so that the absence indicator is no longer an indicator, but a rate between 0 and 1. In the second column, the time span is two consecutive working days, so that observations are worker-two day periods. (If the data include an odd number of days for a given worker, one observation for the worker is a worker-one day period.) In the third column, the time span is a calendar week (typically four working days), and in the last column the time span is two calendar weeks. Standard errors are clustered at the store level. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Figure A.1. Effect of Manager Bias on Overtime

Effect on Staying at Least the Given Number of Minutes after Shift End



Note: Regressions control for date and time of day fixed effects.

Part II

The Value of a Vacancy: Evidence from a Randomized Evaluation with Local Employment Agencies in France

joint with Yann Algan and Bruno Crépon¹

Abstract

We study the effect of a Public Employment Service's (PES) intensive firm prospection campaign in which free recruitment services were proposed to a large number of small and medium sized firms. We find large impacts on vacancy postings with the PES and overall employment creation by firms. Confronting a simple model of firm search for candidates against data on vacancy characteristics and services delivered, we find strong evidence that candidate prescreening by the PES may be a key component of the intervention. These results suggest that active labor market policies that focus on firm labor demand may be a valuable addition to the labor policy toolkit.

¹We would like to thank Pierre Cahuc for his very helpful input as well as Barbara Petrongolo. We are grateful for comments and suggestions during seminars and conferences at Barcelona GSE, CREST, CEPR/IZA in London, Pôle Emploi, Sciences Po, Stockholm School of Economics and the University of Gothenburg. A big thank you to Agathe Pernoud for her superb research assistance. Also, special thanks to Pôle Emploi, notably François Aventur, Yannick Galliot, Marie-Jose Rabner and Stéphane Ducatez at the DSEE for their input and guidance during the project and to Aude Busson and Catherine Touati at the DSE for their excellent collaboration in the implementation phase of the project. We gratefully acknowledge funding from the Chaire de Sécurisation des Parcours Professionnels.

1 Introduction

The Mortensen and Pissarides (1994) equilibrium job search and matching framework explicitly models recruitment costs as a key parameter in determining labor demand. Yet active labor market policies, and the studies that try to measure their effects, have focused almost exclusively on the relationship between a firm’s labor demand and the expected productivity and costs that are incurred after hiring. Indeed, many studies have explored the impacts of assisting jobseekers through training programs or in their job search and there is strong evidence that these programs can be effective (Card et al., 2015). In this paper we explore whether a symmetric intervention that assists firms in their recruitment might also provide added value in the labor market.

The broad themes of matching and costly search in the face of asymmetric information for both the employer and employee form the basis of the economic literature, yet far less attention has been paid to how employers approach the hiring process (Oyer et al., 2011). For instance, Manning (2011) surveys the literature and shows that estimates of recruitment costs, as a part of the wage bill, vary from 1.5% to 11% depending on context and job-type.² Within the recruitment process, monetary job advertising costs may be quite low or even zero when posted through the public employment service, as Dickens et al. (2001) highlight, but Manning (2011) stresses that the bulk of the cost comes not from generating applicants, but in screening and then training them.³

Having devoted resources almost exclusively on assisting jobseekers since 2008, the public employment service (PES) in France, known as Pôle Emploi, revamped their firms services program for 2015.⁴ The new firm services were based upon a more intensive and dynamic treatment of vacancies, most notably the introduction of preselection, or prescreening, services to help firms more efficiently find the right candidate for the job. Delivery of these new services was to be complemented by prospection: agency counselors were to actively study firm needs and proactively contact and offer the new recruiting services in order to improve

²Table 2 in Manning (2011) summarizes percentages taken from eight different studies across several different countries and points out that we cannot fully determine if these are marginal or average recruiting costs. He suggests that the bulk of the cost is associated with training new employees, but that the cost of the hiring process itself remains an important question.

³As of mid 2017, a standard, 30 day vacancy posting on Monster.fr costs 855 euros. This posting cost increases to 1,390 euros for a 60 day, full service posting. Thus it is not readily apparent that vacancy flow costs are negligible in this context.

⁴Pôle Emploi has over a thousand local agencies throughout mainland France and its overseas territories. In 2008, it was created as the result of a merger between the ANPE (Agence nationale pour l’emploi), the government agency concerned with job counseling and recruitment services, and l’Assedic (Association pour l’emploi dans l’industrie et le commerce) the agency that dealt with the distribution and oversight of unemployment insurance benefits.

jobseeker placement.⁵

We study the effects of these new recruitment services by randomizing which firms were prospected. 8,232 firms participated in the study with half of them prospected intensively during a 10 week window during late 2014. Caseworkers were then instructed to incubate and keep relations with the prospected firms until April 2015, while no proactive action was taken towards the control group.⁶ By imposing random variation in firms expected recruiting costs we thus attempt to “shock” the value of a vacancy to the firm and examine how variation in this value affects firm demand for labor.

The PES’ partial shift away from an exclusive focus on jobseekers parallels recent work on the potential for displacement effects that may negate the effectiveness of active labor market policies that focus on improving the quality of the labor supply (Crépon et al., 2013). Thus a motivation for this study is not only to examine firm labor demand when we shock the underlying parameters of vacancy valuation, but also to test the effectiveness of an innovative policy that may avoid the mitigating effects of queue jumping between jobseekers vying for the same post.

Using our most conservative estimates, we find that this shock led to a 30% increase in vacancy postings with the PES, and more importantly, it also led to a 9% increase in permanent contract hires, translating into 48 more workdays created by firms, on average over the sixth month sanctuary period. We find impacts on the intensive and extensive margins, with prospected firms significantly more likely to hire a registered jobseeker. We are able to rule out that the strong positive effects on vacancy creation and hires are being driven by intertemporal substitution, i.e. an acceleration of the recruitment process. We also argue that these impacts are unlikely to be mollified by potential ‘firm displacement effects’ whereby the intervention could have led employees in control group firms to move disproportionately to jobs created in the treatment group.⁷ Nor do we believe that treatment firms likely displaced the recruitment processes either of control firms, or of firms outside our sample. We argue that because the intervention took place in low-tightness labor markets, the pool of available jobseekers is large and thus competition over candidates is low. Additionally, sample firms make up only a very small proportion of local firms, thus general equilibrium effects linked to program implementation are likely to be small. Finally, while we find strong impacts on vacancy creation with the PES on all firms within the sample, we find that impacts on hires are exclusively centered on firms that had been in contact with

⁵In qualitative interviews conducted in the feasibility phase of the study in early 2014, local job counselors who had previously worked for the ANPE highlighted anecdotally that the loss of relations with firms since the merger was a loss to job placement opportunities.

⁶Control firms were not denied service if they requested it during this “sanctuary” period.

⁷This would entail a potential for zero net creation in real employment.

the PES in the nine months preceding the intervention.

Finally, to structure our understanding of the results, we present a multi-channel recruitment model in which a firm’s valuation of a vacancy is based not only on the profitability margins through wage and vacancy flow costs, as is standard in the literature, but also on the costly selection process available in each channel. Though only suggestive, we believe it provides insight into how a firm’s demand for labor can shift given a shock to non-profit margins. Though non-experimental, we find empirical evidence to support the predictions of the model.⁸ We find no difference in the wage profitability margin between treatment and control vacancies, but find that free cost-reducing services were delivered to treatment vacancies. These services were centered on the implementation of prescreening candidates by caseworkers, which led to a reduction in firm search effort and also a drop in the average number of candidates sent to the firm for examination. These results are robust to controlling for selection effects correlated with the distribution of vacancy characteristics. This suggests that the boost in vacancy creation and hires may result from a reduction in matching frictions in the recruitment process itself and to our knowledge, this is the first study to demonstrate this strong link between recruitment costs and labor demand.

2 Description of intervention and heterogeneity dimension

2.1 Description of intervention

The public employment service’s new firm services or “nouvelle offre de services aux entreprises” is based on providing more comprehensive support to firms for their recruitment needs. The overall objective of the PES is to place jobseekers but in 2015, it moved towards a more balanced approach between aiding both jobseekers and firms. To accompany this renewed focus on firm relations, the PES elaborated two new services: Reinforced vacancy monitoring and follow-up (‘accompagnement à l’offre’) and prospection. This reinforced vacancy monitoring included specific services:

- Support with vacancy creation, drafting and posting
- Candidate prescreening and selection
- Interview support

⁸We use the data set of vacancies posted by control and treatment firms to explore impact mechanisms. Vacancies are not randomly assigned to treatment, only firms. Thus selection bias must be taken into account.

- Profile promotion (spontaneous candidatures sent by counselors to firms)
- Creation of a personal online recruitment web site and access to the PES résumé bank
- Adaptation and testing periods subsidized through continued UI benefits

After having been given the list of treatment and control firms (we explain in detail the sampling and randomization procedures below), the 129 participating local employment agencies were instructed to prospect the treatment firms intensively for 2-3 months starting on September 15th, 2014 while the control group was to be ‘sanctuarized’ for 6 months: no proactive action was to be taken towards these firms, but they were not refused service if they requested it. Agencies were required to have an in-depth interview with treatment firms during the intense period either through a face-to-face visit by a counselor or over the phone. During the interview, counselors were required to market the new and existing services to firms and understand the firm’s recruitment needs. Over the following 3-4 months, agencies were instructed to continue to nurture relations with the treatment firms. After the sanctuary period, agencies were free to contact and propose services to the control group.⁹

2.2 Heterogeneity

Firm prospection was put into place as a vehicle to promote and present the new PES recruitment services to firms that may, or may not, already have existing relations with the PES. Thus the intervention entailed actively engaging with two types of firms to learn about their recruiting needs and discover placement opportunities for jobseekers. We define this dimension of heterogeneity as a PES caseworker having made at least one successful phone call to the firm between January and August 2014. It is primordial to understanding the extent of the impact on our sample firms because the marginal impact of any introduction of new services provided to firms may be very different across “in-contact firms” and “no-contact firms.” Our sampling strategy ensured that we had a substantial proportion of both types of firms (36% were in contact with the PES during 8 months preceding the intervention) Hence, we will systematically display aggregate impacts along with heterogeneous effects across this dimension throughout the analysis.

⁹Agencies did not, in fact, massively switch to prospecting control firms after the sanctuary period. In informal exchanges with agencies, it appeared the time allotted to firm relations and prospection during the evaluation were not immediately preserved in the day-to-day workings of our sample agencies. But in 2016, the PES implemented a nation wide program in which certain agencies maintained caseworkers dedicated to helping firms with their recruitment needs.

3 Data, sampling and randomization

3.1 Data

We have access to rich historical administrative data from the public employment service at the firm and jobseeker levels. This includes vacancies posted with the PES as well as vacancy and jobseeker characteristics and the recruitment services provided to them. It is important to note a limitation of the vacancy data. We do not have an exhaustive measure of vacancy creation because the PES is only one outlet for vacancies among many other actors present on the market. This is counterbalanced by our data on hires. We have an exhaustive measure of hires through legally required hiring declarations called DPAAE, “Déclaration préalable à l’embauche.” All firms are required to submit a hiring declaration before, or shortly after the contract start date.¹⁰ Interim, or, temp-work contracts also require a declaration, but this is done by the temp agency. Thus, we exploit a separate data set created by the PES that documents the final employer (“using employer”) of the temp contract and append this to our data set of permanent and fixed-term contracts.

Vacancy data includes the posting date and type of contract: permanent (open-ended), fixed-term or interim (temp). We also have the posted characteristics such as the minimum annual salary, the profession, the required qualification, the minimum required experience and duration (for fixed-term contracts) and the weekly working hours.¹¹ Importantly for our mechanism analysis, we also have the applications, or, potential matches made through the PES to these vacancies through three different channels: applications initiated by the jobseeker, the firm and by the PES caseworker. We know the date of these applications, the origin and the result of the potential match.¹²

Data from the hiring declarations are not as rich as the vacancy data. We only have the contract type and its start and end dates (for fixed-term and temp contracts) and whether the person hired was a registered jobseeker with the PES in the 30 days preceding the hire date. Using the start and end dates for fixed-term and temp work contracts that ended during the observation period we calculate the number of workdays created for each contract. For

¹⁰Exceptions to the requirement for this hiring declaration concern internships and volunteer contracts and for the recruitment of private child care professionals and some public sector jobs. In the sampling phase, described in section 3.2, we target firms that were unlikely to make hires that do not require a declaration.

¹¹Unfortunately, vacancy closure dates are not reliable as the PES does not systematically follow-up on whether a vacancy has been satisfied. Hence we cannot reliably measure vacancy filling rates. This is even more of a problem given that the intervention led to a more intensive monitoring of the vacancy, thus exploiting vacancy closing dates could be subject to differential measurement error between groups.

¹²The result of the potential match is either a hire or a rejection by the jobseeker, the caseworker or the firm depending on the initiator of the application. The reliability of this match outcome variable may also suffer from differential measurement error between groups for the same reason as the vacancy closing date variable, thus we do not exploit these data for this paper.

permanent contracts, and fixed-term/temp-work contracts that ended after the observation period, we censor the end-date at January 31st, 2016, the end of the observed hiring period. We do this because these declarations are contract flows and for a large proportion of them, we have no personal identifiers due to the individual privacy constraints faced by the PES: personal identifiers are only available for individuals who were registered with the PES in the 3 years preceding the date of hire.¹³ This allows us to have a standardized measure of employment creation. For example, a week of one-day (Monday to Saturday) hires for the same individual would be counted as 6 fixed-term contracts flows, but as only one contract if it were a fixed-term contract that ran for the week. Thus, calculating workdays allows us to compare overall employment creation across contract types.¹⁴

Finally, we have rich data on the specific services that were provided to vacancies posted with the PES as listed in subsection 2.1.

3.2 Sampling and Randomization

It was important for the study’s external validity that the intervention targeted firms that were representative of local agencies’ firm portfolios and, at the same time, targeted firms that might recruit in low-tightness and low-job finding rate professions. As highlighted above, the public employment service’s main goal is to get jobseekers off their rosters, thus making sure that the intervention attacked low-tightness and job-finding rate professions was primordial.¹⁵ Thus the research team collaborated with the Firm Services Department at the PES to develop a sampling algorithm to target pertinent firms attached to the 129 local employment agencies participating in the study. We started by calculating labor market tightness over the 12 months preceding the randomization using jobseeker rosters and vacancy postings for each profession within each local agency. We also calculated the job-finding rate within these “micro-markets” for the same period. We then created a priority ranking of professions per agency using these two parameters as well as the stock of jobseekers registered in the agency.¹⁶ Using a profession-sector correspondence table, we then aggregated the weights per sector and ranked them. This gave us a ranked list of sectors in which firms were most likely to recruit within the prioritized professions.

¹³We count these anonymous hires as non-registered job seekers in the analysis.

¹⁴If we had individual identifiers for all the contract flows we could also measure the proportion of churn i.e. new hires versus renewed contracts for the same person. Unfortunately, we only have this for registered job seekers and as we’ll see below, our treatment affects the hires of both registered job seekers and non-registered job seekers.

¹⁵For example, it was important for the PES that any publicity of services made the distinction that they were provided to help jobseekers get back to work and not simply help firms recruit.

¹⁶The function used to assign the weights to the professions was convex in the stock of the job seekers and concave in tightness and the job finding rate.

Finally, these sector identifiers were linked to local firms that had responded to the PES’ annual survey *Besoin en Main d’Oeuvre* (BMO) or “Labor Needs” survey. Roughly 400,000 firms are surveyed in France each year to gauge their recruitment needs for the following year. The results are entered into an online platform used by the agencies to follow-up on potential hirings declared in the survey. We sampled in BMO 2014, a survey conducted in autumn 2013 on recruitment needs for 2014. This ensured that a significant portion of sample firms would have had contact with the PES in 2014, preceding the intervention. Each agency was then given a list of “priority firms” to potentially prospect drawn out of the BMO survey (those that were at the top of the sector rankings). They were then instructed to select roughly half of the list using their own local expertise. The final agency-trimmed list was sent back to the research team for randomization.

We stratified the final sample by indicators for the agency, if the firm intended to recruit in 2014 and by the number of employees on the firm’s payroll (in four categories). We were unable to stratify by the in-contact heterogeneity dimension because we did not have access to the administrative data for contacts at the time of the sampling and randomization. Within each stratum we randomly assigned treatment with probability one-half. For strata with odd numbers of firms we re-randomized the last firm within the stratum with probability 0.5 and did the same for singleton strata.¹⁷

3.3 Empirical Specification

We follow Imbens and Rubin (2015) and measure average treatment effects as the sum of the weighted difference in means within strata,

$$\widehat{ATE} = \sum_{s=1}^S q_s (\hat{\mu}_{1,s} - \hat{\mu}_{0,s}) \quad (1)$$

where $\hat{\mu}_{1,s} = \bar{y}^{s,T=1}$ and $\hat{\mu}_{0,s} = \bar{y}^{s,T=0}$ and q_s is equal to the sample share of observations in stratum s . The benefit of equation (1) is that it exploits the stratified sampling design of our study to the fullest. It is numerically analogous to estimating a fully interacted model that is centered at the mean rate of treatment assignment within strata,

$$y = a + bT + \sum_{s>1} \alpha_s 1_s + \sum_{s>1} \beta_s T(1_s - q_s) + u \quad (2)$$

Given the large number of strata in our study, equation 2 is computationally heavy,

¹⁷For the analysis, these singleton strata are reabsorbed into the closest stratum based on size, local agency and 2014 recruitment in that order of priority to have a minimum of 4 firms per stratum.

thus we directly calculate equation 1.¹⁸ We estimate the variance of our estimate following the influence function methodology developed in Hirano et al. (2003). This also allows us to cluster the standard errors at the local employment agency level (see the appendix for details) to account for correlation in outcomes among firm attached to the same agency.

4 Balance, sample description and compliance

Table 1 shows distribution statistics and balance checks for the final 8,232 firms retained in our sample. We also show these statistics and balance check estimates across our dimension of heterogeneity for the 7,859 firms for which we have within-strat variation in baseline contact with the PES. Each row presents the weighted control group mean and the treatment group difference as defined in equation (1). All dependent variables are indicators. Firm characteristics are collected from the BMO survey. For hires, vacancy postings, contacts and use of PES services, we sum the variables from January 2014 to August 2014 and create an indicator for the sum being larger than zero.

Across the board we see treatment coefficients close to zero and insignificant at the 10% level for all but four specifications out of a total of 69 regressions, showing that the stratified randomization was successful.

Examining the baseline characteristics of firms, we see that 72% of firms have less than 26 full time employees and that they are predominantly in the service (42%) and business sectors (27%) while manufacturing and construction make up 27% of the sample. 50% of firms hired a short-term contract (1 day to 6 months in duration) and 43% hired at least one employee in a permanent contract during this time period. Yet, relatively few firms post vacancies with the PES compared to the proportion that hire. For example, only 9% of firms posted a permanent contract vacancy over the eight month pre-intervention period.

In-contact firms are a bit larger in size, but are similar to no-contact firms on the other dimensions of firm characteristics that we measure. In contrast, in-contact firms display larger vacancy posting and hiring rates and, unsurprisingly, they receive visits and emails, and benefit from PES services at a significantly higher rate.

¹⁸Other estimates could be (a) the simple difference in treatment group means across the whole sample or (b) a regression equation simply including dummy variables for the strata. Because the assignment rate is not exactly 0.5 due to uneven and singleton strata, the first estimate (a) cannot be exactly written as a weighted average of estimated impacts within strata. Estimate (b) can be written in such a way, but the weights also involve the assignment rate in each strata (see Imbens and Rubin (2015)). We prefer to consider the *ATE* in (1) because it accurately reflects the experimental design. Estimates obtained using either (a) or (b) give very similar results.

5 Compliance and treatment intensity

Figure 1 plots the monthly cumulative evolution for visits, phone calls, candidate promotion and emails by firm from January 2014 through January 2016 using unconditional binned averages. The shaded region denotes the intense treatment period in which all treatment firms were expected to undergo an in-depth interview with a PES counselor. We see an upward linear evolution in all forms of contact and a sharp discontinuity for the treatment group at the beginning of the treatment phase. The figures show a jump of about a half a visit per firm on average and an increase of about one and a half more telephone calls made to the treatment group, representing 488% and 152% increases off of the baseline level.

A key strategy of the PES was the promotion, by counselors, of spontaneous candidatures adapted to firm needs. We consider this a form of compliance that demonstrates the implication of the counselors in the intervention.¹⁹ Again we see that treatment firms received close to one additional spontaneous candidature, on average, emanating from caseworkers compared to the control group which received almost none during the initial months of the treatment.

It is important to note the uninterrupted linear trajectory of the control group. As highlighted above, firms were free to contact the PES and request recruitment services. We do not observe a sudden change in the evolution of the control group curves. They do not suddenly go flat starting in September 2014. Thus our counterfactual represents simply what would have happened in absence of the prospection campaign, not what happens when firms are severed from PES services. We also note that contacts do not substantially change on average after the sanctuary period end date, March 31, 2015. One could imagine that when agencies were permitted to proactively encourage the control group firms to take advantage of PES services, we might see a jump in the contact and service levels given to control group firms after this date. This is not the case and thus permits us to explore whether effects persist over time.²⁰

¹⁹In the feasibility stages of the project, we attempted to convince the PES to allow for two distinct treatment groups. One that contained just the marketing of services and in-depth interviews and another which would additionally use the spontaneous candidature promotion service. This experimental design was deemed infeasible for the PES.

²⁰For example, we will show that the impacts are not a result of short-term intertemporal substitution i.e. a “speeding-up” of the recruitment process during the sanctuary period.

6 Impacts

6.1 Vacancy and hiring flows

Table 2 displays results for the average treatment effect on flows for each type of contract during the sanctuary period. We top-code all vacancy and hiring variables at the 99th percentile of their distribution.²¹ Panel A displays impacts on aggregate, over the whole sample, while Panel B exhibits heterogeneous impacts using our in-contact indicator. All estimations come from equation 1 and standard errors are clustered at the local employment agency level. For all regressions we show the strata weighted control mean of the dependent variable, in order to gauge effect magnitudes, and the p-value for a test of equality of the average treatment effect between in-contact and no-contact firms. In examining columns 1-3 of Panel A, we see that prospection leads to large increases in vacancy flows posted at the PES. On average, treatment firms posted 0.064 and 0.048 more job offers for open-ended and fixed-term contracts, respectively. This is a bit larger than a 30% increase off the baseline mean for both types of contracts. During this period, we see no significant impact on the posting of temp work vacancies. Across all contract types (column 4) the intervention led to an increase of 0.11 vacancies posted with the PES, an increase of around 19% when compared to the control mean.

Column 5 in panel A of Table 2 shows that this increase in vacancy posting is accompanied by an impact on hires in permanent contracts. Treatment firms create 0.12 additional open-ended contracts, on average, equivalent to an increase of 8.8%. In contrast, we see non-significant point estimates for fixed-term and temp hires.²² When looking over hires in all contract types the point estimate is positive, but insignificant.

Panel B of Table 2 displays results for the heterogeneity analysis. Firms that had previous contact with employment agencies in the months leading up to the intervention drive a significant portion of the effect. Though both types of firms significantly increase their vacancy postings we see larger point estimates for in-contact firms in columns 1-3 that are about twice the size. The relative percentage change is also larger for in-contact firms (25.2%) compared to no-contact firms (17.6%) when looking over all vacancies, though we cannot

²¹Table A.1 in the appendix presents results using non top-coded data. Effect sizes are larger with the non-transformed dependent variables. We believe it is important to show these effects because this is administrative data so the data points in the upper distribution are unlikely to be errors. In addition, the PES naturally tries to create and improve relations with “large-recruiters”, what they call *les grands comptes*. Thus, prospection of these firms could have also led to increased vacancy and hiring flows in these firms. Nevertheless, we prefer to display our most conservative estimates in the main text.

²²The substantially higher level of flows for fixed-term and temp contracts per firm as seen in the control means are due primarily to very short and mostly one-day contracts in which the same person may be hired multiple times by a firm in a short period of time. We address this below in our analysis of workday creation.

formally reject the hypothesis that the ATE is equal between the two types of firms at a reasonable significance threshold (p-value of 0.195). But these heterogeneous effects are most striking when we look at hiring flows. We see in column 5 of panel B that the effect seen on the whole sample is centered entirely on in-contact firms. Firms who had previous relations with the PES see a 24% increase in permanent contract hires. Even with the loss in statistical precision by splitting the sample along this dimension, this estimate is still significant at the 1% level. When summing overall all contract types we find a positive effect (column 8): in-contact firms create almost 3 more contracts than their control counterparts. For firms with no previous contact, the estimated impact on permanent contract flows is much smaller and insignificant. Considering all contracts, we estimate a small and negative coefficient for no-contact firms with the hypothesis of the equality in the ATE between the two types of firms being rejected at the 5% level.

The evidence in Table 2 strongly suggests that the program generated substantial vacancy postings with the PES and that, at least, some of these postings were vacancies that would not have been created in absence of the program. We can make this inference because we also see large impacts on real hires. Thus, the impact of the program cannot be reduced to a simple vacancy substitution effect whereby firms could have either duplicated their postings that would have been posted elsewhere anyway, or simply substituted away from their existing posting medium to the PES.

6.2 Employment creation

Even though we see strong positive impacts on permanent hire contracts, it is insufficient to only examine contract flows to determine if the intervention led to more job creation. It could be that the treatment leads to a substitution between contract types. For example, if the treatment leads to more permanent contracts and less fixed-term or temp work contracts being emitted by the firm, it could be that actual job creation is null. Hence, in order to estimate the impact on real employment creation we sum the workdays within each contract. This allows us to compare employment creation over the different contract types. We use the start and end dates of the contract available in the hiring declaration and impute the end-date of contracts that terminate after the study period or that are open-ended (i.e. permanent and some temp contracts) as January 31st, 2016. This allows us to measure impacts on total workdays created by sample firms.²³

²³As noted above, we only have personal identifiers in the hiring data for individuals who were registered at least once at the PES in the three years preceding the start date of the contract. Thus, we are unable to determine if contracts emitted by the firm concern the same individual (i.e. renewal) or different individuals for a significant portion of the hiring flows. Focusing on workday creation circumvents this issue of rotation on the same post.

Table 3 provides evidence on employment creation. The dependent variable is the sum of workdays created in new posts for each firm over the 6 month sanctuary period. We group the estimates from the model by type of jobseeker. Registered jobseekers are defined as individuals who were registered with the PES within the last 30 days of the contract start date. We consider hires who were not registered with the PES within the last 30 days or for whom we have no personal identifiers in the hiring declaration data as non-registered jobseekers.²⁴ For these jobseeker types, we display results for employment creation within permanent contracts and aggregated across all contracts (omitting the specific results for fixed-term and temp contracts). Descriptively, we see that control group firms created, on average, 838 workdays during the sanctuary period of which 526 days were in permanent contracts. Interestingly, the majority of this employment was created for non-registered individuals (523 days versus 315 days for registered jobseekers). This illustrates the fact that much of the employment creation, destruction and turnover in the job market happens outside of the PES' purview.

In turning to program impacts, columns 1 and 2 of Panel A of Table 3 provide supportive evidence that the treatment increased net job creation. Treated firms created 33 more days of employment on average and this overall positive effect on employment creation is driven by the creation of 48 more workdays in permanent contracts, on average per firm. The heterogeneity analysis in Panel B again shows the striking differential effect between in-contact and non-contact firms: all of the strong positive impact on employment creation is centered on firms who were already in contact with the PES prior to the beginning of the intervention. On average, treated in-contact firms created 155 more work days than their control equivalents while we see no effect on no-contact firms. We even measure a negative average effect for the hires of registered jobseekers as can be seen in columns 3 and 4 of panel B, though the effect on workdays in all contract types is only marginally significant at the 10% level.²⁵ All-in-all, across panel B we strongly reject the null hypothesis that impacts on employment creation are equal between in-contact and no-contact firms for all models

²⁴See section 2 for more details on the hiring data.

²⁵There are several interpretations of this possible negative effect. One is that, given the estimate's marginal significance, this is just statistical noise and the true population parameter for the no-contact firms is 0. Another could be that this is possible evidence of displacement among firms in which in-contact treatment firms partially exhaust the stock of pertinent registered jobseekers (we will discuss possible displacement effects in a dedicated section below). Finally, it could be that caseworker efforts to incite the firm to use PES services could actually have a negative effect on no-contact firms' willingness to hire registered jobseekers. In qualitative interviews during the treatment period, caseworkers noted that, for certain firms, the PES has an image problem. These caseworkers evoked the idea that some firms believe that the PES is only interested in placing registered jobseekers regardless of firm satisfaction with the hire. This idea connects more broadly to the constraints the PES faces in their mission to both place jobseekers that have difficulty getting back to work, and to satisfy firm needs. We will return to these themes in the mechanisms section.

except for employment creation for non-registered jobseekers.

In comparing columns 3 and 6 of panel B in Table 3, we also see that the strong, positive impact on workday creation for in-contact firms is roughly evenly split between registered and non-registered jobseekers. The ratio of these point estimates is $64.7/76.9 \approx .84$ and when we compare this to the ratio of control means for these firms $220.6/383.4 \approx .58$ it provides strong evidence that the intensive prospection campaign led not only to more hires, but a shift to hiring more registered jobseekers. Examining the distribution of the effect on hiring flows provides further evidence of this move towards hiring more registered jobseekers.

6.3 Distribution of impacts

Overall, treatment firms are 1.7 percentage points more likely to make at least one hire in a permanent contract (significant at the 10% level) compared to the control group. And as discussed above, this impact is driven by the in-contact firms in our sample. Looking at the left-hand side graph in the second row of column 1 of Figure 2, we see that this increase on the extensive margin is driven by a massive increase in the probability to recruit registered jobseekers. Each bar represents the treatment impact on the probability to make at least the given number of hires as denoted by the horizontal axis. We overlay the bars with 95% confidence intervals and also highlight the quantiles of the underlying distribution of permanent contract hires with vertical red lines. We also systematically report the p-value for tests of the equality of distributions of hires between groups (Mann-Whitney test). We see large treatment impacts on the probability to hire at least one, two, three, four or five registered jobseekers. Specifically, the treatment is associated with a 5.8 percentage point increase to hire at least one registered jobseeker for in-contact firms. In addition to showing the stark differences in impact across our dimension of heterogeneity, these figures provide a key robustness check: They show that the impact on mean hires is not being driven exclusively by firms on the far right-hand side of the hiring distribution. Though only 25% of firms make at least one permanent contract hire, a large part of the impact is centered on an extensive margin of registered jobseekers, thus permitting the unemployed to get back to work in a stable contract.

Finally, Figure 2 also provides evidence that the intervention could have also incentivized an increases in permanent contract hires for large recruiting firms. Interestingly, this effect on the intensive margin is driven by the hiring of non-registered job seekers among in-contact firms as can be seen in the graph in the third row of column 1. This suggests in addition to triggering effects on the extensive margin, that prospection may also have incentivized

relatively large scale recruiters to shift their recruitment a bit to the right.²⁶

6.4 Simply intertemporal substitution?

In order to rule out the possibility that impacts are being driven by an intertemporal substitution effect in which the treatment causes a simple acceleration of an existing or future recruitment process, we also explore cumulative vacancy and hiring flows for the post sanctuary period and over the entire 17 month observation period. Table A.2 compares the impacts on cumulative flows during three periods: the sanctuary period, the post sanctuary period and the entire study period. If we look at a possible post-sanctuary impact of the treatment on vacancy posting, we see that the coefficients are close to zero and insignificant for all contract types except for temp vacancies: we estimate a negative effect on temp contract vacancy flows shown in panel A, column 3, which is significant at the 10% level.²⁷ Looking over the entire study period we see overall similar effect sizes that are no longer statistically significant due to the additional statistical noise introduced by the wider observation period. This suggests that all of the impact on vacancies occurred during the sanctuary period. This conclusion is supported by examining the vacancy flows per month over the entire study period on the left column of graphs in Figure A.1 in the appendix. These figures show that all of the positive effect on vacancies occurred during the first six months of the intervention.

Impacts on hiring flows in Panel B of Table A.2 tell a similar story to vacancies except that the impact on permanent contract flows is larger and still statistically significant at the 10% level when looking at cumulative flows over the entire 17 months. This suggests that over the nine months following the sanctuary period, the control group did not “catch-up” to the treatment group, supporting the idea that the effect is not due to inter-temporal substitution, but real job creation.

6.5 Displacement?

We make the distinction between two types of displacement that could affect the interpretation of our results. The first is displacement among jobseekers. Crépon et al. (2013) show that active labor market policies aimed directly at jobseekers that do not take into account general equilibrium effects may over estimate impacts due to the effect of “queue jumping:”

²⁶One possible reason for this could be the increased vacancy exposure received when a vacancy is posted with the PES. Treatment firms are thus more exposed to applicants engaging in on-the-job search. Unfortunately, we cannot test this hypothesis directly.

²⁷Rather than intertemporal substitution, this could be evidence of a substitution between contract types. It is plausible that the large increase in permanent and fixed-term vacancy postings is a partial result of firms moving away from temp work contracts. Nevertheless, the evidence is weak

policies only change the order in which vacant posts are filled, not actual employment levels. Indeed, one of the motivations of this study was to focus specifically on job creation because of these potential externalities related supply-side interventions. The simple theoretical framework derived in their paper clearly shows that displacement effects associated with a shift in the labor supply are quite large in weak labor markets, i.e. markets for which the equilibrium tightness is low. The intervention we consider in this paper instead consists in shifting the labor *demand* curve to the right.

Clearly a shift in a firm's demand curve is also associated with a second form of displacement effects at the firm-level, whereby firms compete to hire the best candidate and the intervention changes not only labor demand, but also the ordering of access to candidates. It is beyond the scope of this paper to investigate such equilibrium effects, but we believe that in our case they are of limited importance. This claim stems not from the intervention itself, but from the labor market context: French labor market tightness at the end of 2014 was at its lowest level since 2000 with a value of 0.4 to compare to the 0.6 long run average and the 0.8 high reached in 2000.²⁸ And contrary to a shift in the labor supply, displacement effects associated with a shift in labor demand are small in weak labor markets, a point made by Michaillat (2012) who argues that if local tightness is low, firms have a plethora of choices among candidates to fill their post and thus suffer less from recruitment competition by other firms.

Another reason why displacement effects due to the intervention may be very small is that the share of firms involved in our experiment at the micro market level is quite small. Sample firms make up only 1% of agencies' portfolio of local firms, on average. It follows that the treatment firms would have very little influence on equilibrium outcomes and as Figure A.3 shows, the labor market tightness that these sample firms experience is very similar to the average low-tightness that non-sample firms face. Finally, returning to panel B of Table A.2, even when aggregating permanent contract flows over the entire 17 month observation period (Sept. 2014 - Jan. 2016), a positive impact persists. If treatment firms were displacing control firms during the sanctuary period, we might expect this difference to return to zero (or at least drop) when aggregating over the longer period.

Finally, another type of displacement linked to job destruction might occur. It could be the case that the impacts that we see are, in part, due to the movement of personnel from control firms to treatment firms (or vice versa). If this were the case we would over estimate the benefits of the intervention in equilibrium because we do not observe employment destruction. We provide evidence on the flows of hires between treatment groups and the starting situation of the recruited person in Table A.3 in the appendix. Row titles correspond

²⁸See <http://dares.travail-emploi.gouv.fr/IMG/pdf/2016-012.pdf> for more details.

to the origin of the hired individual, column titles to where the hired person was placed and in which type of contract. We thus categorize all flows within, coming in- and going out of our sample firms during the sanctuary period. The proportions are displayed above the number of total flows for each type of jobseeker.

We find that there are relatively inconsequential flows between our sample firms. They represent roughly 3% of all flows for permanent contracts, and are almost all rehires or change of contracts within the same firm. Reassuringly, we measure almost no flows to and from other firms, and this regardless of the treatment status of the firm. In terms of flows coming into the sample, unsurprisingly a large proportion come from unemployment i.e. registered jobseekers, thus there is no issue of employment destruction elsewhere. We also note a significant proportion of contract flows for jobseekers whom we do not know the origin or who were employed elsewhere. Unfortunately we are unable to obtain a clear picture of what these really are. Thus it is possible that some of these individuals are leaving vacant posts behind them as they move into employment in our sample.²⁹

In sum, we are confident that our estimated parameters truly measure the impact on treated firms. This doesn't imply, however, that potential displacement effects would not become an issue in the case of a large scale-up of the policy.

7 Potential Mechanisms

We have shown that the intervention led to unambiguously large impacts on firms vacancy posting with the PES. This was accompanied by a substantial increase in permanent contract hires leading to an increase in the number of workdays created by treatment firms as compared to control firms. This increase was, however, not observed for all firms. We do not detect any impact on either contract flows or workday creation for firms which were not previously in contact with the PES. In this section we present a simple model that we believe elucidates some potential mechanisms driving these results. It predicts that providing search and screening services to firms leads them to reduce their own search effort and to be more picky in their hiring decisions. It also predicts that providing this assistance will have a positive impact on vacancy postings, but the impact on the number of hires will be ambiguous.

We then use the rich administrative data set of the 2,052 permanent contract vacancies posted with the PES during the 6 month sanctuary period to test the model's predictions and complement our experimental evidence. We underscore that this part of the empirical

²⁹As noted above, the hiring declaration data only contains personal identifiers for the individual if they were registered at least once in the preceding three years before the hiring date.

analysis is non experimental and therefore only suggestive. Indeed, the treatment increased the number of vacancies posted with the PES and it would therefore be imprudent to consider the two sets of vacancies as identical: comparisons between the two sets confound a selection effect linked to new types of vacancies posted and a treatment effect on the efficiency of vacancy filling.

We begin with a simple description of the way vacancies are treated at the PES. This involves two important aspects: (1) the different search channels available and (2) the range of services provided to the vacancy. There are three main channels through which vacancies can be filled and our model will account for all of them. The majority of vacancies are posted online (see Table A.7) and, depending on the firm's preference, jobseekers may apply directly to the firm or only through the counselor. The caseworker can then filter candidates and send them on to the firm. If the vacancy is not publicly posted, then only the caseworker is responsible for both generating and prescreening applicants. Lastly the firm itself can search for candidates in the PES resumé bank and contact them directly, or through the caseworker.

The PES vacancy services target different parts of the recruitment process. We start by presenting evidence on the delivery of these recruitment services. As our analysis will make clear, the treatment almost exclusively involves search and screening assistance, and, as a result, our model will focus on these two aspects. Table 4 shows results from simple OLS regressions of indicators for different types of service provision, as noted in the column titles, on a treatment indicator. These results totally ignore the selection issue linked to the fact that the set of vacancies posted by firms in the treatment and control groups are not the same. Later we produce additional results showing that the results are stable when controlling for vacancy characteristics, as a way to deal with this selection issue.

Looking at column 1 of Table 4, we find that treatment status is highly correlated with vacancies being tagged for intensive follow-up support. On average, treatment vacancies are 41% more likely to receive intensive follow-up support (known as "offre en accompagnement") within the agencies, an increase of around 12 percentage points off the baseline mean. The counselor initiated act of categorizing the vacancy for follow-up support effectively opens the door to apply the whole gamut of new services. Looking at columns 2 we see that the tagging for follow-up support entailed an almost systematic implementation of prescreening services and that this prescreening involved two additional key services: special preselection and verification. Special preselection involves working with the employer to establish specific criteria, or, a maximum of 5 prerequisites, on which to prescreen candidates that are sent to the employer for an interview. Verification entails that a maximum of 5 to 10 candidates per post are sent to the employer and that the way in which the candidates apply is appropriate.

For example, the PES might recommend that the firms choose to have applicants apply only through the counselor. Finally, verification requires counselors to negotiate a time frame with the firm for the delivery of the applicants and ways in which to adapt the vacancy if there is an insufficient number of applicants. Interestingly we find that the service of valorization, in which counselors put special effort into highlighting specific jobseekers assets and abilities is not widely used. We see a small control group mean and no difference between the groups of vacancies.

In examining the remaining columns in Table 4 we find very low levels of other service provision and no difference between treatment and control vacancies in the application of services that might influence the costs associated with vacancy posting or interviewing applicants. For example, we see no difference in the provision of services that might reduce costs associated with creating and drafting a vacancy appropriately (columns 7 and 8). Nor do we find that counselors applied their services after the preselection phase: we see nothing in terms of interview support (column 9) nor the implementation of helping jobseekers adapt to the job (column 10).

7.1 Outline

We model the firm recruitment process through its valuation of a vacancy. To summarize, we consider each vacancy requires a specific skill set and workers have heterogeneous skill sets so that they are an imperfect quality match to the job. We decompose the hiring process into several steps including the search for candidates, their screening, interviews and hiring. Following our previous description, we introduce several channels through which applications are made and these channels may or may not involve prescreening. The firm maximizes its valuation of a vacancy over these channels with respect to its recruitment effort and hiring threshold. The model predicts unambiguously that, under preselection treatment, more vacancies will be posted and that these vacancies will receive less applicants due to reduced firm effort and jobseekers that are more stringently filtered by the PES. It also predicts an ambiguous effect on the number of applicants coming from PES counselors. Finally, the effect on hires is also ambiguous because firms reduce their own search effort and become more picky even as they post more vacancies.

7.2 Setup

The firm has an opportunity to produce output y during a period of time that ends at an instantaneous rate s . For this production it offers a reference wage w . The value to the firm of this activity is $v = (y - w)/(r + s)$ where r is the discount rate. The firm must recruit

somebody to realize the production and there are three channels through which candidates arrive: (1) jobseekers apply on their own directly to the firm or through the PES platform, (2) the firm expends effort to search on its own. Finally, candidates come through (3) a PES caseworker channel. The arrival rate is δ for jobseekers, μ for the caseworker channel and e for the firm which incurs a cost for its search effort $c(e)$. The candidate and caseworker channels are free. The firm may decide however to only use its own channel ($d = 0$) or to also consider applicants arriving through the caseworker and jobseeker channels ($d = 1$).

The firm looks for different skills and has imperfect knowledge of the labor market, meaning candidates in all channels are more or less suitable to the needs of the firm. Hired jobseekers can provide the firm an instantaneous profit $t \times v$, with t being an applicant specific random draw from a uniform distribution over $[1 - 1/\gamma, 1]$.³⁰ This captures the idea that the job requires certain skills and that jobseeker skills are an imperfect match. This also implies that it might be difficult to find employees willing to do the job for the targeted wage w .

The free services offered by the PES as part of the intervention include preselection or, prescreening, of candidates. Under preselection, jobseeker applications are no longer direct and are filtered by the caseworker. With the implementation of preselection, candidates arriving through the PES (either through δ or μ) are drawn from the same distribution but only the θ top quantile $t > 1 - 1/\gamma + \theta/\gamma$ get through. We assume that the firm always prescreens candidates. Once received in an interview the characteristic t is revealed to the firm. This is a final screening phase with cost κ in all channels. After interviewing a candidate, the firm decides whether to hire. This decision is based on the characteristic t being above a threshold \bar{t} .

- Absent the program, we consider jobseekers arrive at a rate μ_0 and δ under the caseworker and jobseeker channels, and are only prescreened under the firm channel. The firm decides the optimal search effort e_0^* and selects applicants with skills above t_0^*
- With the intervention, we consider jobseekers arrive at a rate μ_1 and δ under the caseworker and jobseeker channels and are prescreened under *all three channels*. The firm makes optimal search effort e_1^* and selects applicants with skills above t_1^*

The value of a vacancy thus depends on the four dimension parameter $\nu = (v, \delta, \kappa, \gamma)$. This is an important aspect of the model because the decision to post a vacancy will depend on this parameter. It involves the profitability of the vacancy, but also three other parameters. These parameters are more specifically linked to labor market characteristics: the

³⁰ $\gamma > 1$, the larger γ is the larger the range of skills of potential candidates.

parameter δ measures the quality and size of the market - some job offers might be atypical and attract only few workers; κ measures screening or interview costs - the firm might have needs outside the scope of its recruiting expertise hence the possibility of large screening costs; and lastly γ measures the quality of a randomly selected applicant.

We derive the value of vacancies characterized by ν under the two regimes, with the program $\Pi(\nu, \mu_1, 1)$ and without the program $\Pi(\nu, \mu_0, 0)$. Without the intervention the firm evaluates the value of a vacancy $\Pi(\nu, \mu_0, 0)$ by optimizing with respect to its effort and hiring threshold:

$$r\Pi(\nu, \mu_0, 0) = \max_{e, \bar{t}, d \in \{0,1\}} \left\{ -c(e) + e\theta \left(\gamma \int_{\bar{t}}^1 (tv - \Pi(\nu, \mu_0, 0)) \frac{dt}{\theta} - \kappa \right) + d(\delta + \mu_0) \left(\gamma \int_{\bar{t}}^1 (tv - \Pi(\nu, \mu_0, 0)) dt - \kappa \right) \right\}, \quad (3)$$

and with the program,

$$r\Pi(\nu, \mu_1, 1) = \max_{e, \bar{t}, d \in \{0,1\}} \left\{ -c(e) + e\theta \left(\gamma \int_{\bar{t}}^1 (tv - \Pi(\nu, \mu_1, 1)) \frac{dt}{\theta} - \kappa \right) + d(\delta + \mu_1) \theta \left(\gamma \int_{\bar{t}}^1 (tv - \Pi(\nu, \mu_1, 1)) \frac{dt}{\theta} - \kappa \right) \right\}. \quad (4)$$

The following propositions demonstrate key results:

1. The value of a vacancy always increases when the intervention is implemented as long as $\mu_1 \geq \mu_0$: $\Pi(\nu, \mu_1, 1) > \Pi(\nu, \mu_0, 0)$
2. Firms always use caseworker and jobseeker channels when the intervention is implemented: $d_1^* = 1$
3. Firms reduce their search effort and are more picky in their choice of candidates: $e_1^* \leq e_0^*$ and $t_1^* \geq t_0^*$
4. The value function under preselection is increasing in v , δ , γ and decreasing in κ

Proofs: See appendix

7.3 Predictions

We further assume that opportunities arrive at a rate λ per unit of time and that they are drawn from a distribution with density $\phi(\nu)$ and that there is a fixed cost F in posting a vacancy.³¹

³¹There is a condition for the value of the vacancy to be positive which writes as $v(1 - 1/(2\gamma)) > \kappa\theta$. Because we introduce the fixed cost for vacancies to be posted, we assume the condition is always satisfied.

This allows us to derive expressions for the number of vacancies posted, number of applicants to these vacancies and finally, the number of hires.

7.3.1 Number and types of vacancies posted

Given that opening a new vacancy has a fixed cost F , a new opportunity ν will lead to the opening of a vacancy absent the program when $S_0(\nu) = \mathbf{1}(\Pi(\nu, \mu_0, 0) > F) = 1$ and when $S_1(\nu) = \mathbf{1}(\Pi(\nu, \mu_1, 1) > F) = 1$ when the program is implemented. Because $\Pi(\nu, \mu_1, 1) > \Pi(\nu, \mu_0, 0)$, all opportunities such that $\Pi(\nu, \mu_1, 1) > F > \Pi(\nu, \mu_0, 0)$ will be opened as a result of program implementation. Therefore the flow of new vacancies opened by the firm, $N_v(\mu_1, F, 1)$ under the program and $N_v(\mu_0, F, 0)$ absent the program can be expressed as,

$$\begin{aligned} N_v(\mu_1, F, 1) &= \lambda P_\nu(S_1(\nu)) \\ N_v(\mu_0, F, 0) &= \lambda P_\nu(S_0(\nu)) \end{aligned} \tag{5}$$

and we have

$$N_v(\mu_1, F, 1) - N_v(\mu_0, F, 0) = \lambda P_\nu\left(\Pi(\nu, \mu_1, 1) > F > \Pi(\nu, \mu_0, 0)\right) > 0.$$

One important dimension is that there are several margins on which the number of vacancies is increasing and new vacancies are not necessarily vacancies of smaller profitability. Consider an opportunity $\nu_0 = v_0, \delta_0, \kappa_0, \gamma_0$ at the margin absent the program, i.e. $\Pi(\nu_0, \mu_0, 0) = F$ in conjunction with Proposition 4. We can describe how these margins are affected

- Profitability margin: less profitable vacancies are posted. v_1 such that $\Pi(v_1, \delta_0, \kappa_0, \gamma_0, \mu_1, 1) = F$, satisfies $v_1 < v_0$
- Jobseeker arrival rate margin: vacancies with lower application rates are posted. δ_1 such that $\Pi(v_0, \delta_1, \kappa_0, \gamma_0, \mu_1, 1) = F$, satisfies $\delta_1 < \delta_0$
- Skills assessment cost margin: vacancies with larger screening costs are posted. κ_1 such that $\Pi(v_0, \delta_0, \kappa_1, \gamma_0, \mu_1, 1) = F$, satisfies $\kappa_1 > \kappa_0$
- Skills signaling margin: vacancies that attract more heterogeneous applicants are posted. γ_1 such that $\Pi(v_0, \delta_0, \kappa_0, \gamma_1, \mu_1, 1) = F$, satisfies $\gamma_1 < \gamma_0$

We now explore the selection of vacancies posted by firms in the treatment and the control groups. It is impossible to get complete measures of all of the underlying parameters thus we are not able to fully document the selection issue. Rather, we rely on key vacancy

characteristics recorded in PES administrative data: the minimum wage offered, hours, the skill and experience requirements as well as the occupation. Our results show that the sets of vacancies posted by firms in the treatment and control groups are not the same, but we do not find any evidence that these differences are related to differences in profitability. This suggests, following the model, that differences in the sets of vacancies in the two groups are more related to the labor market parameters δ , κ and γ . To document this point, we simply run OLS regressions of vacancy characteristics on a treatment indicator.

Table 5 presents these results. In columns 1-3 we look at a key job search parameter, the posted wage.³² In the first column we see small and insignificant coefficients on the treatment indicator for the log of the annual posted wage. In the second column, we predict the wage on a sample of 1,921,148 permanent contract vacancies posted with PES by firms outside of our sample during the treatment period. We construct this prediction by regressing the log of the posted annual wage on indicator variables for the number of hours in 8 categories and indicators for required experience (in years) in 6 categories. Finally we include 95 indicators for the profession and the required qualification in 9 levels along with their interactions. The main motivation behind using this prediction is that it proxies for the output linked to the job and thus the underlying components of its profitability.³³ Again we see no reasonable selection effects on the predicted wage. Column 3 shows results of regressing the difference between the real posted wage and the predicted value on our treatment indicator. Again, we see small and insignificant coefficients indicating that wage determinants between treatment and control vacancies are similar. In addition, we see no significant differences when splitting the population by our in-contact indicator. In sum, we cannot provide credible evidence that vacancies in the treatment group were selected for posting (or not) based on a wage profitability margin.

In contrast we see in columns 4-9 of Table 5 that treatment vacancies posted by in-contact firms require less experience, less qualification and more hours worked per week. On aggregate, we only see significant correlation between treatment status and the type of required qualification posted. Treatment vacancies are 11 percentage points more likely to

³²We use the log of the minimum posted wage in the vacancy data. The max wage is missing for a large percentage of vacancies.

³³Formally, the predicted wage takes the form

$$\hat{w} = \hat{\beta}_0 + \sum_{h=2}^8 \hat{\gamma}_h + \sum_{e=2}^6 \hat{\alpha}_e + \sum_{p=2}^{140} \hat{\delta}_p + \sum_{q=2}^9 \hat{\lambda}_q + \sum_{p=2}^{95} \sum_{q=2}^9 \hat{D}_{pq} \quad (6)$$

Where w is the log of the wage posted and β_0 is a constant. $\hat{\gamma}_h$, $\hat{\alpha}_e$, $\hat{\delta}_p$ and $\hat{\lambda}_q$ are predicted coefficients on dummy variables for the weekly hours, experience, profession and qualification categories, respectively, and \hat{D}_{pq} are the coefficients on the interaction of the profession qualification indicators from an OLS regression. We then apply these estimated parameters to predict the wage within our sample of vacancies.

be jobs requiring lower skills.³⁴ It is difficult to link these differences to the three other parameters characterizing vacancies in our model. However, we could consider that low skill jobs without experience are jobs for which the relevant skills might be difficult to read and might correspond to a lower γ . Similarly we observe that, for in-contact firms, treatment vacancies more often require hours above the 35 hour legal limit. One interpretation in line with the model would be that this corresponds to jobs with atypical hours that are therefore less easy to fill (smaller δ).

In the rest of the paper, when comparing vacancies in treatment and control group, we will account for these differences in characteristics using inverse probability weighted regressions (IPW). We do not believe this fully solves the selection bias, but we believe this makes the comparisons more meaningful. We apply this method to first test the robustness of the service delivery results. Appendix Table A.4 replicates results from Table 4 using IPW regressions.³⁵ We find that the IPW results are very similar to our simple regression results.

7.3.2 Number of applications

The model also allows us to obtain the instantaneous probability that an application arrives through a channel on opportunity ν under the two regimes. The average number of applications for vacancies posted by firms in treatment and control groups directly follows. We summarize these average predicted flows in Table 6. Given these expressions we clearly observe that the model predicts a reduction in the average number of applications to treatment firm vacancies through two of the channels: a smaller application rate generated by the firm itself as $e_1^* < e_0^*$ while jobseeker applications are subject to caseworker screening ($\delta\theta$). The only potential positive impact passes through an increase in counselor effort μ_1 , but this is mitigated by the preselection of the best applicants. Moreover these predictions also hold in the presence of selection bias. For example, vacancies posted at the margin have lower jobseeker application rates, all else being equal.

Table 7 reports the difference in application rates between firms in treatment and control groups using IPW regressions.³⁶ The table shows the the number of applications per vacancy along the three channels (employer, caseworker and jobseeker) during the two weeks following the posting date of the vacancy.³⁷ We see that treatment vacancies receive significantly

³⁴We define low qualification as laborers, production workers and unqualified employees. High qualification jobs are defined as supervisors, technicians and management.

³⁵See table A.4 notes for details on the specification.

³⁶See Table 7 notes for details. Appendix Table A.5 shows that unweighted regression results are very similar to the IPW estimates.

³⁷Appendix Table A.6 shows results at 8 weeks. They are very similar suggesting that the majority of

fewer applicants through the firm and jobseeker channels and that these represent very large percent changes off the control mean. There are about 2.4 less jobseeker initiated matches in the treatment group of a baseline of 6.5 for the control group. Employers make roughly 0.78 less potential matches off an average of 1.14 in the control group. Interestingly, we see no effect of the treatment on mean applications coming from the counselors themselves, but it appears that, though insignificantly different (with a p-value of 0.21), treatment vacancies for in-contact firms had a reduced mean application rate through μ while the no-contact firm rate was higher.

This empirical evidence provides support for the model’s predictions. Using expressions in Table 6 for firm and caseworker effort ratios and prescreening intensity, we now apply the estimates from Table 7. Results are reported in Table 8. For the application rate ratio through the employer channel in treatment and control groups, we obtain $E(e_1^*)/E(e_0^*) \approx 0.31$.³⁸ Using the application rate ratio through the jobseeker channel in treatment and control groups, we obtain an estimate for $\theta E(\delta)/E(\delta) \approx 0.66$. Lastly, if we consider the ratio of caseworker effort we get a ratio of $\mu_1/\mu_0 \approx 1.53$. Thus we observe evidence that the intervention led to recruitment conditions in which $e_1^* \leq e_0^*$, $\mu_1 > \mu_0$ and $\theta < 1$. Put into words, this implies that firms reduced their own search effort, caseworkers increased their effort and that there was a significant increase in the level of prescreening.³⁹

Our ability to use the model to derive the ratios of effort and prescreening rates relies on strong modeling assumptions. Indeed, we would not be able to identify the caseworker channel ratio if we allowed the screening rate to be channel dependent. However, we believe these results are highly illustrative. They provide strong evidence that prescreening took place and (given that the prescreening rate does not differ substantially by channel) that there was a significant substitution between firm and caseworker search effort.

7.3.3 Number of hires

To close discussion of the model we now turn to its prediction on hiring rates. The instantaneous probability of a hire on a specific opportunity ν both with and without the program

‘action’ on vacancies happens within the first couple of weeks of creation.

³⁸We no longer use the terms that explicitly condition for selection due to the different vacancy sets, as we have already addressed this issue using IPW estimates. Of course, the caveat being that this is only an approximate solution for the selection bias issue.

³⁹Appendix Table A.7 presents additional results on selection showing that part of the prescreening toolkit appears to be that treatment vacancies are less likely to be posted publicly on the internet. Column 1 shows that treatment vacancies are about 8 percentage points less likely to be posted on the PES website. As noted above, when the vacancy is not posted publicly, the caseworker is responsible for generating and screening applicants. This interpretation is also supported by column 1 in Table A.8 which shows that counselors dramatically intensified their vacancy monitoring.

can be formalized as,

$$\begin{aligned} P(\text{Hire}|\nu, \mu_1, 1) &= (e_1^* + \mu_1 + \delta) \gamma(1 - t_1^*) S_1(\nu) \\ P(\text{Hire}|\nu, \mu_0, 0) &= (e_0^* + d_0^*(\mu_0 + \delta)) \gamma(1 - t_0^*) S_0(\nu) \end{aligned} \quad (7)$$

and we can also derive the corresponding number of hires at the firm level:

$$\begin{aligned} N_h(\theta, \mu_1, 1) &= \lambda E_\nu [(e_1^* + \mu_1 + \delta) \gamma(1 - t_1^*) S_1(\nu)] \\ N_h(\theta, \mu_0, 0) &= \lambda E_\nu [(e_0^* + d_0^*(\mu_0 + \delta)) \gamma(1 - t_0^*) S_0(\nu)]. \end{aligned} \quad (8)$$

Clearly the impact on hires is ambiguous. The initial positive effect of a broader set of vacancies posted ($S_1 \geq S_0$) and the increase in caseworkers' effort ($\mu_1 > \mu_0$) is counterbalanced by the reduction in firm search effort $e_1^* < e_0^*$ and also by its more selective behavior $t_1^* > t_0^*$.

We now discuss the divergent impacts on hires across our contact heterogeneity dimension. Recall that Table 3 shows a large and significant increase in workday creation for registered jobseekers in in-contact firms, but that the point estimate is small and negative for no-contact firms. One possibility could be that the program is simply implemented differently for the two sets of firms. Table 8 also presents the estimates of effort ratios and the screening parameter for the two different types of firms. As can be seen in the table, the underlying parameters we estimate are broadly identical between in- and no-contact firms. We see little difference overall albeit with a caseworker effort ratio slightly higher for no-contact firms. Hence, it does not appear that this is the driving factor for the difference in hiring rates between the two groups.

Another possible explanation could be that the distribution of vacancies in the two sets of firms is different. Looking again at Table 5 we see that indeed the distributions are not the same. Vacancies opened by in-contact firms require less experience, lower qualification and have atypical working hours. It is unclear how these specific characteristics relate to the hiring impact. To answer this question we would need to know for which vacancies the intervention leads firms to reduce their effort the most or for which they are more picky about applicants. We are unable to detail the actual candidate hiring threshold t^* . However, we see that, at the very least, the firm effort ratio $E(e_1^*)/E(e_0^*)$ is almost identical between the two types of firms, suggesting that the difference in vacancy characteristics may not be the main factor behind the differential impact.

A simpler and potentially more straight forward explanation for the differential hiring effect could simply be that the program was not implemented effectively for the no-contact firms, despite the PES' best efforts. There might be two reasons for this. The first is that

the distribution of skills the firm has access to through the PES might have little overlap with the the set of skills it needs. This might be the reason the firm was not in contact with the PES in the first place. The second reason may be linked to difficulty counselors face in effectively prescreening for firms they are just getting to know.

8 Conclusion

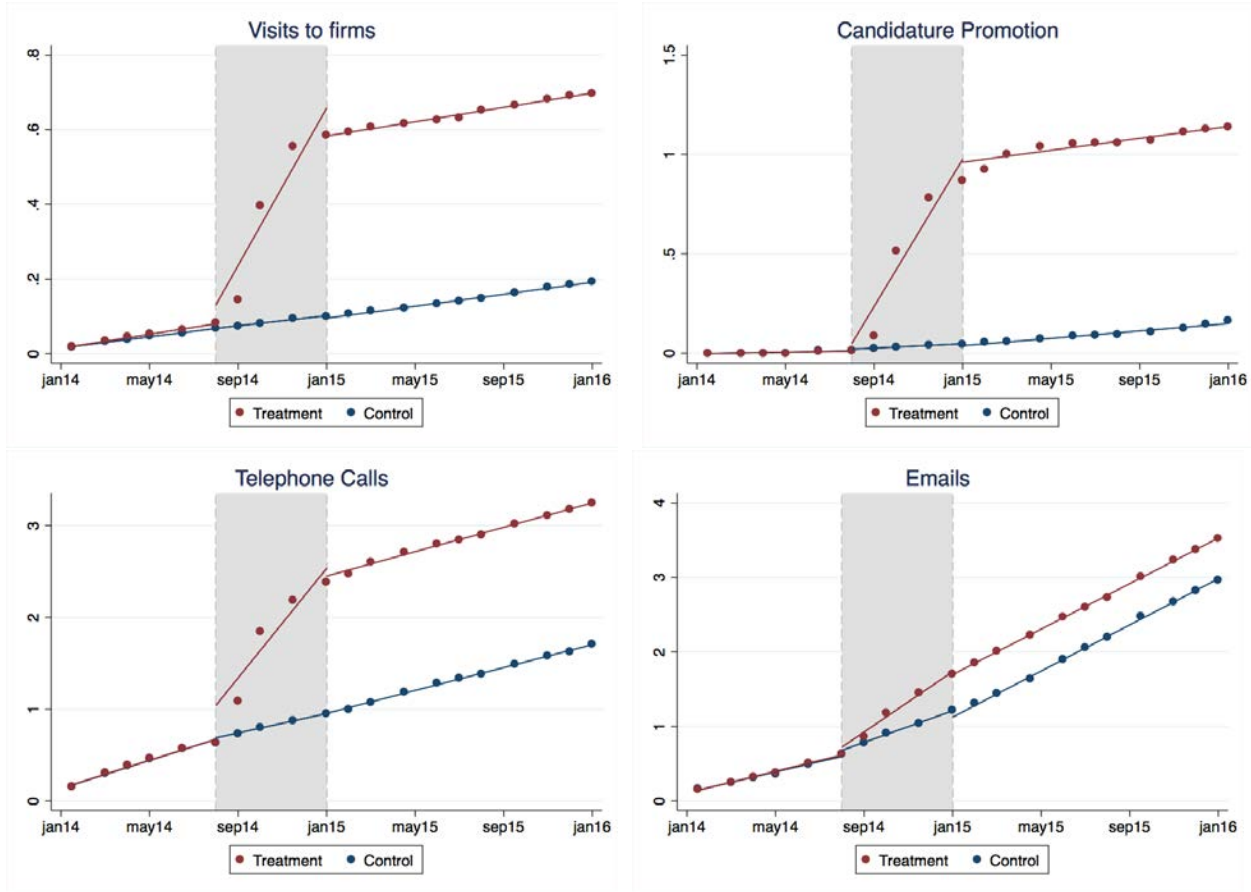
We present evidence on the impacts of an intervention that targets firm labor demand by supporting its recruitment practices. We study the effect of a Public Employment Service's (PES) intensive firm prospection campaign in which free recruitment services were proposed to thousands of firms. We find large impacts on vacancy postings, permanent contract hiring flows and the number of workdays created by firms, suggesting that active labor market policies that focus on firm labor demand can be very effective.

We examine the characteristics of job postings that were created by sample firms and find that the vacancies created by treatment firms were subject to much higher levels of candidate prescreening by the PES, but do not differ from vacancies created by control firms on typical wage profitability margins. We develop a multi-channel firm search model to better understand the potential mechanisms driving the effect and find suggestive evidence that the delivery of these additional prescreening services may be a key component of the intervention because it transfers search and screening costs away from the firm. In turn it predicts that firms put forth less effort and also become more picky, leading to potentially ambiguous effects on hires.

We find strong heterogeneous effects in the experimental results. Though we find impacts on vacancy creation across all types of firms, only firms that were previously in contact with the PES increase their hiring rates. This suggests that the public employment service's relationships with firms, in light of their primary responsibility to place marginalized jobseekers, is an important area of future research.

Figures and Tables

Figure 1: Compliance and treatment intensity



Note: Figures illustrate the average number of counselor initiated visits, phone calls, emails and jobseeker résumés spontaneously sent to firms. The numbers are averaged into bins corresponding to each month during the observation period for treatment and control firms along. The shaded region indicates the intensive treatment period (September - December 2014) in which caseworkers were supposed to engage in in-depth interviews with firms to promote services and learn about their recruitment needs.

Table 1: Balance check and descriptive statistics

	Total		Contact= 1		Contact= 0	
	(1) Control Mean	(2) Treatment	(3) Control Mean	(4) Treatment	(5) Control Mean	(6) Treatment
<u>Heterogeneity</u>						
Contact with PES	0.357	-0.003 (0.009)				
<u>Firm Characteristics</u>						
≤ 10 employees	0.402	-0.001 (0.001)	0.301	0.004 (0.003)	0.463	-0.003 (0.002)
> 10 & ≤ 25 employees	0.324	0.004 (0.005)	0.330	0.002 (0.010)	0.320	0.004 (0.006)
> 25 & ≤ 50 employees	0.169	-0.003 (0.005)	0.210	-0.003 (0.010)	0.145	-0.003 (0.006)
> 50 employees	0.106	0.000 (0.001)	0.158	-0.002 (0.004)	0.072	0.002 (0.002)
Manufacturing	0.113	-0.004 (0.006)	0.108	0.004 (0.011)	0.114	-0.007 (0.008)
Construction	0.157	0.002 (0.007)	0.180	-0.020 (0.015)	0.150	0.007 (0.008)
Commerce	0.269	-0.010 (0.008)	0.250	-0.004 (0.016)	0.278	-0.002 (0.011)
Service	0.418	0.008 (0.010)	0.420	0.015 (0.021)	0.416	0.001 (0.012)
Other sectors	0.043	0.004 (0.005)	0.043	0.005 (0.009)	0.042	0.001 (0.006)
<u>Hires by contract type</u>						
Fixed-term < 6 months	0.500	-0.005 (0.011)	0.593	-0.010 (0.021)	0.443	-0.001 (0.014)
Fixed-term ≥ 6 months	0.157	0.004 (0.007)	0.209	-0.002 (0.015)	0.128	0.006 (0.009)
Permanent	0.434	0.010 (0.011)	0.517	0.015 (0.019)	0.385	0.003 (0.014)
Temporary	0.220	-0.006 (0.008)	0.275	-0.007 (0.015)	0.194	-0.006 (0.010)
<u>Vacancies posted at PE</u>						
Fixed-term	0.073	-0.003 (0.006)	0.132	-0.021* (0.013)	0.039	0.002 (0.007)
Permanent	0.087	-0.008 (0.005)	0.137	-0.009 (0.012)	0.058	-0.008 (0.005)
Temporary	0.107	0.002 (0.006)	0.138	0.002 (0.011)	0.088	0.003 (0.007)
<u>Contact with PE</u>						
Emails	0.218	0.008 (0.008)	0.351	0.015 (0.016)	0.137	0.004 (0.009)
Visits	0.060	0.009* (0.005)	0.103	0.006 (0.013)	0.037	0.012** (0.005)
<u>PE services</u>						
Jobseeker initiated match	0.147	-0.008 (0.007)	0.200	-0.004 (0.014)	0.112	-0.007 (0.008)
Counselor initiated match	0.177	-0.009 (0.007)	0.266	-0.019 (0.015)	0.120	-0.005 (0.008)
Employer initiated match	0.022	0.002 (0.003)	0.031	-0.002 (0.006)	0.015	0.007* (0.003)
Successful match	0.057	-0.001 (0.005)	0.098	-0.002 (0.010)	0.031	-0.002 (0.005)
Spontaneous candidature	0.013	-0.001 (0.003)	0.022	0.003 (0.006)	0.008	-0.003 (0.003)
N		8232		2686		5173

Note: Rows display results from separate estimates of equation 1 for the given dependent variable. Columns 1 and 2 display results for all firms while columns 3-6 show results across our dimension of heterogeneity. All dependent variables are {0, 1} indicators for which we display the weighted control mean along with the difference in the treatment group. Standard errors for the treatment group difference are in parenthesis and are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 2: Impact on flows

	Vacancies				Hires			
	Permanent (1)	Fixed-term (2)	Temp (3)	All (4)	Permanent (5)	Fixed-term (6)	Temp (7)	All (8)
Panel A: Overall								
Treatment	0.064*** (0.017)	0.048*** (0.013)	-0.001 (0.019)	0.110*** (0.032)	0.118** (0.051)	0.296 (0.380)	-0.084 (0.584)	0.331 (0.643)
Control Mean	0.202	0.142	0.249	0.592	1.335	5.257	7.685	14.277
N	8232	8232	8232	8232	8232	8232	8232	8232
Panel B: Heterogeneity								
<u>Contact= 1</u>								
Treatment	0.106*** (0.036)	0.077*** (0.029)	0.002 (0.038)	0.185*** (0.068)	0.366*** (0.124)	1.047 (0.711)	1.368 (1.182)	2.782** (1.323)
<u>Contact= 0</u>								
Treatment	0.051** (0.021)	0.035** (0.015)	-0.001 (0.025)	0.085** (0.039)	0.028 (0.057)	-0.067 (0.330)	-0.542 (0.713)	-0.581 (0.742)
Control Mean(Contact= 1)	0.291	0.171	0.271	0.733	1.530	6.847	8.561	16.938
Control Mean(Contact= 0)	0.159	0.107	0.216	0.482	1.213	4.336	7.227	12.777
p-value Equality of Coefs.	0.184	0.204	0.937	0.195	0.020	0.130	0.160	0.025
N	7859	7859	7859	7859	7859	7859	7859	7859

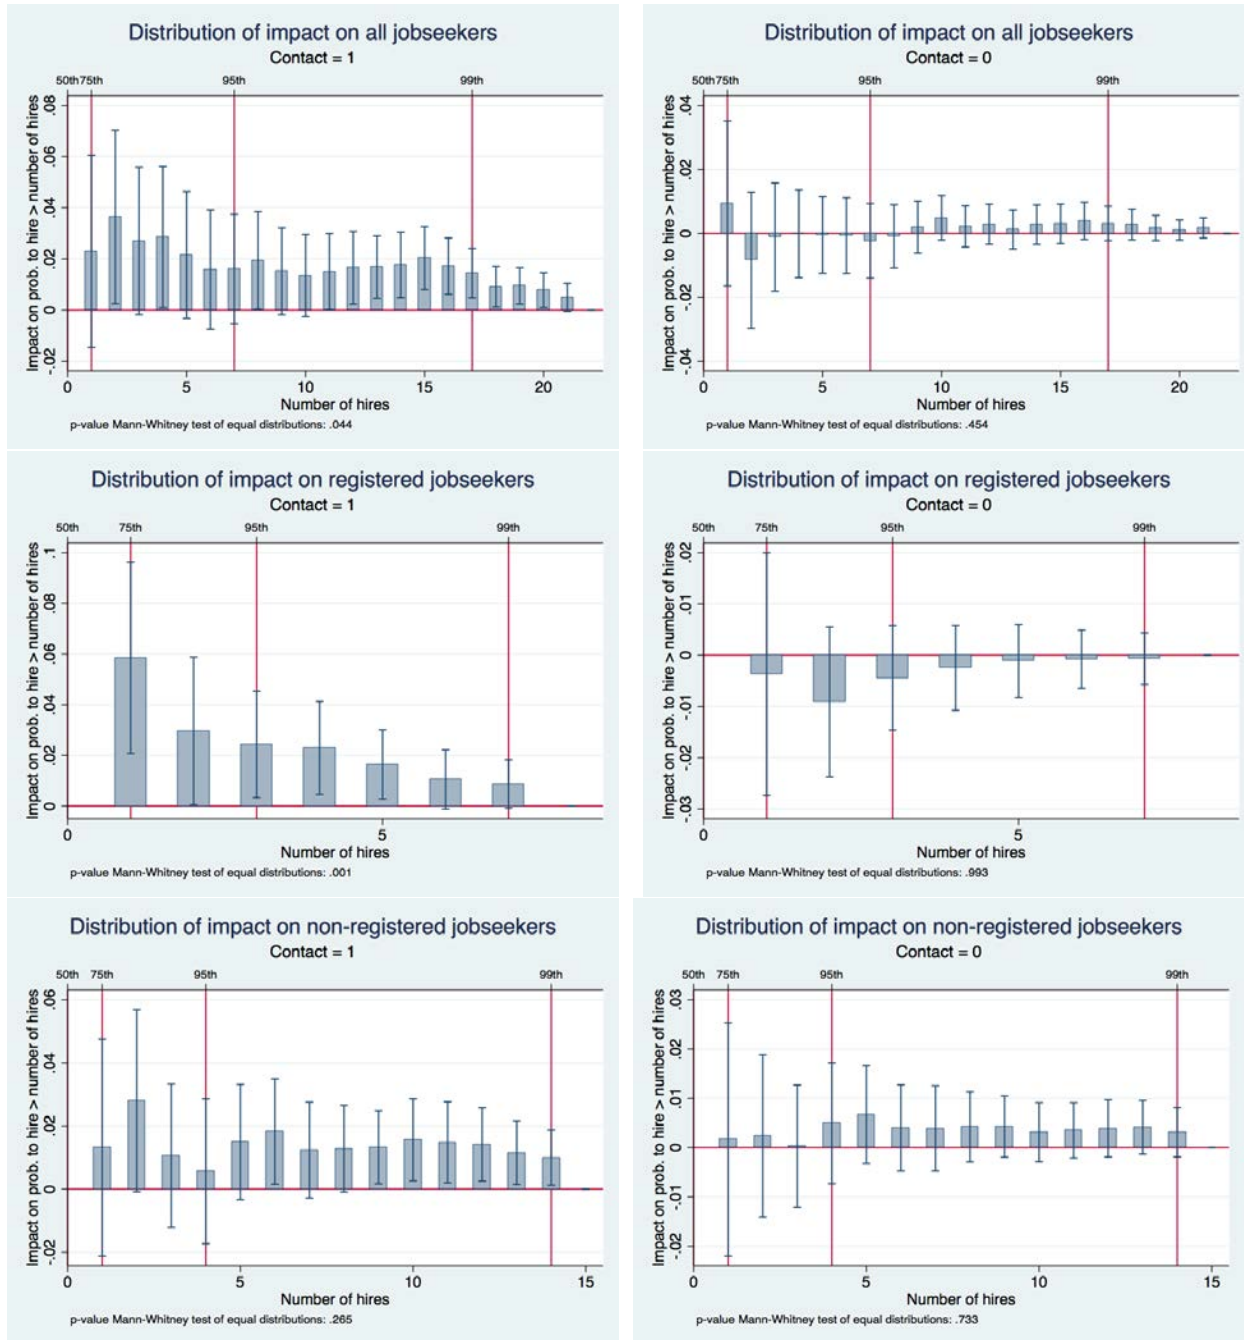
Note: This table presents impacts on vacancy postings with the PES and hiring flows for the three contract types (columns 1-3 and 5-7) as well as total flows across all contracts (columns 4 and 8). Panel A presents average treatment effects on the whole sample while Panel B displays impacts across our heterogeneity dimension (having previous contact with the PES) along with the p-value for a test of equality of treatment effects between in-contact and no-contact firms. Only firms that have within-stratum variation in contact status are used in the heterogeneity analysis. Average treatment effects are estimated per equation 1. Strata weighted control group means are also shown. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 3: Impact on workdays

	All Jobseekers		Registered Jobseekers		non-Registered Jobseekers	
	(1) Permanent	(2) All	(3) Permanent	(4) All	(5) Permanent	(6) All
Panel A: Overall						
Treatment	48.1** (20.4)	33.8 (23.4)	17.8* (9.4)	10.0 (10.9)	30.2** (14.6)	23.8 (17.2)
Control Mean	525.6	837.5	190.4	314.7	335.1	522.7
N	8232	8232	8232	8232	8232	8232
Panel B: Heterogeneity						
<u>Contact= 1</u>						
Treatment	141.6*** (49.5)	154.7*** (55.7)	64.7*** (21.6)	65.9** (26.5)	76.9** (33.9)	88.8** (37.8)
<u>Contact= 0</u>						
Treatment	14.7 (22.7)	-6.2 (25.1)	-8.2 (9.9)	-20.0* (10.6)	23.0 (17.4)	13.9 (20.4)
Control Mean(Contact= 1)	604.0	965.4	220.6	370.4	383.4	595.0
Control Mean(Contact= 0)	477.0	749.3	171.4	274.2	305.6	475.1
p-value Equality of Coefs.	0.028	0.009	0.002	0.002	0.185	0.091
N	7859	7859	7859	7859	7859	7859

Note: This table presents impacts on workday creation within permanent contracts and over all contract types. Columns 1 and 2 present results for all hires while columns 3-4 and columns 5-6 display results for registered and non-registered jobseeker hires, respectively. Workdays are calculated using the start and end dates of the contract, with end dates for permanent and fixed-term contracts censored at 31 January 2016 (this concerns fixed term contracts that end after this date). Panel A presents average treatment effects on the whole sample while Panel B displays impacts across our heterogeneity dimension (having previous contact with the PES) along with the p-value for a test of equality of treatment effects between in-contact and no-contact firms. Only firms that have within-stratum variation in contact status are used in the heterogeneity analysis. Average treatment effects are estimated per equation 1. Strata weighted control group means are also shown. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Figure 2: Distribution of impact on permanent hiring flows



Note: Figures illustrate the distribution of impacts on permanent contract hiring flows for different types of jobseekers for in-contact firms (figures in the left column) and no-contact firms (right column). Vertical bars show the percentage point impact on an indicator for making at least the number of hires as denoted by the horizontal axis. Bars are overlaid with 95% confidence intervals. Vertical lines mark the quantiles of the underlying distribution of hires. Impacts are estimated per equation 1 and standard errors are clustered at the agency level.

Table 4: Selection on provision of vacancy services

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Intensive Follow-up	Preselection	Special Preselection	Verification	Valorization	Evaluation	Analysis of post	Drafting support	Interview support	Adaptation support
Panel A: Overall										
Treatment	0.113*** (0.033)	0.109*** (0.033)	0.115*** (0.032)	0.106*** (0.032)	0.008 (0.008)	0.001 (0.004)	0.007 (0.011)	0.003 (0.020)	0.003 (0.004)	-0.003 (0.006)
Control Mean	0.283	0.270	0.235	0.261	0.025	0.007	0.022	0.113	0.005	0.014
N	2052	2052	2052	2052	2052	2052	2052	2052	2052	2052
Panel B: Heterogeneity										
<u>Contact= 1</u>										
Treatment	0.098** (0.046)	0.092** (0.046)	0.089** (0.043)	0.093** (0.045)	0.003 (0.010)	0.001 (0.007)	0.015 (0.015)	-0.005 (0.029)	-0.004 (.)	0.000 (0.008)
<u>Contact= 0</u>										
Treatment	0.133*** (0.039)	0.129*** (0.037)	0.148*** (0.037)	0.123*** (0.037)	0.014 (0.011)	0.002 (0.005)	-0.004 (0.012)	0.012 (0.028)	0.011 (.)	-0.006 (0.010)
Control Mean(Contact= 1)	0.285	0.273	0.244	0.263	0.027	0.008	0.015	0.123	0.008	0.015
Control Mean(Contact= 0)	0.279	0.265	0.224	0.260	0.022	0.005	0.033	0.101	0.000	0.014
p-value Equality of Coefs.	0.562	0.534	0.303	0.599	0.452	0.856	0.313	0.680	.	0.621
N	2052	2052	2052	2052	2052	2052	2052	2052	2052	2052

Note: This table presents selection effects of PES services applied to permanent contract vacancies during the sanctuary period using linear probability models. Indicator variables for service type (displayed in the column headers) are regressed on a treatment indicator. Panel A presents these effects on the whole vacancy sample while Panel B displays effects across our heterogeneity dimension (having previous contact with the PES) along with the p-value for a test of equality of selection effects between in-contact and no-contact firms. Control group means are also shown. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 5: Selection on vacancy characteristics

	Vacancy Characteristics								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	w	\hat{w}	$w - \hat{w}$	Experience	Low Qualif.	Hours	Hours < 35	Hours = 35	Hours > 35
Panel A: Overall									
Treatment	-0.032 (0.035)	-0.015 (0.032)	-0.008 (0.014)	-0.200 (0.127)	0.106*** (0.030)	0.291 (0.483)	-0.017 (0.024)	-0.018 (0.030)	0.034 (0.026)
Control Mean	9.916	9.892	0.022	2.179	0.632	33.874	0.136	0.664	0.200
p-value rank-sum (Mann-Whitney)	0.405	0.124	0.887	0.241	0.000	0.040	0.267	0.411	0.067
N	1921	2052	1921	2052	2052	2052	2052	2052	2052
Panel B: Heterogeneity									
<u>Contact= 1</u>									
Treatment	-0.050 (0.056)	-0.019 (0.049)	-0.021 (0.018)	-0.298** (0.147)	0.141*** (0.043)	0.315 (0.718)	-0.002 (0.031)	-0.087** (0.042)	0.089** (0.036)
<u>Contact= 0</u>									
Treatment	-0.007 (0.038)	-0.011 (0.037)	0.009 (0.018)	-0.073 (0.206)	0.059 (0.040)	0.258 (0.659)	-0.035 (0.039)	0.072* (0.043)	-0.037 (0.037)
Control Mean(Contact= 1)	9.923	9.889	0.031	2.244	0.615	33.810	0.125	0.702	0.173
Control Mean(Contact= 0)	9.907	9.896	0.010	2.094	0.656	33.958	0.150	0.615	0.235
p-value Equality of Coefs.	0.520	0.895	0.245	0.373	0.161	0.953	0.508	0.008	0.014
N	1921	2052	1921	2052	2052	2052	2052	2052	2052

Note: We display characteristics for permanent contract vacancies during the sanctuary period and their correlation with treatment status. w and \hat{w} are the log of the posted minimum yearly wage and the log of its outside sample prediction (see section 7.3.1 for details on the prediction). Only 1,921 permanent contract vacancies have usable wage data. Experience is defined as the minimum required experience for the post in years. Low qualification, Hours₃₅, Hours = 35, Hours_{>35} are indicator variables. Panel A presents these correlations on the whole vacancy sample while Panel B displays correlations across our heterogeneity dimension (having previous contact with the PES) along with the p-value for a test of equality of coefficients between in-contact and no-contact firms. Control group means are also shown. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 6: Application rates under the two programs

	Firm	Counselor	Jobseeker
Application rate for a given application			
With the program	$e_1^*\theta$	$\mu_1\theta$	$\delta\theta$
Absent the program	$e_0^*\theta$	μ_0	δ
Average application rate for vacancies posted under the two programs			
With the program on all opened vacancies	$E_\nu(e_1^*\theta S_1 = 1)$	$E_\nu(\mu_1\theta S_1 = 1)$	$E_\nu(\delta\theta S_1 = 1)$
With the program on vacancies opened absent the program	$E_\nu(e_1^*\theta S_0 = 1)$	$E_\nu(\mu_1\theta S_0 = 1)$	$E_\nu(\delta\theta S_0 = 1)$
Without the program on vacancies opened absent the program	$E_\nu(e_0^*\theta S_0 = 1)$	$E_\nu(\mu_0 S_0 = 1)$	$E_\nu(\delta S_0 = 1)$

Note: The theoretical application rate derived from the model and its average are shown for each channel. S indicates selection into the PES vacancy services program. $E_\nu(e_1^*\theta | S_0 = 1)$, $E_\nu(\mu_1\theta | S_0 = 1)$ and $E_\nu(\delta\theta | S_0 = 1)$ are unobservable counterfactuals.

Table 7: Match selection at 2 weeks by channel

	Counselor $E(\mu_1\theta) - E(\mu_0)$ (1)	Employer $E(e_1^*\theta) - E(e_0^*\theta)$ (2)	Jobseeker $E(\delta\theta) - E(\delta)$ (3)	Refusals by counselor (4)
Panel A: Overall				
Treatment	-0.190 (0.367)	-0.781*** (0.279)	-2.407** (0.963)	-0.058 (0.132)
Control Mean	3.502	1.137	6.556	0.519
p-value rank-sum (Mann-Whitney)	0.141	0.036	0.000	0.540
N	1921	1921	1921	1921
Panel B: Heterogeneity				
<u>Contact= 1</u>				
Treatment	-0.566 (0.531)	-0.927* (0.491)	-2.361 (1.521)	-0.242 (0.180)
<u>Contact= 0</u>				
Treatment	0.307 (0.451)	-0.584* (0.304)	-2.454*** (0.906)	0.185 (0.151)
Control Mean(Contact= 1)	3.776	1.320	6.883	0.642
Control Mean(Contact= 0)	3.136	0.894	6.120	0.355
p-value Equality of Coefs.	0.210	0.553	0.958	0.069
N	1921	1921	1921	1921

Note: This table presents inversely propensity weighted (IPW) regression results for the intervention's impact on the number of applicants coming through each channel. We predict vacancy selection into treatment S using our observable vacancy characteristics. $\Pr(S = 1 \mid \text{wage, pred. wage, hours, experience, qualification}) =$

$$F\left(\beta_0 + \beta_1 w + \beta_2 \hat{w} + \beta_3 w * \hat{w} + \beta_4 \text{Low Qual.} + \sum_{h=2}^8 \gamma_h 1(\text{Hours}_h = 1) + \sum_{e=2}^6 \alpha_e 1(\text{Exper}_e = 1)\right)$$

with F the logistic function. We then run an OLS regression of the number of applications in each channel on a treatment indicator with observations weighted by $\frac{T}{\Pr(S=1)} + \frac{1-T}{1-\Pr(S=1)}$. The p-values for Mann-Whitney tests of the equality in distributions are displayed for the overall sample of vacancies. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 8: Estimates of screening rate and caseworker and firm effort ratios

	$\hat{\theta}$	$\widehat{\mu_1/\mu_0}$	$\widehat{e_1^*/e_0^*}$
Overall	0.66	1.53	0.31
Contact=1	0.65	1.32	0.32
Contact=0	0.60	1.90	0.32

Note: These estimates are calculated by applying the estimates from Table 7 to the expressions derived in Table 6.

Appendix Figures and Tables

Figure A.1: Vacancy and contract flows per month



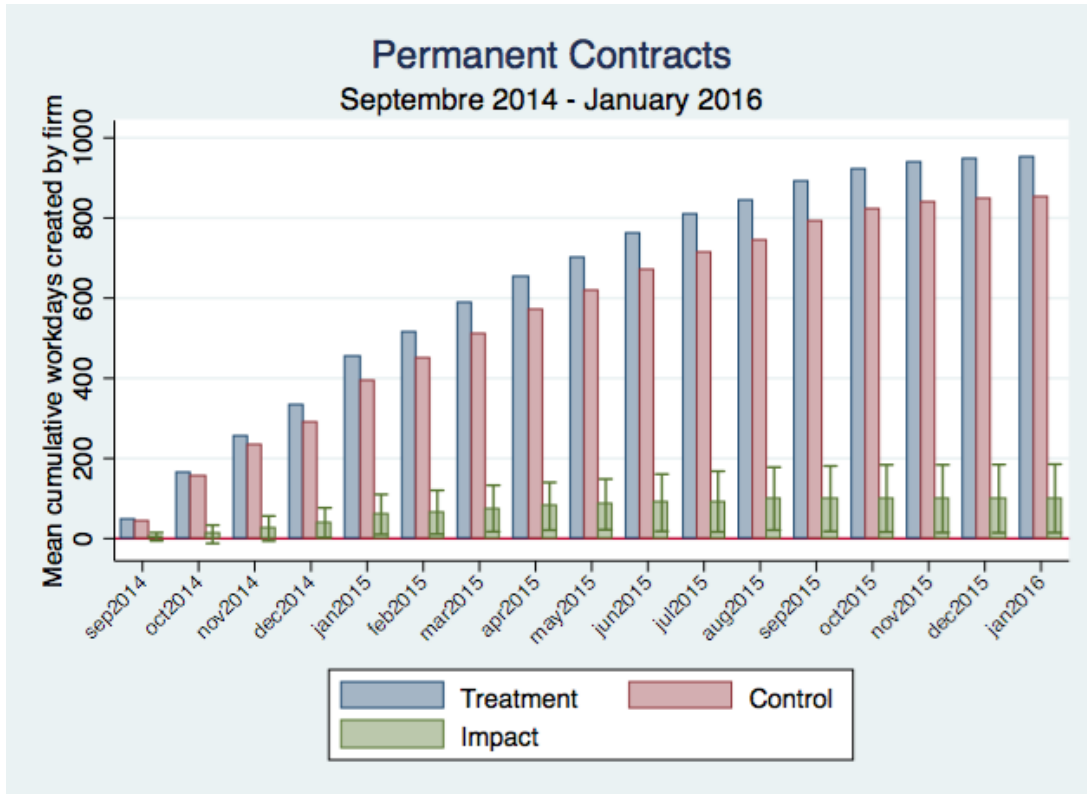
Note: Figures illustrate the dynamic impact of the intervention on vacancies flows (figures in the left column) and contract flows (right column). Vertical bars show the treatment and control group means along with the difference for each month during the 17 month observation period. The differences are overlaid with 95% confidence intervals. Impacts are estimated per equation 1 and standard errors are clustered at the agency level.

Table A.1: Impact on flows using non top-coded data

	Vacancies				Hires			
	Permanent (1)	Fixed-term (2)	Temp (3)	All (4)	Permanent (5)	Fixed-term (6)	Temp (7)	All (8)
Panel A: Overall								
Treatment	0.101*** (0.025)	0.062** (0.024)	0.013 (0.037)	0.177*** (0.057)	0.214** (0.083)	-1.338 (2.910)	0.779 (1.375)	-0.344 (2.994)
Control Mean	0.219	0.166	0.315	0.699	1.460	11.489	10.131	23.080
N	8232	8232	8232	8232	8232	8232	8232	8232
Panel B: Heterogeneity								
<u>Contact= 1</u>								
Treatment	0.177*** (0.059)	0.161*** (0.048)	0.018 (0.078)	0.355*** (0.125)	0.637*** (0.199)	-2.464 (5.094)	6.079* (3.659)	4.251 (6.184)
<u>Contact= 0</u>								
Treatment	0.074*** (0.025)	0.029 (0.022)	0.014 (0.042)	0.117** (0.058)	0.096 (0.083)	2.116 (2.827)	-1.131 (1.587)	1.081 (2.703)
Control Mean(Contact= 1)	0.332	0.196	0.350	0.878	1.738	14.699	10.539	26.976
Control Mean(Contact= 0)	0.161	0.124	0.256	0.541	1.320	8.258	9.743	19.321
p-value Equality of Coefs.	0.120	0.011	0.969	0.086	0.013	0.357	0.070	0.609
N	7859	7859	7859	7859	7859	7859	7859	7859

Note: This table replicates results in Table 2, but uses non top-coded data. Average treatment effects are estimated per equation 1. Strata weighted control group means are also shown. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Figure A.2: Cumulative impact on workday creation in permanent contracts



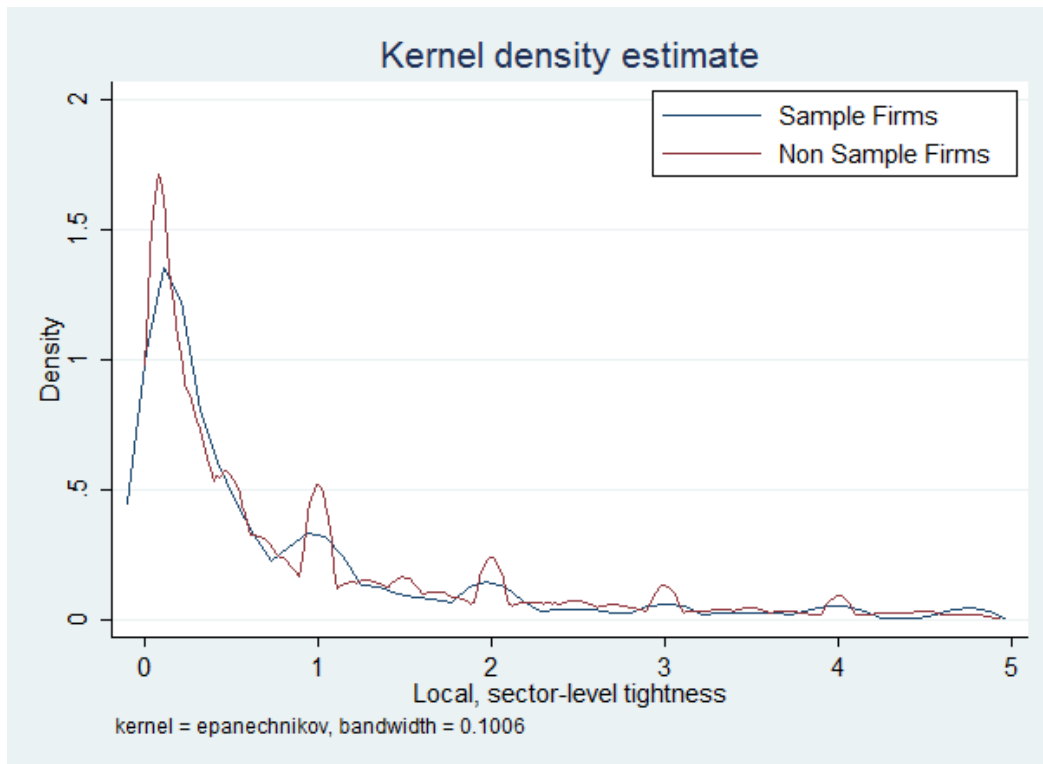
Note: This figure illustrates the dynamic impact on the cumulative number of workdays created in permanent contracts by sample firms. Vertical bars show the treatment and control group means along with the difference for each month during the 17 month observation period. The differences are overlaid with 95% confidence intervals. Impacts are estimated per equation 1 and standard errors are clustered at the agency level.

Table A.2: Impact on flows by period

	(1)	(2)	(3)	(4)
	Permanent	Fixed-term	Temp	All
<hr/>				
Panel A: Vacancies				
<hr/>				
<u>September - March</u>				
Treatment	0.064*** (0.017)	0.048*** (0.013)	-0.001 (0.019)	0.110*** (0.032)
Control Mean	0.202	0.142	0.249	0.592
<u>April - January</u>				
Treatment	-0.021 (0.025)	-0.008 (0.018)	-0.062* (0.035)	-0.091* (0.049)
Control Mean	0.377	0.279	0.509	1.165
<u>September - January</u>				
Treatment	0.052 (0.037)	0.041 (0.026)	-0.062 (0.050)	0.030 (0.073)
Control Mean	0.582	0.423	0.776	1.781
N	8232	8232	8232	8232
<hr/>				
Panel B: Hires				
<hr/>				
<u>September - April</u>				
Treatment	0.118** (0.051)	0.296 (0.380)	-0.084 (0.584)	0.331 (0.643)
Control Mean	1.335	5.257	7.685	14.277
<u>May - January</u>				
Treatment	0.054 (0.068)	0.535 (0.535)	-0.678 (0.694)	-0.089 (0.777)
Control Mean	1.779	7.515	10.407	19.702
<u>September - January</u>				
Treatment	0.188* (0.112)	0.863 (0.916)	-0.714 (1.226)	0.336 (1.370)
Control Mean	3.147	12.854	18.094	34.095
N	8232	8232	8232	8232

Note: This table presents impacts on vacancy postings with the PES (Panel A) and hiring flows (Panel B) for the three contract types and on aggregate. For each panel, average treatment effects on cumulative flows are displayed for the sanctuary period (Sep. 2014 - Mar./Apr. 2015), post-sanctuary period (Apr./May 2015 - Jan. 2016) and all periods (Sep. 2014 - Jan. 2016). We define the end month of the sanctuary period for hiring flows as April 2015 to account for hires linked to vacancies posted in March. Average treatment effects are estimated per equation 1. Strata weighted control group means are also shown. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Figure A.3: Local market tightness pre treatment



Note: This figure displays probability density functions of labor market tightness for firms attached to the 129 local employment agencies' portfolios that participated in the study. Tightness is calculated at the local-sector level as the ratio of vacancies to the stock of registered jobseekers who's searched-for profession corresponds to the sector (see section 3.2 for a description of how we link professions to sectors). Tightness is averaged over the 12 months preceding the intervention. Densities are displayed for sample firms as well as non sample firms attached to the 129 agencies.

Table A.3: Movements between firms

	Fixed-term		Permanent		Temp	
	Control	Treatment	Control	Treatment	Control	Treatment
<u>Individuals hired coming from:</u>						
Unemployment	0.569	0.570	0.353	0.344	0.883	0.839
	61895	57023	5053	5400	74179	91449
Employment in Sample Firm:	0.106	0.093	0.031	0.029	0.003	0.005
	11494	9271	444	448	262	531
From same firm	0.105	0.091	0.030	0.028	0.003	0.005
	11393	9144	432	442	260	521
From treated firm to -	0.000	0.001	0.000	0.000	0.000	0.000
	38	84	7	5	1	5
From control firm to -	0.001	0.000	0.000	0.000	0.000	0.000
	63	43	5	1	1	5
Inactivity or Employment Elsewhere	0.051	0.050	0.218	0.209	0.003	0.049
	5553	5055	3123	3285	243	5335
Unknown	0.274	0.287	0.398	0.418	0.110	0.107
	29808	28763	5693	6550	9277	11662

Note: Data used are the hiring declarations made by all sample firms during the 6 month sanctuary period. Row titles correspond to the origin of the hired individual, columns titles where the hired person was placed and in which type of contract. The proportion of total flows by column and total volume of flows are displayed for each category. Unemployment is defined as jobseekers registered with the PES within the 30 days preceding the hiring date. Employment in sample firms is broken down into three categories: flows within the same firm, flows coming from treatment firms and flows coming from control firms. Inactivity and employment elsewhere is defined as hiring flows for people entering the labor market or who were employed in another firm outside the sample. Unknown is defined as flows for individuals for whom we have no identifiers and thus cannot trace their hiring or unemployment history.

Table A.4: Selection on provision of vacancy services (IPW)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Intensive Follow-up	Preselection	Special Preselection	Verification	Valorization	Evaluation	Analysis of post	Drafting support	Interview support	Adaptation support	EMTPR	PMSMP
Panel A: Overall												
Treatment	0.099*** (0.033)	0.094*** (0.033)	0.102*** (0.031)	0.093*** (0.033)	0.005 (0.008)	0.000 (0.004)	0.009 (0.012)	0.001 (0.021)	0.003 (0.004)	0.002 (0.005)	0.002 (0.003)	-0.002 (0.002)
Control Mean	0.272	0.260	0.223	0.249	0.026	0.007	0.017	0.109	0.004	0.010	0.005	0.003
N	1921	1921	1921	1921	1921	1921	1921	1921	1921	1921	1921	1921
Panel B: Heterogeneity												
<u>Contact=1</u>												
Treatment	0.089* (0.046)	0.084* (0.046)	0.080* (0.043)	0.087* (0.046)	0.002 (0.011)	-0.001 (0.006)	0.015 (0.015)	-0.009 (0.029)	-0.004 (.)	0.000 (0.008)	0.000 (0.005)	-0.001 (0.003)
<u>Contact=0</u>												
Treatment	0.113*** (0.044)	0.109** (0.043)	0.131*** (0.040)	0.101** (0.042)	0.008 (0.009)	0.002 (0.006)	0.002 (0.013)	0.015 (0.030)	0.012 (.)	0.004 (0.006)	0.005 (0.005)	-0.003 (0.003)
Control Mean(Contact=1)	0.284	0.272	0.241	0.254	0.030	0.008	0.014	0.121	0.008	0.015	0.006	0.003
Control Mean(Contact=0)	0.256	0.245	0.198	0.242	0.019	0.006	0.021	0.093	0.000	0.004	0.003	0.003
p-value Equality of Coefs.	0.706	0.692	0.389	0.822	0.705	0.776	0.523	0.562	.	0.692	0.513	0.606
N	1921	1921	1921	1921	1921	1921	1921	1921	1921	1921	1921	1921

Note: This table presents inversely propensity weighted (IPW) regression results for the selection of services to vacancies. We predict vacancy selection into treatment S using our observable vacancy characteristics. $\Pr(S = 1 \mid \text{wage, pred. wage, hours, experience, qualification}) =$

$$F\left(\beta_0 + \beta_1 w + \beta_2 \hat{w} + \beta_3 w * \hat{w} + \beta_4 \text{Low Qual.} + \sum_{h=2}^8 \gamma_h 1(\text{Hours}_h = 1) + \sum_{e=2}^6 \alpha_e 1(\text{Exper}_e = 1)\right)$$

with F the logistic function. We then run an OLS regression of indicators for various services on a treatment indicator with observations weighted by $\frac{T}{\Pr(S=1)} + \frac{1-T}{1-\Pr(S=1)}$. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Table A.5: Match selection at 2 weeks (non IPW estimates)

	Counselor (1)	Employer (2)	Jobseeker (3)	Refusals by counselor (4)
Panel A: Overall				
Treatment	-0.001 (0.358)	-0.729*** (0.254)	-2.014** (0.877)	-0.024 (0.123)
Control Mean	3.435	1.060	6.048	0.515
p-value rank-sum (Mann-Whitney)	0.141	0.036	0.000	0.540
N	2052	2052	2052	2052
Panel B: Heterogeneity				
<u>Contact= 1</u>				
Treatment	-0.394 (0.518)	-0.905** (0.445)	-1.893 (1.393)	-0.176 (0.165)
<u>Contact= 0</u>				
Treatment	0.508 (0.432)	-0.496* (0.265)	-2.157*** (0.811)	0.171 (0.140)
Control Mean(Contact= 1)	3.679	1.260	6.406	0.592
Control Mean(Contact= 0)	3.115	0.798	5.579	0.415
p-value Equality of Coefs.	0.182	0.429	0.870	0.108
N	2052	2052	2052	2052

Note: This table is analogous to Table 7, but uses simple OLS regression estimates with no IPW. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Table A.6: Match Selection at 8 weeks (IPW estimates)

	Counselor (1)	Employer (2)	Jobseeker (3)	Refusals by counselor (4)
Panel A: Overall				
Treatment	-0.158 (0.588)	-0.857*** (0.296)	-2.613** (1.000)	-0.073 (0.153)
Control Mean	4.741	1.260	7.094	0.609
p-value rank-sum (Mann-Whitney)	0.009	0.064	0.000	0.594
N	1921	1921	1921	1921
Panel B: Heterogeneity				
<u>Contact= 1</u>				
Treatment	-0.726 (0.884)	-0.975* (0.524)	-2.591 (1.608)	-0.286 (0.205)
<u>Contact= 0</u>				
Treatment	0.593 (0.617)	-0.697** (0.323)	-2.627*** (0.949)	0.208 (0.183)
Control Mean(Contact= 1)	5.147	1.437	7.465	0.733
Control Mean(Contact= 0)	4.200	1.024	6.598	0.444
p-value Equality of Coefs.	0.221	0.651	0.984	0.072
N	1921	1921	1921	1921

Note: This table is analogous to Table 7, but counts applications made up to 8 weeks after the vacancy posting date. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Table A.7: Vacancy distribution

	(1)	(2)	(3)	(4)
	Internet	Anonymous	Firm name and Contact info	Firm name only
Panel A: Overall				
Treatment	-0.078** (0.038)	0.053 (0.035)	-0.044 (0.043)	0.004 (0.025)
Control Mean	0.584	0.311	0.619	0.091
N	2052	2052	2052	2052
Panel B: Heterogeneity				
<u>Contact= 1</u>				
Treatment	-0.059 (0.056)	0.054 (0.047)	-0.038 (0.063)	0.002 (0.041)
<u>Contact= 0</u>				
Treatment	-0.102** (0.049)	0.052 (0.043)	-0.051 (0.045)	0.006 (0.018)
Control Mean(Contact= 1)	0.588	0.294	0.617	0.115
Control Mean(Contact= 0)	0.579	0.333	0.623	0.060
p-value Equality of Coefs.	0.567	0.971	0.867	0.935
N	2052	2052	2052	2052

Note: This table presents results from simple OLS regressions of indicators for firm selection on the way vacancies are posted on a treatment indicator. Internet signifies whether the vacancy is posted publicly on pole-emploi.fr. Anonymous indicates that the vacancy does not show any identifying information on the firm. "Firm name and contact info" indicates that the firm displays all contact information in the vacancy. "Firm name only" indicates that the vacancy only displays the name of the firm and no contact info. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Table A.8: Intensity and actor

	Number of modifications made to vacancy				Actor		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Counselor	Firm	Automatic	By internet	Vacancy created by employer	Vacancy posted by 3rd party	Posted with PES web space
Panel A: Overall							
Treatment	0.859*** (0.312)	-0.047 (0.051)	-0.206 (0.160)	-0.201* (0.116)	0.028 (0.017)	-0.033 (0.039)	0.031 (0.025)
Control Mean	3.324	0.368	2.708	1.741	0.096	0.264	0.232
N	2052	2052	2052	2052	2052	2052	2052
Panel B: Heterogeneity							
<u>Contact= 1</u>							
Treatment	0.773* (0.437)	-0.110 (0.069)	-0.041 (0.223)	-0.203 (0.155)	0.027 (0.021)	-0.080 (0.060)	0.024 (0.035)
<u>Contact= 0</u>							
Treatment	0.971*** (0.370)	0.035 (0.071)	-0.418** (0.188)	-0.197 (0.191)	0.029 (0.026)	0.027 (0.048)	0.040 (0.034)
Control Mean(Contact= 1)	3.375	0.390	2.698	1.788	0.081	0.271	0.223
Control Mean(Contact= 0)	3.257	0.339	2.721	1.680	0.115	0.254	0.243
p-value Equality of Coefs.	0.729	0.145	0.197	0.981	0.955	0.165	0.750
N	2052	2052	2052	2052	2052	2052	2052

Note: Modifications during the life of the vacancy can either be made by the counselor or the firm. Automatic signifies that it is simply PES computer system that automatically cancels the vacancy after a certain length of time with no activity. Modified by internet means that the firm made its modification by internet. The vacancy can also be created and posted by the firm or through a 3rd party actor on behalf of the firm via its personal web space. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Variance Computation Appendix

In their paper on the efficient estimation of average treatment effects using propensity scores, Hirano et al. (2003) derive an influence function of the estimate which is very useful to compute the standard error of propensity weighted estimates. Their framework can be directly used here. The case we considered is, however, far more simple than the general case considered in their paper in which the propensity function is a complicated function of the covariates entering the conditional independence assumption. In our case, the influence function is simply

$$\psi(y, s, T) = (y - \hat{\mu}_{1,s}) \frac{T}{e_s} - (y - \hat{\mu}_{0,s}) \frac{1-T}{1-e_s} + \hat{\mu}_{1,s} - \hat{\mu}_{0,s} \quad (9)$$

The sample variance of the function $\psi(y, s, T)$ can be rewritten as $\psi(y, s, T) = \widehat{ATE} + \varepsilon(y, s, T)$ and an estimate of the variance is

$$\widehat{V}(\widehat{ATE}) = \frac{1}{N^2} \sum_i \varepsilon(y_i, s_i, T_i)^2 \quad (10)$$

Throughout our paper we cluster the standard errors to correct for possible correlation in outcomes at the local agency level a , hence our estimate of the variance will simply be,

$$\widehat{V}(\widehat{ATE}) = \frac{1}{N^2} \sum_{a=1}^A \left(\sum_{i \in a} \varepsilon(y_i, s_i, T_i) \right)^2 \quad (11)$$

As noted in the main text, we will systematically display heterogeneous treatment effects along the baseline in-contact dimension $c \in \{0, 1\}$. This estimation is straight forward as it simply involves separately estimating the ATE and influence functions for the two subsamples.⁴⁰ Nevertheless, there will be some strata s in which we have no variation in the heterogeneity dimension. Hence $\hat{\mu}_{T,s,c=1}$ or $\hat{\mu}_{T,s,c=0}$ may be undefined and observations for which this is the case are dropped from the heterogeneity analysis.

Proof Appendix

Propositions:

1. $\Pi(\nu, \mu_1, 1) > \Pi(\nu, \mu_0, 0)$ as long as $\mu_1 \geq \mu_0$

⁴⁰We compute the simultaneous covariance matrix and report results of a Wald test $H_0 : ATE_{c=1} = ATE_{c=0}$

2. $d_1^* = 1$

3. $e_1^* \leq e_0^*$ and $t_1^* \geq t_0^*$

4. Value function under preselection is increasing in v , δ , γ and decreasing in κ

Proofs:

Consider the function $\Pi(\nu, \mu, x)$:

$$r\Pi(\nu, \mu, x) = \max_{e, \bar{t}, d \in \{0,1\}} \left\{ -c(e) + e\theta \left(\gamma \int_{t>\bar{t}}^1 (tv - \Pi(\nu, \mu, x)) \frac{dt}{\theta} - \kappa \right) + d(\delta + \mu)(1 - (1 - \theta)x) \left(\gamma \int_{t>\bar{t}}^1 (tv - \Pi(\nu, \mu, x)) \frac{dt}{1 - (1 - \theta)x} - \kappa \right) \right\}, \quad (12)$$

with FOCs w.r.t. \bar{t} and e :

$$\begin{aligned} e: \quad & -c'(e) + \left(\gamma \int_{t>\bar{t}}^1 (tv - \Pi(\nu, \mu, x)) dt - \theta\kappa \right) = 0 \\ \bar{t}: \quad & \bar{t}v - \Pi(\nu, \mu, x) = 0 \\ & \text{or } \bar{t} = 1 - \frac{1}{\gamma} \text{ if } (1 - \frac{1}{\gamma})v \geq \Pi(\nu, \mu, x) \end{aligned} \quad (13)$$

Looking at the condition whether to include the counselor and jobseeker channels into the search process,

$$\gamma \int_{t>\bar{t}}^1 (tv - \Pi(\nu, \mu, x)) dt - (1 - (1 - \theta)x)\kappa \geq 0 \quad (14)$$

under preselection, $x = 1$, this condition rewrites

$$\gamma \int_{t>\bar{t}}^1 (tv - \Pi(\nu, \mu, x)) dt - \theta\kappa \geq 0 \quad (15)$$

This shows point 2: under preselection, the jobseeker and caseworker channels are always active: $d_1^* = 1$

Without preselection, $x = 0$, the condition rewrites

$$\gamma \int_{t>\bar{t}}^1 (tv - \Pi(\nu, \mu, x)) dt - \kappa \geq 0 \quad (16)$$

\Rightarrow Without preselection, there might be cases in which the jobseeker or caseworker channels are not active: $d_0^* = 0$

Let us show point 1:

We easily get the following expressions for the derivatives with respect to x and μ : For x , we get

$$r\Pi'_x = -(e + d(\delta + \mu))\gamma\Pi'_x(1 - \bar{t}^*) + d(\delta + \mu)\kappa(1 - \theta)$$

Thus if $d = 0$, $\Pi'_x = 0$ and if $d = 1$ $\Pi'_x > 0$, this is enough to ensure $\Pi(\nu, \mu_0, 1) > \Pi(\nu, \mu_0, 0)$
For μ , we get

$$r\Pi'_\mu = -(e + d(\delta + \mu))\gamma\Pi'_\mu(1 - \bar{t}^*) + d(\gamma \int_{t > \bar{t}}^1 (tv - \Pi(\nu, \mu, x)) dt - (1 - (1 - \theta)x)\kappa)$$

Thus if $d = 0$, $\Pi'_\mu = 0$ and if $d = 1$, $\Pi'_\mu > 0$, this is enough to ensure $\Pi(\nu, \mu_1, 1) > \Pi(\nu, \mu_0, 1)$
These results show $\Pi(\nu, \mu_1, 1) > \Pi(\nu, \mu_0, 0)$

Point 3 directly follows:

The first order condition in \bar{t} (see equation (14)) writes as: $\bar{t}v = \Pi(\nu, \mu, x)$, as $\Pi(\nu, \mu_1, 1) > \Pi(\nu, \mu_0, 0)$ we directly get $t_1^* > t_0^*$

The first order condition in e (see equation (14)) rewrites as $c'(e) = \gamma(1 - t^*)^2/2 - \theta\kappa$, thus the previous point implies $e_1^* < e_0^*$

Turning to point 4, we can now derive the value function with respect to each of the components of ν :

$$[r + (e_1^* + \delta + \mu)\gamma(1 - t_1^*)] \begin{pmatrix} \pi'_\delta \\ \pi'_v \\ \pi'_\kappa \\ \pi'_\gamma \end{pmatrix} = \begin{pmatrix} \int_{t > \bar{t}}^1 (tv - \Pi(\nu, \mu, x)) dt - \kappa\theta > 0 \\ (e_1^* + \delta + \mu)\gamma(1 - t_1^{*2})/2 > 0 \\ -(e_1^* + \delta + \mu)\kappa < 0 \\ \int_{t > \bar{t}}^1 (tv - \Pi(\nu, \mu, x)) dt > 0 \end{pmatrix} \quad (17)$$

Part III

Job Search and Intermediation under Discrimination: Evidence from Terrorist Attacks in France¹

Abstract

Using detailed, high frequency data on potential job matches made through the French Public Employment Service, I present evidence showing that search intensity both by minority jobseekers, as well as by their job counselors and by firms, is highly sensitive to a shock that increases discrimination against their type. On average, minority jobseekers – defined as having a first name of Maghreb or Mashriq origin – reduce their job search effort by 11% in the 10 weeks following the January 2015 “Charlie Hebdo” attacks compared to majority jobseekers, defined as those with French sounding first names. Employers also reduce their search effort for minority candidates, but only for their high quality contracts. This drop is offset by an increase in counselor matching effort made for minorities after the shock, but only in areas with low latent levels of discrimination. In addition this counselor effect is much stronger for counselors who are themselves minorities and for majority counselors who specialize in getting the most marginalized jobseekers back to work. Looking at total work creation, I find little evidence that minorities found less work than majorities in the weeks that followed the attack. This suggests that labor market intermediaries may play an important role in mitigating the adverse effects of discrimination.

¹I would like to thank Yann Algan and Bruno Crépon for their valuable input and encouragement. Also a very special thank you to Anita Bonnet and Cyril Nouveau at Pôle Emploi for their support and tenacity. I have also greatly benefited from helpful discussions with Guillaume Blache, David Buchner, Johana Carrier, Amanda Pallais, Arnaud Philippe, Victor Pouliquen, Alexandra Roulet and Benoît Sadrin.

1 Introduction

A broad literature on the effects of discrimination on the labor market outcomes of minorities has been expanding since the seminal work of Becker (1957). This work, *The Economics of Discrimination*, viewed the effects of discrimination, defined as differential treatment of equally productive minority and non-minority workers, as a result of the disutility employers experience when employing minorities. Due to this “distaste,” minority workers must compensate the employer’s bias by either being more productive than non-minorities or by accepting inferior wages for equal productivity. This conception of discrimination was followed by Phelps (1972) and Arrow (1973), who tackled discrimination not as problem of taste, but one of imperfect information. Employer’s beliefs or priors about the average productivity of a minority group may lead to the unequal treatment of equally qualified workers. To date, these works, and the research that built on their insights, have focused primarily on the interaction of workers with the employer, either during the hiring process (start with Lang and Lehmann (2012) for review of the theory literature and Bertrand and Duflo (2017) for applied field experiments) or, more recently, on-the-job (Hjort (2014) and Glover, Pallais, and Pariente (2017)). A topic that has received less attention is how discrimination, or the perception of it, might affect job search itself.

In this study, we examine the effects on labor market outcomes of a shock that potentially dramatically increased discriminatory tastes towards a specific minority group. Though there is ample evidence that discrimination against Muslim minorities exists in France (Adida, Laitin, and Valfort, 2014), we exploit the terrorist attacks in and around the Paris region on the Charlie Hebdo satirical newspaper, the police and the Hypercacher supermarket between the 7th and 9th of January 2015, as an exogenous shock that may have have substantially increased bias against Muslim minorities in France. To illustrate the link between these terrorist attacks and an increase in potential discrimination we refer to Figure 1 which plots Google search interest in France for the word “islamophobie,” the equivalent of islamophobia in English, over the 18 months spanning the January 2015 attacks and those of November 13th.² We see that the search interest is close to null until the week of the January attacks to which it quickly jumps to 100 in the week of the attack, the date by which all other points are normalized. And though we are not able to ascertain exactly why or who searches for this term, the figure indicates that bias against Muslim minorities came prominently into

²Search interest is calculated as (number of searches for term) / (total Google searches). The search interest score is then normalized to the date with the highest search interest. To understand the relative volume of the search for islamophobie we refer to Figure A.1 in the appendix which plots search interest in islamophobie compared to the search interest for “trouver un emploi”, translated as “find a job” and to Christine and the Queens, a popular French pop star who’s debut album was released in early 2015.

the public consciousness in the days and weeks that followed the attack.

We focus our analysis on the effects of this shock on the search intensity of three labor market actors: jobseekers registered with the Public Employment Service (PES), their job counselors and employers. Using a double difference-in-differences strategy whereby we control for existing differential seasonal minority and majority group trends, we find that the shock led to a large transitory decline in minority search effort compared to majorities, as measured by the average number of applications per jobseeker, or potential matches, they make to offers posted with the PES. On average minority jobseekers, defined as individuals with Maghreb/Mashriq-sounding first names reduce their search intensity by 11% compared to majorities – those with French-sounding first names – in the 10 weeks following the terrorist attacks. In contrast, we find that job counselors increased the number of potential matches for minorities compared to majorities in the weeks that followed the attack, an increase of 13%. Importantly, this counselor effect is stronger for counselors who are themselves minorities and for majority counselors whose job entails working with the most marginalized jobseekers. We also find that the shock induces employers to make less overtures to minorities for their job offers, but only for jobs that are in the most prized type of contract, a permanent contract. We find that they increase their potential matches to minorities for their fixed-term and temp work. Over all contract types we see no impact on employer initiated matches due to the shock.

Given these large impacts on key channels of job search and matching, we make the case that attributing impacts measured on hiring outcomes exclusively to increased employer discrimination may be imprudent. Hence throughout the paper, we make a conceptual distinction between the effects of an increase in discrimination on job search and its effect on employment outcomes. We make the distinction between these “first” and “second order” outcomes because our identification strategy does not allow us to distinguish between a change in employment outcomes for minorities due to the increased discrimination they may face by employers from changes in their employment outcomes that pass through a change in their own search behavior or that of their counselors.

This paper thus contributes to the literature that concerns itself with how to measure, and indeed, the very existence of discrimination in the labor market. Correspondence studies (see Riach and Rich (2002) for an introduction) generally find large discrepancies in call back rates for minorities with identical resumé as compared to majorities (see Bertrand and Mullainathan (2004) in the US and Behaghel et al. (2015), Petit et al. (2011), Berson (2012) and Adida et al. (2016) for the French context). Yet, Heckman (1998) has made the point that the results of correspondence studies are, at the very least, difficult to interpret as indicating the presence of discrimination in the labor market because they are based on

the premise that workers apply randomly to jobs. The results presented in this paper on these first order outcomes show that the impact of discrimination on search behavior may be substantial and thus it is difficult to disentangle the direct effect of discrimination on employment outcomes due to employer bias from the search intensity of, or for, minorities. A secondary contribution of this paper is hence based on the fact that we actually observe real employer search behavior for firms that recruit through the PES.

We illustrate the difficulty of separating the effect of changes in search intensity from increased employer discrimination in our analysis using the shocks effect on hires. We find that after the shock minorities enter into very short-term contracts disproportionately to majorities, but this effect is driven by a relative drop in the use of these contracts by majorities. Hence a possible explanation could be that due to the increased effective labor market tightness associated with the drop in the minority application rate for permanent contracts, majorities have a larger job finding rate in these contracts and thus rely less on short-term contracts. Another explanation could simply be that because employers are less willing to hire minorities in permanent contracts, they substitute to majorities and that this is the driving mechanism for the differential in short-term hires. Unfortunately we cannot distinguish between the two, and possibly many other scenarios driving employment dynamics, therefore throughout the paper we will remain agnostic about the channels driving changes in hiring dynamics following the shock. The one conclusion for which we can provide supporting evidence, is that when aggregating work days over all contract types, we find no evidence that the shock led to less job creation for minorities as compared to majorities.

On a more basic level, this study also adds to the literature on the effects of ethno-religious terrorist attacks on market outcomes for Muslim minorities. Previous results are mixed when it comes to the effects on employment outcomes. Åslund and Rooth (2005) find that attitudes towards Muslim minorities in Sweden changed after the 9-11 terrorist attacks, but find no evidence that this translated into worse employment outcomes for these minorities. Similarly Kaushal et al. (2007) find no employment effect on Arab minorities in the US after 9-11, but find that they may have experienced lower wages compared to majorities in the years that followed. In addition, Gautier et al. (2009) find that local housing prices diminish in minority neighborhoods and residential segregation increases following an ethno-religious terrorist attack in the city of study, reflecting similar results found by Ratcliffe and von Hinke Kessler Scholder (2015) who also show that these types of attacks may have had a negative effect on employment levels in neighborhoods with a higher percentage of minority residents. We hope that this paper thus contributes to a better understanding of how to interpret the results on the second order outcomes analyzed in these studies.

We refine the analysis in this paper using a measure of existing, or latent discrimination

at the local level. We use the municipality level vote share for the far-right political party in France, the Front National (FN), in the 2012 French presidential elections. We use these data to proxy for the existing discrimination that jobseekers may face in their job search. We report strong heterogeneous effects of the discrimination shock for jobseekers and counselors across this dimension. While we see a negative effect on a minority’s own search effort in both high and low FN areas, the impact is about three and half times as large in *low* FN areas. This heterogeneous effect across existing discrimination is counterbalanced by the effort put forth by counselors, but only in low FN areas. We motivate the interpretation of these, perhaps, striking results by modeling the marginal effect of discrimination as diminishing over the support of latent discrimination. Jobseekers’ cost to search is endogenous to the job finding rate which is affected by real and/or perceived discrimination in the market. Since discrimination has diminishing returns, the marginal effect of a large shock to discriminatory tastes should be larger in areas with low initial levels of discrimination. For counselors, we interpret their effort as an input in the efficiency parameter of a labor market matching function. If their effort is a tension between increasing effort for jobseekers who are having a harder time finding a job as well as the cost of their own discrimination, such heterogeneous empirical findings are quite plausible.

To support the use of the FN vote share as a proxy for existing discrimination we return to the use of Google search trends data. We first look at search trends in the year before the shock, disaggregated by French region. We find that the FN vote share is indeed strongly correlated with Google searches that connote the prevalence of discrimination and discriminatory animus towards our minority group. We then measure the change in search trends around the date of the January attacks. We show that the volume for these negative search terms sharply increases due to the shock and continue to be highly positively correlated with the FN vote share. We then look at terms that connote social cohesion and see a large jump in interest during the shock, but in stark contrast, these terms are negatively correlated with the FN vote share.

Though only speculative, this may be evidence of the important role public employment services play in mitigating the effects of group-level adverse shocks if they can internalize the increased difficulties that minority jobseekers face. And that this “compensatory effect” made by job counselors may be linked to preferences for integration and social cohesion.

2 Data

Administrative data

We have access to rich historical administrative data at the daily level for just under 4.4

million jobseekers registered with the PES over the 10 weeks before and after the shock that we can categorize as either majority or minority. We have their personal characteristics, their hires and the potential matches, i.e. applications made to job offers posted with the PES. Importantly, we can separate the potential matches that concern these jobseekers into three channels: potential matches initiated by the jobseeker, a PES counselor or by an employer. And though these potential matches made through the PES are not the only way candidates and employers match in the French labor market, the hiring data is a near exhaustive measure of job creation in the market. All firms are required to submit a hiring declaration before, or shortly after, the contract start date, thus we have a reliable measure of job creation for the entire population of jobseekers.³

We exploit the date, channel and the group status of the jobseeker concerned by the potential match as our key outcome variables for the analysis: We thus, measure each labor market actor’s search effort in the weeks before and after the shock.

Using the hiring declarations, we extract the contract type, its start- and end-dates (for fixed-term contracts) and the personal identifier to link the hire to jobseekers on the PES roster. Using the start and end dates for fixed-term contracts we calculate the number of workdays created for each contract. For permanent contracts, we censor the end-date at the end of the observation period. We do this because these declarations are contract flows and thus are not directly a measure of employment. For example, a week of one-day (Monday to Saturday) hires for the same individual would be counted as 6 fixed-term contract flows, but as only one contract if it were a fixed-term contract that ran for the week. By calculating workdays, it allows us to have a standardized measure of employment creation over all contract types, and thus measure the shock’s effect on real job creation.

Names

To assign minority status we use the first names of jobseekers available in the PES administrative data. We do this because it is illegal to collect data on ethnicity in France. The first name data are the same used in Algan, Mayer, Thoenig et al. (2013) and Behaghel, Crépon, and Le Barbanchon (2015) who, like in this study, use it as a proxy for the origin or ethnic background of an individual.⁴ This data set links 23,388 first names to nine etymological origins: French, Maghreb/Mashriq, sub-Saharan African, Asian, British, Germanic, Jewish,

³These hiring declarations, required by French law, are called “Déclaration préalable à l’embauche”. Exceptions to the requirement for this hiring declaration concern internships and volunteer contracts and for the recruitment of private child care professionals and some public sector jobs.

⁴Glover, Pallais, and Pariente (2017) also use names to identify minority status in France. For a discussion about how naming is also related to preferences for social integration and cultural transmission, start with Algan et al. (2013) who estimate the economic penalties associated with having a first name of Maghreb/Mashriq origin.

Southern European and Eastern European. The categorization was compiled using register data on birth names given to French babies from 2003-2007 by Algan et al. (2013). We define majority status as those jobseekers with French first names and minority status as those with Maghreb/Mashriq first names and will also exploit British and Southern European names as a robustness check.

A measure of existing discrimination

We contrast our results through a local measure of discrimination. To do this, we use the vote share for the Front National (FN) party in the first round of the 2012 French presidential elections. We use the 36,565 commune-level, or municipality-level, vote shares for the FN. Each commune is administratively attached to a PES agency. We thus aggregate the total vote share for the FN in all municipalities in the agency's purview and assign it this score. We thus have local-level variation in a proxy for existing discrimination.⁵

Data structure

We take PES administrative data on the number of potential matches and hires at the daily-jobseeker level, combined with names and vote data, and aggregate it to the week and employment agency level for the 10 weeks before and 10 weeks after the shock.⁶ We thus have two observations per week for each local employment agency, one for the minority and one for the majority population in the agency. In total we follow 810 local employment agencies throughout mainland France and Corsica over the 20 weeks spanning the attack, providing 68,400 minority/majority-week-agency observations.

Google Trends

Finally, to determine whether the shock and FN vote are correlated with discriminatory attitudes, we use Google trends data on search volume over time for all of France and comparatively across French regions. Specifically, we look at key search terms before and after the date of the shock that may be correlated with the prevalence of discrimination and animus against minorities as well as search terms that connote social cohesion. We then look at the correlation of these search scores and the FN vote at the regional level to understand what our local measure of discrimination proxies for.

⁵We will discuss at length and provide empirical evidence on the link between the FN vote share and discriminatory attitudes in the sections below.

⁶We aggregate because it is computationally more efficient given that the shock is aggregate and, at most, the variation in discrimination we will be utilizing is at the agency level i.e variation in the FN vote share.

2.1 Distribution statistics

Table 1 displays distribution statistics for jobseekers at the employment agency level. Statistics are averaged over the 10 week pre-shock weeks by agency. Column 1 shows the overall average proportion per agency while columns 2 and 3 show the relative proportion within majority and minority populations, respectively. In examining the typology of jobseekers registered with the PES we see that around 71% are currently unemployed and looking for full time work in a permanent contract. Comparing across groups, minorities are about 8 percentage points more likely to be in this category than majorities. The next most frequent type of jobseeker are those looking for part-time work followed by individuals looking for fixed-term, temp or seasonal work and those engaging in on-the-job search. Minorities are relatively less likely to fall into these last two categories.

Perhaps unsurprisingly, we see that while 98% of jobseekers with French sounding first names are French nationals only 62% of jobseekers with Maghreb/Mashriq first names are. We also see that we categorize zero jobseekers who are born in the Maghreb as majorities while 30% of our minority jobseekers are Moroccan, Algerian or Tunisian nationals.

Turning to demographic characteristics, we see that minority jobseekers are more likely to be male, less likely to have a college degree or be categorized as skilled labor (high qualification), but are about 18 percentage points more likely to live in a Sensitive Urban Zone (Zone Urbaine Sensible or ZUS). These are residential zones that have been prioritized by the French government to receive additional resources and funding because they exhibit significantly higher levels of unemployment and poverty.⁷

We also see significant differences between groups in the type of professions jobseekers are looking to work in, suggesting that there is rather large occupational segregation in the French context.⁸ Minorities are less likely to search in commerce and sales, IT, accounting, human resources and secretarial work, while they are more likely to look for jobs in construction, transport and “personal services” which may include child care services or living assistance for the elderly.

In examining jobseeker applications, or potential matches, we also see stark contrasts between groups. The average jobseeker application rate (number of applications divided by number of jobseekers) is about 20 points higher for minorities compared to majority jobseekers. The average rate for counselor initiated matches is also four points higher for minorities.⁹ Finally, perhaps surprisingly, we see very little pre-shock difference between

⁷In 2014, the Observatoire National des Zones Urbaines Sensibles recorded a total of 751 ZUS in France in which 4.4 million people live. See http://publications.onz.us.fr/rapport_2014 for more information.

⁸See Aeberhardt et al. (2010) for a discussion on the types of jobs taken-up by minorities and its relationship with wage differentials between groups.

⁹This baseline difference is interesting and may speak to the underlying differences in demographic charac-

groups in the average rate of potential matches initiated by employers.

For hires we see no difference across groups in the rate of being hired in the most prized type of contract, a permanent contract, but see large differences for the other types of contracts: -13 points for fixed-term contracts and +12 points for temp, or interim, contracts for minorities compared to majorities. These last statistics are also interesting when we keep in mind the level of pre-shock discrimination in the market. Fixed-term contracts are very hard to break legally before the end-date by either the employer or the employee unless a better contract has been found in the mean time, while temp work requires no contractual agreement between the individual and the employer, only between the temp agency and the employer. Thus personnel can be easily changed at the request of the employer. Unfortunately, these hiring dynamics are beyond the scope of this paper.

Finally in describing the sample, we look at the average number of jobseekers registered per local employment agency. On average there are 972 minorities per 5428 registered jobseekers. And looking through our main dimension of discrimination heterogeneity, we see that this proportion is lower for agencies that have below-the-median vote shares for the FN, but not dramatically so.

3 Identification and empirical specification

We will assume that the terrorist attacks were completely unforeseen by the French population, thus we will not concern ourselves with labor market actors modifying their behavior in anticipation of the shock. But since the shock could have had aggregate economic impacts, we want to control for this. A difference in differences approach is advantageous because it allows us to isolate the effect of a discrimination shock on minorities as compared to compared to a control group over time. Adopting the framework of Angrist and Pischke (2008) we assume that the effect of the discrimination shock, β , is additive to a group and time effect,

$$E(y_1|m, t) = E(y_0|m, t) + \beta = \theta_m + \delta_t + \beta$$

teristics between majority and minority jobseekers. But it may also be evidence of pre-existing discrimination in the market, i.e. more search effort is required for minorities to find a job, on average, a point made by the correspondence studies cited above. At any rate, it appears that minority jobseekers are more reliant on the the PES platform for their job search. The differential in counselor initiated matches is also intriguing and could be linked to the fact that skills searched for in vacancies posted with the PES are, on average, better matched to the skills of minority jobseekers. Yet we cannot exclude the hypothesis that counselors treat minorities differently, a point we will we come back to in the discussion on the impact of the shock.

where y_0 is the outcome in absence of the shock and θ the specific group effect: $m = 1$ for minority and $m = 0$ majority. The time effect is δ for the pre- or post-shock periods, $t \in 0, 1$.

Difference in differences will identify the effect of the shock on the minority group,

$$\begin{aligned} DD &= [E(y|m = 1, t = 1) - E(y|m = 1, t = 0)] \\ &\quad - [E(y|m = 0, t = 1) - E(y|m = 0, t = 0)] \\ &= \beta, \end{aligned}$$

if the time trend ($\delta_{t=1} - \delta_{t=0}$) is assumed to be non group specific. But if it is period specific by group, γ_{mt} , we have,

$$DD = \beta + (\gamma_{m=1,t=1} - \gamma_{m=1,t=0}) - (\gamma_{m=0,t=1} - \gamma_{m=0,t=0}).$$

The DD estimator will be biased whenever $(\gamma_{m=1,t=1} - \gamma_{m=1,t=0}) \neq (\gamma_{m=0,t=1} - \gamma_{m=0,t=0})$. This simply formalizes the standard common trends assumption that groups would need to evolve in the same way absence the shock.

To see if a DD approach is appropriate in our context we refer to Figure 2. This figure plots the evolution of key outcome variables: the average number of potential matches made by jobseekers, counselors and employers for our time period of interest, but in the previous year: 2013-2014. Outcomes are binned at the weekly level for majority and minority populations and these points are fitted using a regression with a polynomial time trend of order 3. In this “placebo year” we see that $t = 0$ weeks exhibit strong parallel trends. But in the $t = 1$ period we see stark diverging trends between groups for jobseeker and counselor initiated potential matches. This suggests that there may be a strong seasonal effect that impacts minorities differentially to majorities in the beginning of the year regardless of the presence of any discrimination shock. In Table 1, we highlighted large differences in key characteristics between the majority and minority populations of jobseekers as well as significant levels of occupational segregation. Hence it may be that job search and employment exhibit differential time effects between groups that are correlated with these observable characteristics, but also other, potentially important, unobservables. Thus it appears that the implementation of a standard DD approach may be hazardous because it will confound any impacts due to a discrimination shock with natural transitory variation between groups in the $t = 1$ period. Furthermore, conditioning parametrically using pre-shock outcomes and characteristics in order to improve the credibility of counterfactual parallel trends assumption is only a partial solution due to the problem of unobservables. Furthermore, since we observe parallel trends

in the expectation function in the pre-shock $t = 0$ period, it is not readily apparent how we might gauge the validity of controlling parametrically for group differences that are correlated with the time effect.

Because we want to compare outcomes between these two groups to isolate a discrimination effect, we can relax this assumption non parametrically by doubling the difference in differences (DDD) with the placebo year. Thus imposing the structure as

$$E(y_1|m, t, T) = E(y_0|m, t, T) + \beta = \theta_m + \gamma_{mtT} + \beta$$

where $T = 1$ indicates that we are in the year where the shock takes place and $T = 0$ the preceding “placebo” year. The first *DD* gives

$$DD_{T=1} = \beta + (\gamma_{m=1,t=1,T=1} - \gamma_{m=1,t=0,T=1}) - (\gamma_{m=0,t=1,T=1} - \gamma_{m=0,t=0,T=1})$$

and the second from the placebo year,

$$DD_{T=0} = (\gamma_{m=1,t=1,T=0} - \gamma_{m=1,t=0,T=0}) - (\gamma_{m=0,t=1,T=0} - \gamma_{m=0,t=0,T=0})$$

Thus,

$$DD_{T=1} - DD_{T=0} = \beta$$

if either time trends would have been equivalent *between groups* in each year - the standard *DD* assumption - or if

$$\begin{aligned} & (\gamma_{m=1,t=1,T=1} - \gamma_{m=1,t=1,T=0}) - (\gamma_{m=0,t=1,T=1} - \gamma_{m=0,t=1,T=0}) \\ & = (\gamma_{m=1,t=0,T=1} - \gamma_{m=1,t=0,T=0}) - (\gamma_{m=0,t=0,T=1} - \gamma_{m=0,t=0,T=0}). \end{aligned}$$

And this expression we can be simply rewritten as a de-trended group specific time effect,

$$\gamma_{m=1,t=1}^{DT} - \gamma_{m=0,t=1}^{DT} = \gamma_{m=1,t=0}^{DT} - \gamma_{m=0,t=0}^{DT} \iff \gamma_{m=1,t=1}^{DT} - \gamma_{m=1,t=0}^{DT} = \gamma_{m=0,t=1}^{DT} - \gamma_{m=0,t=0}^{DT} \quad (1)$$

where $\gamma_{m,t}^{DT} = \gamma_{m,t,T=1} - \gamma_{m,t,T=0}$. Hence, even if trends would have differed in period $t = 1, T = 1$ between groups in absence of any discrimination shock, we can still achieve identification of the shock’s effect if the *de-trended (DT)* evolution between groups would have been similar across periods in absence of the shock. Put another way, if the trend in de-trended time effects is constant between groups, the *DDD* parameter identifies the

shock's effect. Equation 1 also shows, identical to a typical *DD* assumption, that we do not need the de-trended levels in the time effect to be equal between groups, only the difference must be constant moving from period $t = 0$ to $t = 1$.

This gives a new formalization to what Bell, Blundell, and Van Reenen (1999) and later Blundell and Dias (2009) coin as the “Differential trend adjusted difference-in-differences” or more recently explored by Lee (2016) as “Generalized differences-in-differences.” It also shows that, as highlighted in Lee (2016), the *DDD* parameter identifies the same effect as a standard *DD* approach.

Given that we observe multiple periods before the shock our main identifying test will involve examining the evolution of outcomes in period $t = 0$ to give credence to the assumption that the difference in de-meanded trends would have stayed the same in absence of the shock. Outcomes are measured at the local employment agency (i) level over 20 weeks s - ten weeks before the shock and 10 weeks after - thus in examining the de-trended evolution of the two groups in the pre-shock period ($t = 0$), we will be able to provide evidence on whether or not the de-trended difference appears stable between the two groups in the weeks before the discrimination shock:

$$\gamma_{m=1,s,t=0}^{DT} - \gamma_{m=1,s,t=0}^{DT} = \gamma_{m=0,s,t=0}^{DT} - \gamma_{m=0,s,t=0}^{DT}, s \in (1, 10)$$

If “de-trended trends” are stable in the pre-shock period, this will support the assumption that they would have stayed the same in absence of the discrimination shock and we can therefore attribute a causal interpretation to any effect we see.

Using a regression equation to estimates the *DDD* parameter, our baseline specification is,

$$y_{imtT} = \beta_0 + \beta_1(m*t*T)_{imtT} + \beta_2(t*T)_{itT} + \beta_3(m*t)_{imt} + \beta_4(m*T)_{imT} + \beta_5t_t + \beta_6m_i + \beta_7T_T + e_{imtT} \quad (2)$$

Observations are each population m in agency i at period t in year T . The parameter of interest is β_1 . It captures the *DDD* parameter and we can think of it as the impact of the shock on minorities as compared to majorities on outcome y , controlling for potential differences in trends that may be present in periods t regardless of the shock.¹⁰

¹⁰We can also see the benefit of this setup in the regression equation. It is clear that if equation 2 is the true population model and we estimate the simple *DD*₁ through

$$y_{it} = \alpha_0 + \alpha_1(m*t)_{it} + \alpha_2t_t + \alpha_3m_i + \epsilon_{it}, \quad (3)$$

we estimate $\widehat{DD}_1 = \hat{\alpha}_1 + \beta_3$, where β_3 is the change in minority outcomes compared to majorities that happens in period $t = 1$ regardless of a discrimination shock.

4 Impacts

4.1 Potential matches

We begin by examining the impact of the Charlie Hebdo and Hypercacher attacks on the search behavior of jobseekers, counselors and employers. As discussed above, we consider these first order outcomes. To candidate to a job offer, for a counselor to encourage a jobseeker to apply to a job posting, or for an employer to seek out a jobseeker, are individual choices that we would like to assume to be independent from each other in a relatively short period of time. We will then consider actual employment outcomes as second order outcomes and will explore their impacts in the following section. We make this conceptual distinction because first order outcomes may have large cross-group impacts on second order outcomes. For example, if the discrimination shock causes a reduction in the search effort of minorities then this may have an effect on the job finding rate of majorities (or other minorities) due to increased effective tightness. And this increase in employment probability for majorities is not necessarily attributable to increased employer discrimination.

Table 2 presents results from estimating equation 2 using weighted least squares (WLS). Because the dependent variables are agency averages (the number of potential matches divided by the number of registered jobseekers), we weight the regression equation by $\sqrt{n_{imtT}}$ where n_{imtT} is the number of jobseekers contributing to the observation's average.¹¹ Panel A displays results for average potential matches over all contract types while panels B-E break down the *DDD* estimate by contract type. Throughout we cluster our standard errors at the agency level to account for correlation in agency outcomes overtime and for correlation between minority and majority outcomes within agencies (Bertrand, Duflo, and Mullainathan, 2004).¹²

For potential matches made on all contract types displayed in the first column of Panel A, we see a drop of roughly 0.002 matches per minority jobseeker, per week as compared to majorities, significant at the 5% level. This represents a reduction of about 3% off the mean minority potential match rate in the pre-shock period (displayed at the bottom of each

¹¹See the appendix for a discussion on the motivation for using WLS as opposed to OLS. Using OLS regressions on unweighted data and also poisson regression on the pure count data give consistent results as can be seen in Tables A.1 and A.2 in the appendix.

¹²Clustering at higher levels of aggregation such as the employment zone or even at the regional level, to account for correlation in outcomes across larger markets, provides very similar standard errors and does not change inference on the results.

panel). As can be seen in the following columns, this effect masks considerable heterogeneity across matching channels. Minority search intensity drops by 0.005 applications per minority jobseeker, equal to a 10.7% drop in the 10 weeks following the shock. In contrast, we see an augmentation in counselor matching intensity for minorities compared to majorities equal to .0029 more matches, or +13% off the pre-shock mean. For potential matches initiated by employers we see a much smaller and insignificant point estimate.¹³

When examining by contract type, the largest absolute and relative impacts for jobseeker and counselor search behavior are seen for permanent contracts, the most sought after contract. Minority jobseekers make 13% less ($\hat{\beta}_1 \approx -0.0033$) candidatures to jobs offering permanent contracts while counselors increase their effort for minorities by 16.5% ($\hat{\beta}_1 \approx 0.0016$) compared to majorities. We also see in panels C-D that the direction of the effect for jobseekers and counselors is consistent over the different types of contracts. This is not the case for employer search behavior. Column 4 in panels B-D show that the null overall effect of the shock for employer initiated matches on minorities masks differential effects by contract type. Employers reduce their search for minority candidates for the best type of contracts while increasing the number of minorities that they contact for their fixed-term and temp-work jobs. Finally, in panel E we see very small mean levels of potential matches made on seasonal contracts and estimated coefficients are close to zero for all channels.

Returning to panel A in Table 2 is also instructive in judging the utility of the *DDD* specification as opposed to a standard difference in differences approach. The estimates of β_3 are highly significant. This formally tests the change in the difference between minority and majority search behavior in $t = 1$ during the placebo year (we previously examined this visually in Figure 2). And it shows that regardless of the presence of a major discrimination shock, minorities and majorities would have had differential outcomes in $t = 1$ regardless. Thus if model 2 is the true population model, simple *DD* would have given an estimated effect of the shock on all potential matches that would have been significantly biased for both the jobseeker and counselor channels.

We now turn to examining these effects visually in order to gauge the validity of the underlying identification assumption. We focus on potential matches made on permanent contracts because this is where we see the largest effects and because it is the most prized contract. In addition, roughly 86% of salaried employees in France are on permanent contracts.¹⁴ Figure 3 plots bins of the de-trended weighted mean application rate to permanent

¹³We see that the mean of the dependent variable presented in these tables is 1/10th that of the averages presented in the descriptive statistics table. This is simply because observations are now at the weekly level compared to Table 1 where we aggregated data over the total pre-shock period. Thus it suffices to multiply the coefficient and standard error by 10 to obtain the *DDD* effect for the entire pre- and post-shock periods.

¹⁴In 2014, 86.4% of salaried employees were in permanent contracts, 9.7% in fixed-term contracts

contracts for each week in the study period. These weekly averages are overlaid with predictions from an OLS regression with a polynomial time trend of order three. For clarity, these graphs are the visual equivalent to examining the evolution of the *DDD* estimate: estimating a *DD* specification on these de-trended averages will give you the same point estimate as $\hat{\beta}_1$ in Panel B of Table 2. First, we can provide strong evidence for the identifying condition: the evolution of de-trended group averages are quite similar before the shock date (as denoted by the vertical line) for the three matching channels. Concerning candidate, or jobseeker, initiated matches in Figure 3a we do notice a more significant drop in the minority jobseeker application rate during the holiday season starting in week 8 in the pre-shock period, but if we look at the proportional drop between week 7 and 8, the difference is perhaps not as dramatic (0.5 for majorities versus 0.6 for minorities).¹⁵ That being said, we must acknowledge that the pre-shock difference is not perfectly constant going from weeks 7-8 for jobseeker initiated matches.

Nevertheless, we see large changes in the trends for potential matches for minorities as compared to majorities after the shock. The impact is most dramatic on potential matches made by jobseekers and counselors. As can be seen in first the graph, Figure 3a, starting in the third week following the shock we see a substantial drop in search effort. This difference with majorities continues for six weeks before the gap eventually opens up to its pre-shock level in the final week. The impact on counselor behavior is equally dramatic. We see a significant closing of the gap between the average number of potential matches made for minorities compared to majorities so that the difference in average match rates between the two groups becomes almost indistinguishable. Again, we see the difference start to shrink in about the second week following the shock. The impact on employer initiated potential matches is perhaps less visually dramatic, but still apparent: the gap between de-trended employer-initiated matches made after the shock period is significantly larger than the average pre-shock gap.¹⁶

Finally, we look at the identifying condition and evolution of the effect by plotting the estimates of the β_s coefficients from the specification,

$$y_{isT} = \delta_0 + \sum_{s=2}^{20} \beta_s (m * s * T)_{isT} + \sum_{s=2}^{20} b_s (s * T)_{isT} + \sum_{s=2}^{20} a_s s_{is} + d(T * m)_{iT} + T_T + m_i + e_{isT} \quad (4)$$

2.4% in temp work and 1.6% apprenticeships. See <https://www.insee.fr/fr/statistiques/fichier/version-html/1560271/ip1569.pdf> for a snapshot of the French labor market in 2014.

¹⁵This may also be linked to the fact that matches are simply biased towards zero during the holiday season thus we would expect a larger drop for the group with the higher level de-trended match rate.

¹⁶To better understand these de-trended graphs, we refer you to Figure A.2 in which we show these trends side by side by year. The trends for 2014-2015, the year of the attack ($T = 1$), are in the first column of graphs and the placebo year, 2013-2014 ($T = 0$), in the second.

The β_s 's capture the *DDD* parameter for week s in comparison to the the reference week, $s = 1$, thus allowing us to formally examine the effects we described above. Figure 4 plots these coefficients over the weeks of the observation period for search for permanent contracts. The first graph on the left shows the evolution of the impact for jobseeker and counselor initiated matches while the graph on the right shows this evolution for employer search (we do this because of the difference in scales). Coefficients are connected by lines with grey dashes indicating 95% confidence intervals. We see initially that in comparison to the reference week, there is significant between group volatility for jobseeker search effort. Minorities initially have slightly higher self-initiated matches compared to the majorities. We then see a switch in the difference during the holiday weeks that then goes to nil in the first two weeks of the year. After this point, we see a sharp drop in their job search intensity. This impact stays relatively constant for about a month before progressing back to majority levels of search effort in the very last week of observation.

For counselors we have the opposite story. They make relatively less matches to minority jobseekers as compared to majorities throughout the pre-shock period (in comparison to week $s = 1$).¹⁷ After the shock date the sign on the coefficients switches and they make relatively more matches for minorities after about the second week. For employers, we see relatively no differences in their search effort for minorities compared to majorities until the shock date, after which they start to drop their interview proposals to minority jobseekers. In comparing jobseeker and counselor initiated matches with those of employers, we note that employers appear to react much more quickly to the discrimination shock and that the shock is more transitory for them, while the effect on jobseekers and counselors takes a few weeks to “kick-in” and lasts longer. This may suggest that the way in which labor market actors internalize discrimination may vary considerably. A shock to discriminatory priors or tastes for employers may immediately cause them to act on them while it may take time for supply-side actors to realize that they are now facing a new level of discrimination in the market.¹⁸

¹⁷It must be noted that minority and majority matching rates are very similar in reference week $s = 1$. This can be seen in the raw data in Figure 3b. This is why we go from negative point estimates to positive estimates. Using week 10 as the reference group gives much more visually striking results.

¹⁸The weeks following the shock saw the spawning of numerous public debates in the media concerning the integration of French Muslims, such as the debate “Je suis Charlie” versus “Je ne suis pas Charlie”. It is therefore plausible that jobseeker and counselor perception of discrimination was made more and more salient in the weeks that followed the attack. See Todd (2015) for an interesting discussion of these topics.

4.2 Impacts on hires

We now turn to examining the shock's impact on minority hiring rates. These are second order outcomes because, as highlighted above, they cannot necessarily be attributed to increased discrimination faced by minorities in the labor market. As we've seen in the previous section, the shock led to large changes in the search behavior of labor market actors. Thus it is impossible to disentangle the effect of, say, a minority who does not get an interview because the employer is now more biased, versus the effect of a change in the equilibrium job-finding rate for majorities due to lowered minority search effort. Or, for example, that a minority becomes more selective about where they apply, hoping to avoid discriminatory employers. Nevertheless, we can see if the shock led to differences in the types of contracts minorities are hired in, and if they have more or less job creation on aggregate compared to majorities.

Table 3 presents results of the shock's impact on contract flows for minorities as compared to majorities. The table shows impacts over all contract flows, column 1, and then breaks down the impact by contract type in columns 2-4. Perhaps surprisingly, we detect a positive impact on total flows. Minorities are predicted to sign about 1.3% more contracts than majorities, overall. We can immediately see that this impact is entirely driven by an effect on fixed-term contracts for minorities (column 3). We see no impact on permanent contract flows, nor for temp work contracts. If we turn immediately to Table 4 we see that this effect comprises a substantial increase in very short-term contracts. We see positive *DDD* estimates significant at the 1% level for 1 day and 1 week contracts. In contrast we see a negative estimate for fixed-term contracts longer than six months, significant at 10% (column 6).

It is insufficient to examine contract flows to measure potential impacts on employment. Therefore we now turn to Table 5 in which we aggregate the total number of workdays created within contracts signed by our jobseekers, and thus explore whether the shock had an impact on real job creation. To calculate this dependent variable we simply use all contract start and end dates available in the contract flows data. Results in this table are telling. We measure no impact on real workday creation for minorities compared to majorities after the shock. The reason for this is evident in Table 4: though we see strong impacts on short-term contracts this is counter-balanced by the negative point estimate on contracts longer than six months.¹⁹

We now discuss the interpretation of the results on these second order outcomes. We begin by illustrating the *DDD* effect on fixed-term contracts visually in Figure 5. We immediately

¹⁹Unfortunately, start and end dates are not available for temp hire declarations.

see that the positive effect on fixed-term contract flows is being driven by a sharp drop in the majority group’s de-trended average at about the the fifth week after the shock. The minority de-trended average remains stable over the entire post-shock period. Hence, it does not appear that the shock has caused minorities to flock to short-term and poor quality contracts; this could be the case if, for example, a jobseeker’s demand for work is inelastic to discrimination and the only work they can get is in short-term jobs because employers now discriminate more. Instead we think it is important to think about the general equilibrium effects potentially caused by the drop in minority search effort that we have discussed above. The fact that minorities reduce their search effort for permanent contracts due to an increase in either real, or perceived, discrimination, has an impact on the job-finding probability of majorities. Hence another interpretation of Figure 5 could be that because majorities have an increased probability of being hired in a permanent contract due to lower “real tightness” stemming from lower minority search effort, they are now less dependent on short-term fixed contracts. The point of this example is that these dynamics can occur regardless of a change in the actual level of bias of employers. Therefore, it would be imprudent to attribute the hiring effects we’ve seen to the impact of the attack’s effect on employer beliefs. Hence we remain agnostic about the shock’s effect on hiring dynamics and only conclude that, overall, we detect no significant impact on overall employment creation for minorities.

4.3 Robustness

Placebo names

The first way we test the robustness of these results is to substitute another name to connote minority status. We simply reproduce our main results using a name that should not be correlated with discriminatory tastes. Table 6 displays results for this test using names that are etymologically British in Panel B and etymologically Southern European in Panel C for our search behavior variables. Panel A reproduces our estimates for the impact on search behavior for permanent contracts from our main results in Panel B of Table 2 for comparison. We find strong supporting evidence for our hypothesis that the shock disproportionately affects minorities of Maghreb/Mashriq decent. When substituting British or Southern European names for minority status we find much smaller and inconsistently signed coefficients for the *DDD* estimate that are almost all insignificant. We do detect a positive and significant coefficient (at the 5% level) for counselor initiated matches, but the point estimates are between 3 and 5 times as small as our baseline definition of minority status, with the null hypothesis of the equality in coefficients easily rejected at the 1% level for both placebo names.

Table 7 tells a similar story for contract flows. Using British and southern European

names as the minority definition gives insignificant and much smaller coefficients across the board.

Parametric controls

We now turn to adding fixed covariates, interacted with our time and group indicators, to our baseline specification. Though we control for group level differences that our fixed over time in model 2 it is informative to explore if our *DDD* estimates are significantly influenced when we add pertinent controls interacted with time and minority status. Indeed, we want to be able to exclude the possibility that the impacts we have thus far demonstrated are the result of underlying determinants of the agencies' pool of jobseekers that might be correlated with minority status *and* changes over time.

Table 8 presents results for estimates of our baseline specification while progressively adding variables from Table 1 interacted with the period t , year T and $t * T$ indicators as well their pairwise interactions with minority status, $m*t$, $m*T$ and $m*t*T$. As a reminder, these variables are average group proportions within agencies during the $t = 0$ (pre-shock) period. They capture the nationality, the types of jobseeker, demographic characteristics and sector specialization.²⁰ We center these control variables at the group mean level in order to interpret our *DDD* parameter in the same way as in our baseline specification. In the first column we reproduce our results from panel B of Table 2, matches made to permanent contracts, for ease of comparison. In the second column we start to add controls, and for the sake of brevity, across the board as we continue to add these controls we see very small changes in the point estimates.

This is quite heartening in our effort to interpret our results as causal. In effect, we are parametrically matching minority and majority agency populations on important characteristics that might be differentially correlated with changes in time periods t and T . The stability of our results suggests that our results are not being driven by underlying differences in group characteristics that interact with time and minority status. This analysis provides supporting evidence that the majority group is a suitable non parametric control group when using *DDD* in our context.

Compositional changes

Though we exploit panel data at the agency level, jobseekers flow in and out of these agencies. Thus our sample is essentially repeated cross sections of registered jobseekers within the agency and we would therefore like to test if the effects we find on search effort are being driven by changes in the underlying composition of registered jobseekers. This is because the composition itself may be affected by the shock. For example, if the shock leads to an

²⁰Because the proportions add up to one for the categories of jobseeker type, nationality and profession, we exclude one variable from each category as our reference to avoid multicollinearity between regressors.

increase in low qualified minority jobseekers registering with the PES and counselors initiate more matches for this demographic naturally, i.e. in absence of the shock, then the increase in counselor initiated matches could be mechanically driven by a change in the composition of their jobseeker portfolio (that differs over years). Table A.3 in the appendix presents impacts on the composition using our *DDD* specification on the available characteristics of jobseekers. We see that out of a total of 29 regressions, six are significant at at least the 5% level. These significant effects are found for the average number of highly qualified jobseekers, young jobseekers, those with Maghreb nationality and jobseekers searching for jobs in construction, the trades, or in the theater and film industry. Thus it appears that the shock may also be correlated with compositional changes in the types of minority jobseekers registered with the PES as compared to majorities. But if we carefully read the table we see that the estimated impact of these compositional changes are relatively small compared to the impacts we see on job search intensity. For instance, on the three demographic characteristics for which we have significant effects, the coefficients are very small and when taken over the minority pre-period average they represent changes of between 0.2% - 0.3%. This is an order of magnitude smaller than the impacts seen on job search intensity. For the remaining significant estimates in Table A.3 the proportional changes are larger, but only because the baseline levels are so small.

Even if these compositional effects are relatively small, we would like to interpret our job search results holding them constant. Indeed, though these could be impacts in and of themselves, the premise of this paper is to interpret the impacts we find as originating from a change in the decision process of labor market actors. Unfortunately, we cannot simply include them (along with their interactions with minority and time dummies) as controls in the regression because we have seen that they are correlated with the shock and therefore are endogenous, thus potentially making our *DDD* estimate inconsistent (see Frölich (2008) for a discussion on the use of endogenous controls). We thus adopt an instrumental variables approach. Since we have 10 weeks of pre-shock data we can take the average over these 10 weeks and use it to instrument the endogenous control variables. For example, the average number of highly qualified jobseekers in the $t = 0$ pre-shock period is highly correlated with its average in all weeks, but is orthogonal to the shock. Table 9 displays results of the shock on search intensities using our baseline specification while also including these potentially endogenous compositional controls. Each of the controls is centered and interacted with the period, year and minority indicators as well as their pairwise interactions. These endogenous controls are then instrumented with their analogous counterparts using the pre-shock means.

We see that our results are very consistent with our baseline specification with no substantive changes to the *DDD* estimates. It thus appears that the changes in labor market

actor search effort is not likely being driven by the underlying, albeit small, changes in the pool of registered jobseekers. Of course there could be many other compositional changes that we do not observe in the data and are thus unable to control for, but given that our results are so stable when we include these endogenous controls we would like to infer that this is not a large source of potential bias in our impact estimates, nor in their interpretation.

Our final formal robustness check involves assessing the validity of our $T = 0$ placebo year compared to a previous year ($T = -1$) and we explore this in depth in the appendix. Lee (2016) suggests that differencing again by another anterior year (say, $T = -1$) to obtain a quadruple difference (QD). If the estimated QD parameter is similar to the DDD estimate, he argues that it provides credibility to the DDD identification strategy. Yet Bell, Blundell, and Van Reenen (1999) and Blundell and Dias (2009) argue that the previous year is most likely to be the best counterfactual because macro economic trends should be most comparable. Using the previous year is, of course, the strategy that we have adapted here.²¹ And as we have seen, an advantage of our data structure compared to these studies is that we can directly examine the $t = 0$ de-trended trends over the 10 pre-shock weeks with which we were able to present evidence to back up our identifying condition. Nevertheless we would like to develop a formal test to decide which placebo year is the most pertinent to use rather than simply arbitrarily choosing a year to use as the placebo or adding additional differences to equation 2 which may actually add bias to results if macro economic trends differ substantially by group as we go back in time. In the appendix, we develop a simple test that exploits the $t = 0$ data in $T = 0$ and $T = -1$ to explore their relative comparability with the actual shock year, $T = 1$. We find that the $T = 0$ year indeed appears to be the most comparable with our shock year. We then elaborate a very simple scalar weighting system to create a synthetic placebo year where outcomes are a weighted average of the two previous years given the comparability of outcomes in $t = 0$.²² This allows us to include information from previous years in the estimation, but in a data driven, non-arbitrary fashion. We find that using this synthetic placebo year gives comparable results on outcomes for which we have data in the two previous years to the shock.

Finally, if the impacts that we are measuring are truly causal and due to a shock that increases discrimination against our minority population, we might expect effects to be correlated with an existing measure of discrimination or with the minority status of the counselor. This is what we find and we detail this analysis in the following sections.

²¹This is also out of necessity as the potential matches data only started to be collected by the PES in 2013. We thus use the hiring data in this exercise.

²²This is, of course, inspired by Abadie et al. (2010), but instead of the synthetic control being a function of weighted covariates that minimize the distance of the treated entity and other panel subjects, the weights here are constructed within entity over time using pre-treatment outcomes.

5 Heterogeneity in impacts

We now explore these results through a dimension of latent discrimination. Unfortunately we are unable to have individual level measures of discriminatory attitudes, but we try to proxy for this at the agency level using the vote percentage for the Front National (FN) in the 2012 presidential election. The Front National is France’s major far-right political party and it has a long and robust relationship to discriminatory tastes against Muslims (start with Mayer and Perrineau (1996) for a history of the political movement). We therefore link our agency observations using FN vote data at the commune, or municipality level of which there are over 36,000 in France. As described above, each municipality is attached to a local employment agency, thus we create agency level vote shares for the FN as a proxy for latent discrimination. We then split the sample between agencies with above the median level FN vote shares and agencies with below the median vote shares for significant parts of the following analysis.

5.1 Theoretical motivation

Before we present the empirical results on heterogeneous effects, we make use of some theory to motivate our analysis and make predictions. We focus on the behavior of jobseekers and counselors. The theory literature on the effect of discrimination on employer behavior is already expansive (start with Lang and Lehmann (2012) for a summary). In sum, we endogenize jobseeker and counselor search effort with respect to the discrimination in the market.

Jobseekers

We start with a standard value function of unemployment with endogenous search effort (Pissarides, 2000) where we assume that that minority jobseekers only take into account the effects of discrimination on their job-finding rate and that their decisions are independent of what counselors do. The utility of unemployment is formalized as follows.

$$rV_{u,i} = z - c(e_i) + e_i f(\theta)(1 - d^\alpha)(V_{e,i} - V_{u,i}) \quad (5)$$

Where rV_u is the discounted value of unemployment for minorities, z any benefits that accrue from that state, $c(\cdot)$ is a cost of effort function that is assumed to be convex. $f(\cdot)$ the job-finding rate which is an increasing function of tightness θ . Discrimination is captured by $d^\alpha \in (0, 1)$ with higher values of d corresponding to more discrimination, but its effect is diminishing with $\alpha \in (0, 1)$. V_e is the value of moving into the employment state. Thus, the probability of finding a job is related to the effort put forth by minority i , the tightness and

matching technology and discrimination in the labor market. Minority workers optimizes utility with respect to effort e :

$$c'(e_i) = f(\theta)(1 - d^\alpha)(V_{e,i} - V_{u,i}) \quad (6)$$

Assuming $V_{e,i} > V_{u,i}$, we can see that since $c(e)$ is convex, $c'(e_i)$ is increasing in the matching rate. Take the derivative of this function w.r.t. d we have,

$$\frac{\partial c'(e_i)}{\partial d} = -\alpha f(\theta) d^{\alpha-1} (V_{e,i} - V_{u,i}) \quad (7)$$

and we clearly see that search effort is decreasing in d . This reflects recent work by Skandalis and Philippe (2016) who find that information about lower perceived job finding probabilities reduces search effort. Now assume that there are two levels of latent discrimination, low and high, with $d_L < d_H$. We now compare the ratio of this marginal effect on search intensity at high and low levels of discrimination and we find,

$$\frac{\partial c'(e_i)}{\partial d_H} / \frac{\partial c'(e_i)}{\partial d_L} = \left(\frac{d_H}{d_L} \right)^{\alpha-1} < 1 \quad (8)$$

Thus, due to the concavity of d^α , the marginal effect on jobseeker search effort of an increase in discrimination should actually be larger in low discrimination areas.

Counselors

Assume that counselor effort e_j is considered an input in the efficiency parameter of a constant returns to scale matching function, $H(\cdot)$ with arguments v and u , the volume of vacancies and jobseekers, respectively:

$$M = e_j(\mathbf{x}_i, d_i) * H(v, u) \quad (9)$$

where $e_j(x_i, d_i) \equiv p_j(\mathbf{x}_i, d_i) - g_j(d_i)$ in which \mathbf{x}_i is a vector of jobseeker i characteristics and d_i the discrimination the jobseeker faces on the labor market. We assume that $p_j(\mathbf{x}_i, d_i) > 0$ with derivatives w.r.t d_i taking form $p'_j(\mathbf{x}_i, d_i) > 0$ and $p''_j(\mathbf{x}_i, d_i) < 0$. Hence $p_j(\cdot) > 0$ is an increasing concave function that represents the *perception* that counselor j has of the job finding difficulty faced by jobseeker i given a vector of his or her characteristics and the discrimination d_i that the jobseeker faces on the labor market. For simplicity, we assume that $g_j(d_i)$ is a linear cost function: $g'_j(d_i) = 1$. In sum, effort put forth by the counselor in the matching process is thus a tension between augmenting effort as a function of the discrimination that the jobseeker faces *and* the cost of that effort due to the counselor's own

bias. A marginal change in d_i gives us,

$$\frac{\partial M}{\partial d} = (p'_j(\mathbf{x}_i, d_i) - 1)H(v, u) \quad (10)$$

Hence $\frac{\partial M}{\partial d} > 0$, if and only if $p'_j(\mathbf{x}_i, d_i) > 1$. With the diminishing marginal effect of $p_j(\cdot)$ at low levels of discrimination, the slope of $p'_j(\mathbf{x}_i, d_i)$ is positive and relatively large. Thus overall counselor initiated matches may increase, but as we move up the support of d , $p'_j(\mathbf{x}_i, d_i)$ will flatten out and the cost of their own discrimination will lead to a potential null or negative effect on the match rate.

Given very light structure on how discrimination enters the search and matching effort of jobseekers and job counselors, we should expect the marginal effect of a shock to discrimination on jobseeker effort to be higher in areas where latent discrimination is *lower*. And we should see a “compensatory” effect on the behalf of counselors because these counselors take into account the worse labor market prospects that discrimination causes for minorities. Furthermore, this compensatory effect should be stronger in areas where there is relatively lower pre-shock discrimination.

5.2 Empirical results

Warning: This section uses and discusses racially charged words that may be hurtful to some readers.

We begin by testing whether the Front National vote share in the first round of the 2012 French presidential election is a predictor of existing, or latent, levels of bias towards our minority group. We follow the work of Stephens-Davidowitz (2014) and use Google Trend data to examine whether people in areas with high vote shares for the FN search at a higher rate for terms that indicate the presence of discrimination against our minority population and whether this discrimination denotes racial animus. We thus intend to provide evidence to support the hypothesis that the FN vote share is positively correlated with discrimination that existed before the attack. We use the terms “islamophobie” and “bougnoule”, the first term proxying for the existence of prejudice in a vague and larger sense and the second, the nature of the prejudice. Bougnoule is the most common racial slur used in France for people of Maghreb origin.^{23,24} On average, this term is searched for once for every 11 searches for

²³It has the highest comparative search rate in the year preceding the shock compared to other “popular” racial slurs against this population, such as “bicot,” “boucaque” and “meteque.” Among these terms it also appears to be the most generally used across French regions as these other terms exhibit high regional correlation.

²⁴Wiktionary notes the English language equivalent of “bougnoule” would be “Sand Nigger”, “Camel Jockey” and “Camel Fucker”. See <https://fr.wiktionary.org/wiki/bougnoule>.

“find a job” and once for every 14 searches for “bake a cake” (faire un gateau) in the year preceding the shock.²⁵

Though we will use very local measures of the FN vote share (municipality level that are attached to our employment agencies) in our impact analysis, Google Trend data are only available at the regional level, of which there are 22 in metropolitan France. Hence, we aggregate the vote shares to the regional level to study correlations. The trend score per region is calculated as the number of times a term was searched for over total searches within the region. These proportions are then normalized to the region with the highest proportion. Thus the regional scores have meaning when compared against one another.²⁶ We regress these scores on the regionally aggregated FN vote share as well as the proportion of minority jobseekers in the region,

$$score_r = \beta_0 + \beta_1(FN\ Vote\%)_r + \beta_2(Prop.\ Minority)_r + u_r \quad (11)$$

We present these results visually by predicting \widehat{score}_r in region r and plotting them on a graph over the scatter plot of the raw data in Figure 6 in which we also display the p-value for the FN vote share coefficient. We control for the proportion of minority jobseekers so as to proxy for the underlying proportion of the minority population in the region. As best we can, we want to interpret the score’s correlation with the FN vote share holding the number of potential searches by minorities constant. In examining the graphs we see strong positive correlation with the search volume for these terms and the FN vote share in the year preceding the attacks. And even though we only have 22 regions, the p-values using robust standard errors for islamophobie and the racial slur are 0.11 and 0.06, respectively. Because we hold the proportion of minorities constant in these regressions, it is not simply that minorities are searching more for these terms in high FN areas and that this is driving the search volume. It appears that, on average, high FN regions appear to be more associated with both the presence of discrimination and that this discrimination is, at-least in part, taste-based.

We now turn to examining impacts using the agency level of the FN vote share. Table 10 displays results. For ease of interpretation, results are taken from estimating equation 2 separately for agencies that have a vote share below the median (Low FN) and for agencies above the median (High FN). Below the point estimates we also display the p-value for a test in the equality of the DDD estimates between the two sub-samples.

We see in column 2 of Table 10 that jobseekers in both low and high FN areas reduce their search effort after the discrimination shock, but that the reduction is much stronger in

²⁵See the introduction for an explanation on the interpretation of the relative search scores.

²⁶This subsumes that there we not large, differential changes in total search volume by region.

low FN areas with this difference being statistically significant. The heterogeneous impact on counselor behavior is even more striking. We see that the increase in potential matches made to minorities after the shock is completely centered on agencies in low FN areas. We see no increase in counselor matching effort for minorities in areas that exhibit relatively larger levels of existing discrimination.²⁷ For employers we see a point estimate that is about twice as large in high FN areas, but we cannot reject the null hypothesis of equality in impacts in the different areas. Finally, when aggregating across all channels for potential matches we see in column 1 that the reduction in total matches is primarily driven by the reduction in jobseeker matches and that only in low FN areas is this reduction in jobseeker effort offset by the increase in counselor effort. Consistent with these search results we see the effect on contract flows is also entirely centered in high FN areas as can be seen in Table 11.

These results provide evidence to support the hypothesis that there may be diminishing marginal effects of discrimination on the search intensity of jobseekers and counselors. To look closer at whether our modeling assumptions are indeed in-line with the data we can plot the marginal effect of the shock over the entire support of the FN vote share. We do this by interacting the terms in our baseline equation 2 with the continuous measure of the agency-level FN vote share and its square. We then take the derivative of the $(m * t * T)$ term evaluated at increasing levels of the FN vote share. We refer to Figure 7 to examine these results. We see strong evidence supporting our modeling assumptions: the effect of the discrimination shock decreases in magnitude for both jobseeker (subfigure 7a) and counselor (subfigure 7b) initiated matches as we move to the right-hand side of the FN vote share distribution. For jobseekers, the shock's impact is apparent up until roughly the 80th percentile of the vote share and then becomes insignificant. The shock's impact on matches made by counselors also show strong decreasing marginal returns, but in contrast to jobseeker matches we see that the shock's effect is much more correlated with our measure of latent discrimination. The effect dramatically decreases as we move towards the 50th percentile after which we see no effect. These results suggest that levels of existing discrimination play a major role in determining the magnitude of impacts that a shock to discrimination might entail. Put another way, if minorities already face high levels of bias then their perception of the returns to search effort are not as affected as minorities who face relatively low levels of initial bias. Counselors, on the other hand, engage in a trade off between helping jobseekers who now have poorer expected employment outcomes and the cost they incur due to their own discrimination. These results point to a scenario where only counselors in areas with relatively lower levels of latent discrimination internalize the effects of discrimination on their

²⁷Given our theoretical interpretation this suggest that the $p_j(\mathbf{x}_i, d_i)$ curve quickly flattens out above the median level of d .

minority jobseekers as this effect quickly diminishes over the support.²⁸

This striking heterogeneity in impacts across a measure of discrimination also reinforces the interpretation of the main results of this paper as causal. In essence we are contrasting the data by a fourth dimension (after the group, period and placebo year) thus this analysis acts as another robustness test predicated on the idea that a shock that increases discrimination is highly dependent on the level of existing discrimination. Because the basic motivation of this paper is to study the effect of discrimination on the job search of labor market actors, if we had found no heterogeneous effects when comparing across the discrimination proxy it would be a cause for concern. Indeed, if that were the case we might have serious doubts that the effects we are picking up are actually linked to discrimination and not to other group-specific differences in the time effect across years that we can't control for. But the fact that we see such striking differences between high and low discriminating areas, and that the FN vote shows marked correlation with discriminatory attitudes, provides credibility that we are identifying the effect of the shock on the job search behavior of minorities.

Impact of shock on search trends

We now turn to examining the shock's impact on search trends for these two terms and whether the effect size is disproportionately correlated with the FN vote. We do this to anchor our understanding about shifts in attitudes that might have been caused by the attacks and to better understand the heterogeneity analysis.

The graphs in Figure 8 present evidence on this. The top row of graphs display the search volume for the terms "islamophobie" and "bougnoule" in the weeks before and after the shock. As noted in the introduction, this top row of graphs do not display total search volume, only the volume relative to the highest point on the chart during the window of observation.

We see that the reference point for search volume is dictated by the shock. There is a massive increase in the relative search volume for both these terms following the shock that quickly dissipates for the racial slur and slightly less quickly for islamophobie. This provides strong evidence that the shock exogenously triggered interest, not only in the potential existence of discrimination towards minorities, but also increased discriminatory tastes.

In the second row of Figures 8 we try to determine if the large increase in search rate for these terms is correlated with the FN vote share. The p-values using robust standard errors

²⁸Assuming another functional form for the cost function of counselors $g_j(d_i)$ would not contradict these empirical results but would give another, perhaps more nuanced, interpretation of what we're seeing. For example we could assume it is convex. Then the arbitrage that the counselor faces would depend on the comparative rates of change with respect to d_i between $g_j(d_i)$ and $p_j(\mathbf{x}_i, d_i)$. It could be that the steep decreasing marginal effect of the shock over the discrimination measure is the result of the derivative of the cost function dominating the derivative of the perception function.

for islamophobie and the racial slur are 0.053 and 0.037, respectively, reflecting the results using data in the year before the shock.

We repeat this exercise using trend scores for terms that might be considered the antithesis of the negative terms. The results are presented in Figure 9. We look at the total French search volume and relative search volume by region for “solidarite” and “fraternite” around the shock date. Again, we show the score’s correlation with the FN vote share below. As with the negative terms, we see big spikes in the relative search volume for these terms starting at the week of the attacks. But in contrast, the FN vote share by region is negatively correlated with these search terms.²⁹ It appears that our proxy for existing discrimination not only proxies for discriminatory tastes, but also for less interest in terms that connote preferences for social cohesion around the date of the shock.

Finally, we would like to test how the shock affected search trends over time in these regions. We do this using the following regression equation,

$$\begin{aligned} score_{rT} = & \gamma_0 + \gamma_1 Shock_T + \gamma_2 (High\ FN)_r + \gamma_3 (High\ FN * Shock)_{rT} \\ & + \gamma_4 (Prop.\ Minority)_r + \gamma_5 (Prop.\ Minority * Year)_{rT} + u_{rT} \end{aligned} \quad (12)$$

where the *score* is the Google trend search score for region r in year T . *High FN* indicates if the region’s FN vote share is above the median and shock is equal to one for regional scores measured during the shock year and zero otherwise. *Prop. Minority* is defined as above and we center the interaction at the mean level in order to interpret γ_1 and γ_3 as effects at the mean minority population level. u is a normally distributed, mean-zero error term. Region scores for $Year = 0$ are taken from Google trends over the year preceding the attack while scores for $Year = 1$ are taken for the 10 weeks following the shock. Table 12 show results from this specification using OLS with clustered standard errors at the regional level. Each column is a separate regression with the regional score for the search term noted in the column title.³⁰ Examining columns 1 and 2, we see large increases in average search scores (compared to the reference region) associated with the shock. On average regional scores jump 6.9 and 13.3 points for islamophobie and bougnoule in low FN regions ($\hat{\gamma}_1$), respectively. And even with the very small sample size, this effect is significant at the 1% level for the ethnic slur. This reflects the large jump seen in the search volume for these words in the weeks following the shocks (Figure 8). Importantly, we do not detect a qualitatively large

²⁹For fraternite, the negative correlation appears to be largely driven by one region.

³⁰As a recap, the dependent variable is a score going from 0-100. A score of 100 is automatically assigned to the region with the highest search rate for the term calculated over all searches within the region. All other regions are then given their score normalized to this highest score. Hence a score of 50 in another region means that the search rate for the term is half of what it is in the region with the highest search rate.

differential mean across time for high FN regions ($\hat{\gamma}_1 + \hat{\gamma}_3$) as compared to low FN regions, suggesting that the discrimination shock affected perception and animus throughout France, on average. Because we are able to difference the search scores over time, the correlation we see in Figure 8 appears to be driven by the underlying discrimination that existed before the shock. Though this data is highly aggregated, this suggests that the functional form of our simple job search functions for jobseekers and counselors create the large differential impacts we see across low and high FN areas, as opposed to large level changes in discriminatory tastes.

Looking at columns 3 and 4 of Table 12 we see a similar story. We see large increases in the search for these positive terms in the shock period, but the growth in search rate for these terms is present over both low and high FN regions. Hence the correlation we see in Figure 9 is primarily a function of the underlying characteristics of low and high FN areas as we do not see significant differential changes across periods between low and high FN areas.

5.3 Compensatory effects

To further test the theory of compensatory effects of job counselors, we test two final hypotheses. (1) Do minority counselors react similarly to majority counselors in response to the shock in terms of the matches they make towards their minority jobseekers? (2) Do counselors who specialize in assisting the most marginalized jobseekers react similarly to “regular” job counselors? We can test these hypotheses because we know two key things about counselors: their minority status and their specialization. Counselors that specialize in assisting jobseekers the most at risk for staying unemployed are called “Conseiller de parcours renforcé”, and we will refer to them simply as “intensive counselors.” They typically have a much smaller case work load (volume of jobseekers they are responsible for) and are expected to provide much more personalized assistance to their clients.³¹

In order to test these hypotheses we create a data set where observations are now at the counselor level over the 20 weeks of the observation window. For jobseekers that we can match to counselors, we aggregate the number of counselor initiated matches by counselor for each type of jobseeker and take the average. Hence, we create the analog of our main data set, but instead of observations at the agency-group level, we follow counselors. We then run *DDD* specifications where instead of *Minority* indicating the minority status of

³¹The guidelines for the number of jobseekers assigned to intensive counselors at any one time as opposed to regular counselors is as follows. 70 jobseekers maximum versus 100-350 for regular counselors. Regular counselors are defined in the PES jargon as either Conseiller de parcours “Guidé” or “Suivi.”

the population in the agency, it indicates the minority status of the counselor, m^c .

$$y_{jtT} = \lambda_0 + \lambda_1(m^c * t * T)_{jtT} + \lambda_2(t * T)_{jtT} + \lambda_3(m^c * t)_{jt} + \lambda_4(m^c * T)_{jT} + \lambda_5 t_t + \lambda_6 m_j^c + \lambda_7 T_T + e_{jtT} \quad (13)$$

Our dependent variables will be the average match rate they make for majority jobseekers, minority jobseekers and, more importantly, the difference between the two.³² This difference in average counselor initiated matches between minority and majority jobseekers is the analog of our main *DDD* estimate.³³

We begin this analysis by examining Table 13 in which we present estimates for the coefficients of interest: λ_1 and λ_2 . They give the effect of the shock for majority counselors (λ_2) and the differential effect for minority counselors compared to majorities (λ_1). Consistent with our main results, we see that the rise in the overall matching rate is driven by a positive differential increase for minority jobseekers. What is striking is that the point estimates in column 4 indicates that - though we see a compensatory effect for majority counselors significant at 10% - minority counselors increase their matching rate by about four and half times as much.³⁴ It is important to note that the effect seen on aggregate in the previous sections is not being driven by minority counselors as they make up only 9.7% of the sample. Yet, the difference is striking and appears to be strong evidence for a rejection of the hypothesis that majority and minority counselors behaved similarly towards their minority jobseekers following the shock. In light of our simple model where counselor effort increases in the perception of discrimination that the jobseeker faces, these results tell the story that minority counselor effort is more sensitive to the shock that increases discrimination against their own type. This story would be consistent with both Dee (2005) and Behncke et al. (2010) who find that students and jobseekers perform better when paired with teachers and counselors that are similar to them. We argue that this evidence is suggestive of a scenario where minority counselors appear to better internalize a potentially worse labor market environment for their minority jobseekers. Yet, this could of course be explained by the cost that they pay for exerting this effort. If minority counselors pay a smaller cost as a function of d_i this could explain the results just as well. We now present evidence that suggests that the a counselor's job type may also play a significant role in how these intermediaries filter

³²Regression equations are now weighted by $\sqrt{n_{jtT}^m}$ where n_{jtT}^m is the number of minority jobseekers in the counselor's portfolio. Unweighted regressions give very similar results.

³³The analog is actually the combined effect of $\lambda_1 + \lambda_2$ and λ_2 weighted by the relative proportion of majority and minority counselors. Even if weighted by the proportion, we will not have the exact same point estimate because we are not able to link all jobseekers to their counselors, only a subset. Furthermore, we are constrained to looking at all potential counselor initiated matches for all contract types because in this data we cannot link the contract type to the match. But we will see that the point estimates are very comparable to Panel A in Table 2.

³⁴We simply take the ratio $(\hat{\lambda}_1 + \hat{\lambda}_2)/\hat{\lambda}_2$.

the impact of a discrimination.

As mentioned above, we can distinguish between two types of counselors: intensive and normal. These intensive counselors regularly see their jobseekers face-to-face and specialize in getting particularly marginalized individuals back to work. We thus propose that counselors who engage in intensive follow-up of their jobseekers may be more or less reactive to a discrimination shock. This idea is related to Prendergast (2007) who explores the motivation of bureaucrats and the relative bias³⁵ they display either towards a principle or towards their clients. He shows that certain types of bureaucracies such as employment agencies (principle) will tend to hire agents (counselors) that are the most altruistic towards clients (jobseekers). This is because client and principle preferences are aligned, i.e. both want the jobseeker to find work. Thus agents who care most about the client exert the most effort and will be hired into intensive roles accordingly. We take this insight to the data and test whether counselors who should provide more or less direct advocacy for their jobseekers are more or less affected by the shock. Table 14 displays results where our dependent variable is the difference in average match rates between minority and majority jobseekers in the counselor's portfolio as shown in column 4 of Table 13. The first column uses all counselor observations while columns 2 and 3 look at effects isolating the sub-sample of intensive support counselors and normal counselors, respectively. Panels further split the sample by our proxy for discrimination. Panel A presents very interesting results. Intensive counselors appear to change their behavior the most. We see strong effort effects by majority counselors. They increase their match rate for minority jobseekers substantially (compared to majority jobseekers). In addition we see that the effect size for minority intensive counselors is about double the size, even though the standard error is very large due to the much smaller sample size and the fact that minority counselors make up only 10% of the counselor sample. When looking at the decomposition by our latent discrimination variable we see that this effect is primarily driven by *majority* intensive counselors in low FN areas. This contrasts with majority normal counselor behavior for which we see no significant change in their match rate, regardless of the underlying latent discrimination in their area. For normal counselors, only minority counselors react to the shock, and this reaction is only apparent in low FN areas.³⁶

In summary, we have provided evidence that minority counselors appear, on average, to provide a higher level of compensatory behavior, suggesting that the shock may either be

³⁵In this case we talk about bias not as discrimination.

³⁶Interestingly, we also note a large and negative point estimate for minority intensive counselors in high FN areas. Though statistically insignificant, this may speak to the types of minority counselors that are hired in these agencies, a theme related to the alignment of principle and client preferences. Or, regardless of their own preferences, how the cost of their effort may be too high given the environment they work in.

more salient to them and/or they pay a lower cost for their effort towards minority jobseekers. Perhaps, more intriguing, we have also uncovered strong positive effects for majority counselors, but these are isolated to individuals whose job is based on the close support and understanding of the difficulties that their jobseekers face in getting back to work. And, as highlighted in the main results, these effects are only present in low discrimination areas.

5.4 Discussion

We begin by discussing a possible explanation for why we see no significant effects of the shock on overall workday creation for minorities as shown in Table 5. The dramatic drop in minority search effort is only partially compensated in volume by the increase in counselor initiated matches over all contracts, and as a reminder, we see no significant effect on employer initiated potential matches when looking over all contracts as shown in Panel A in Table 2. One explanation, and perhaps the most straightforward, is that counselor initiated matches have a higher probability of resulting in a hire. The PES tries to track the outcome of potential matches, namely, if they resulted in a hire. The mean success rates of potential matches are very different depending on if they emanate from a counselor or a jobseeker: for counselors the hire rate for their matches is 6.2% while for jobseekers it is only 0.5%. The large caveat here is that it is the counselors themselves that check whether the potential matches of their jobseekers resulted in a hire, thus comparing these two means is subject to large measurement error bias.³⁷ Nevertheless, the difference is very large. Using the estimates from Table 2 the real mean success rates needed to produce a null effect on total hires in permanent contracts would actually have to be 2:1. Indeed, this ratio seems much more plausible.

Employer search

We have thus far focused primarily on jobseeker and counselor behavior and only briefly described the effects on employer search behavior. We now turn to the key points we might be able to highlight about the results found on their search behavior and their connection to the literature. As introduced in the beginning of this paper, there is debate about the best way to interpret audit studies (Heckman (1998) and Riach and Rich (2002)) in the presence of directed search.³⁸ Theoretically this has been addressed in the literature by Lang et al. (2005), but empirical evidence is lacking on real firm search behavior. In this paper we actually observe the search behavior of employers so we are able to circumvent this

³⁷Indeed this is a reason why we do not use this data in the analysis. In addition the measurement error could be correlated with the shock if counselors are more/less vigilant about the outcome of their match proposal and that this depends on the group status of the jobseeker.

³⁸The PES requires all posted vacancies to provide a minimum annual salary.

debate and, importantly, (some) of the general equilibrium effects induced by the changes in jobseeker and counselor search because we focus on a relatively small window of time. Hence, this paper arguably provides new and innovative evidence on the effect of an increase in discriminatory tastes on employer search patterns for minority employees. Returning to Table 2 we note that while employers reduce the rate at which they contact minorities for their permanent contracts, they actually increase the rate at which they show interest in minorities for fixed-term and temp work, producing a null overall effect on employer initiated potential matches. Thus it is not clear cut on exactly how employers are reacting to the discrimination shock. One possible explanation for this heterogeneity in effects across contract types is that employers become more risk averse about hiring minorities (Aigner and Cain, 1977) due to the shock and thus propose short-term contracts instead. Yet this is only speculation and does not speak to the change in composition of the types of minorities employers search for, nor if these minorities accept the offer from these employers. Regardless of the underlying aggregate hiring dynamics caused by the shock, it is employers who decide to hire in the end. Thus more work using real, observable employer behavior should be the focus of future research.

6 Conclusion

This paper exploits, rich, high-frequency data during the weeks preceding and following the January 2015 terrorist attacks in France to explore the effect of an increase in discrimination on the job search patterns of three principal labor market actors: jobseekers, job counselors and employers. We present evidence that this discrimination shock significantly reduced minority jobseeker effort in the weeks that followed the shock and that employers also reduce their search for minority candidates, but only for their best jobs. We show that this effect was partially compensated for by an increase in job counselor effort. This “compensatory” effect is only found in areas that exhibit low levels of existing prejudice against minorities, as measured by the local vote for the far-right political party in 2012. Furthermore, the effect is much larger for counselors who are themselves a minority. We also find strong positive effects for majority counselors, but only among those who specialize in getting the most marginalized jobseekers back to work. We see no increase in the effort of “normal” majority counselors. We argue that the mechanisms driving these effects appear to be that minority jobseekers reduce their search effort given the (perceived) lower returns to search, while counselors partially internalize the degraded search prospects of their minority jobseekers by increasing their effort. Overall, we find no significant impact on the number of workdays created for minorities compared majorities.

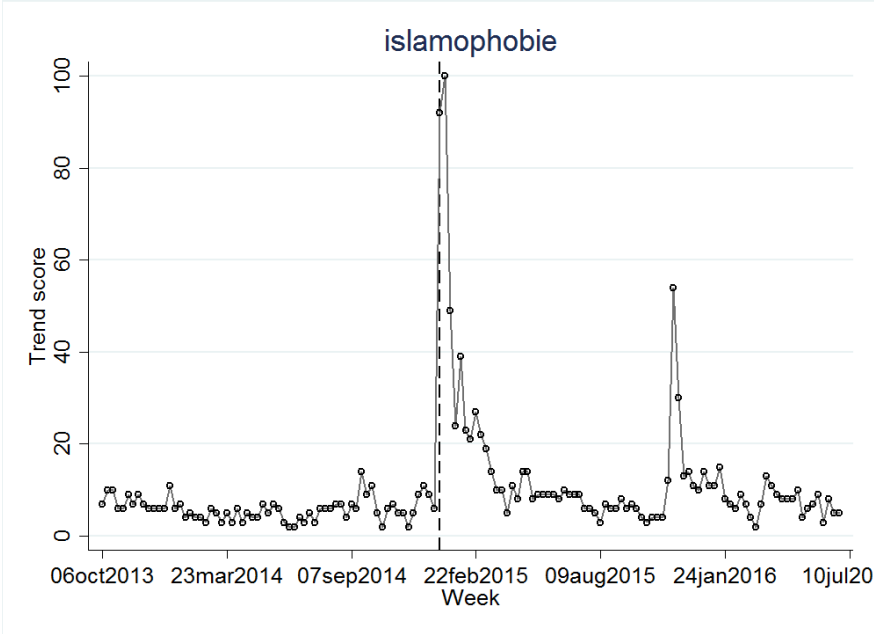
Using Google search trend data before and around the shock date, we try to better understand the presence and nature of the discrimination as well as its relationship to the far-right vote share. We find a significant positive correlation between terms that connote racial animus and the vote share while the relationship is the opposite for terms that connote social cohesion. Integrating this into the main analysis we find that the shock provokes an increase in discriminatory tastes, but that the marginal effect of this shock on job search is highest in areas with relatively lower levels of prior discrimination, suggesting that there may be strong diminishing returns to discrimination shocks.

And though we find impacts on the types of contracts minorities are hired in, we find it difficult to have a precise interpretation about the shock's effect on employer hiring behavior given that we document large and diverse impacts on the search behavior of, and for, minorities. This is due to the potentially large general equilibrium effects brought on by these changes in search intensity. Thus it is difficult to disentangle the effect of a change in minority search behavior from a change in employer bias.

Finally, the existence of a compensatory effect may not be isolated to intermediation from professional caseworkers as many jobs are found through informal channels and personal networks. Future work on this subject is needed to better understand how discrimination filters through labor market actors to build effective policies to combat its adverse effects.

Figures and Tables

Figure 1: Terrorist attack effect on Google searches for “islamophobie”



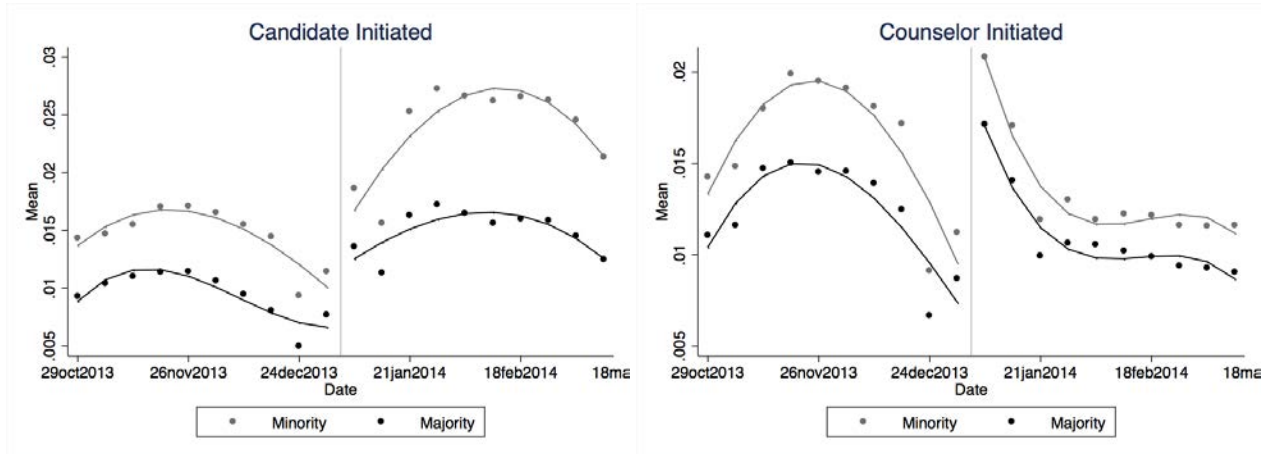
Note: Data are weekly series for the search interest in “islamophobie” in France. The vertical dashed line indicates the date of the January terrorist attacks. See the introduction for the interpretation of the search score.

Table 1: Descriptive statistics

	(1)	(2)	(3)
	All	Majority	Minority
<i>Jobseeker type</i>			
Unemp. looking for full-time work in permanent contract	0.708	0.694	0.772
Unemp. looking for part-time work in permanent contract	0.103	0.104	0.100
Unemp. looking for work in fixed-term, temp or seasonal contract	0.077	0.082	0.055
Unemp. but not immediately available for work	0.039	0.041	0.032
Emp. looking for other work	0.072	0.079	0.042
<i>Nationality</i>			
French	0.914	0.978	0.621
Maghreb	0.055	0.000	0.303
Western Europe	0.010	0.008	0.023
Sub-Saharan Africa	0.017	0.012	0.041
Other	0.004	0.003	0.012
<i>Demographics</i>			
Male	0.518	0.503	0.584
< 35 years	0.416	0.412	0.433
College degree	0.240	0.252	0.188
High qualification	0.596	0.623	0.475
Lives in Sensitive Urban Zone	0.084	0.051	0.234
<i>Profession searched for</i>			
Agriculture	0.041	0.044	0.026
Arts	0.006	0.006	0.004
Banking, insurance and real estate	0.013	0.013	0.010
Commercial and Sales	0.143	0.149	0.117
Communications, marketing and media	0.020	0.023	0.007
Construction	0.093	0.081	0.148
Hotel, restaurants and tourism	0.081	0.081	0.081
Manufacturing industry	0.082	0.083	0.080
Trades	0.041	0.042	0.038
Health	0.036	0.038	0.028
Personal services	0.202	0.195	0.231
Theater and film	0.024	0.028	0.006
IT, secretarial, accounting and RH	0.122	0.129	0.090
Transport	0.096	0.088	0.133
<i>Potential matches by initiator</i>			
Jobseeker	0.303	0.264	0.479
Counselor	0.186	0.179	0.218
Employer	0.037	0.037	0.040
<i>Hiring flows</i>			
Permanent	0.031	0.031	0.030
Fixed-term	0.317	0.341	0.209
Temp work	0.389	0.368	0.488
<i>Number of jobseekers by ALE</i>			
Ave. Num. of Jobseekers	5428	4456	972
-in low FN areas	5454	4333	1122
-in high FN areas	5401	4579	823
Agencies	810	810	810

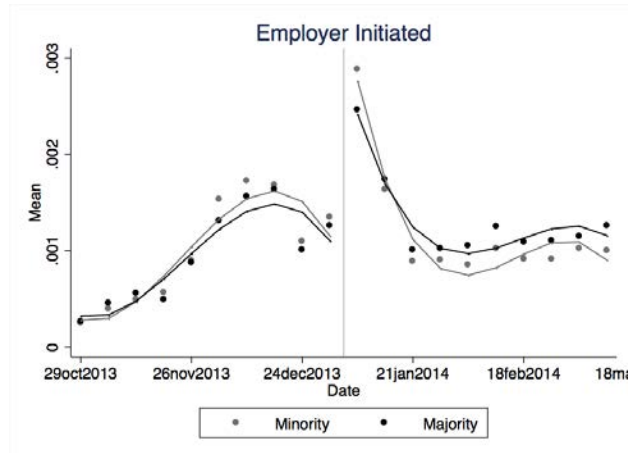
Note: Statistics are the agency average over the 10 week pre-shock period. Column 1 shows the overall average proportion per agency while columns 2 and 3 show the relative proportion within majority and minority populations, respectively.

Figure 2: Potential matches by channel in previous, “placebo” year



(a) Candidate

(b) Counselor



(c) Employer

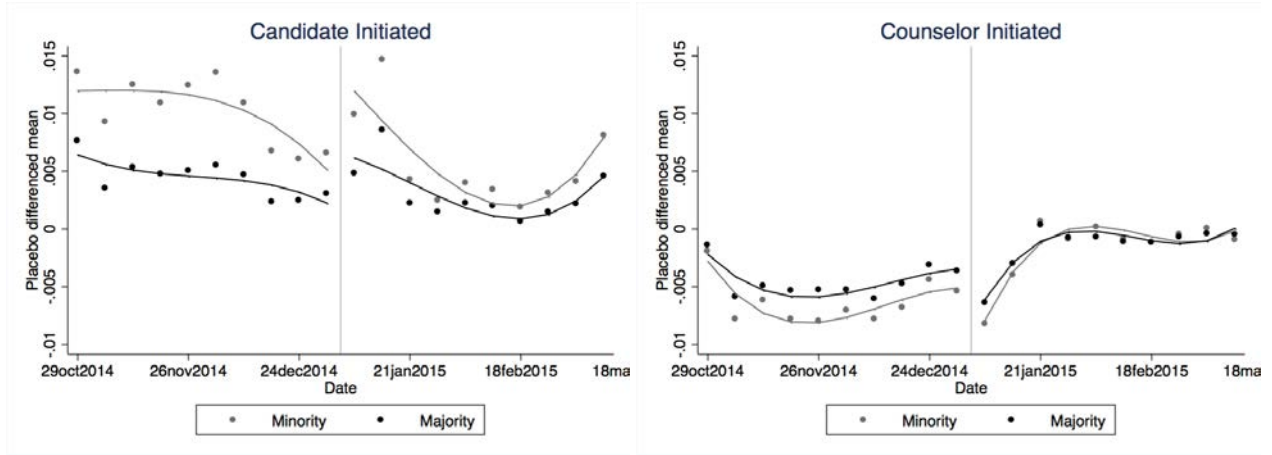
Note: Outcomes are the average number of potential matches made by jobseekers (a), counselors (b) and employers (c) for the same dates, but in the year previous to the January attacks. Observations are binned averages at the weekly level for majority and minority populations. The points are fitted using an OLS regression with a polynomial time trend of order 3. The vertical line indicates the week of the discontinuity date of the attack for the following year.

Table 2: Impact on potential matches by contract type

	(1)	(2)	(3)	(4)
	All potential matches	Jobseeker	Counselor	Employer
Panel A: All contracts				
(Minority*Period*Shock)	-0.00216** (0.00085)	-0.00512*** (0.00076)	0.00288*** (0.00036)	0.00008 (0.00007)
(Period*Shock)	0.00723*** (0.00031)	-0.00190*** (0.00022)	0.00812*** (0.00019)	0.00101*** (0.00003)
(Minority*Period)	0.00463*** (0.00063)	0.00744*** (0.00056)	-0.00260*** (0.00032)	-0.00022*** (0.00004)
(Minority*Shock)	0.00758*** (0.00073)	0.01122*** (0.00069)	-0.00362*** (0.00039)	-0.00002 (0.00006)
Period	0.00829*** (0.00028)	0.01090*** (0.00020)	-0.00324*** (0.00017)	0.00063*** (0.00002)
Minority	0.01814*** (0.00087)	0.01027*** (0.00056)	0.00752*** (0.00045)	0.00035*** (0.00004)
Shock	-0.00075** (0.00031)	0.00908*** (0.00019)	-0.01172*** (0.00022)	0.00189*** (0.00004)
Constant	0.04870*** (0.00049)	0.01733*** (0.00028)	0.02959*** (0.00029)	0.00177*** (0.00002)
Mean Dep. Var. Minority	0.07367	0.04790	0.02177	0.00400
N	64800	64800	64800	64800
Panel B: Permanent contracts				
(Minority*Period*Shock)	-0.00182*** (0.00050)	-0.00327*** (0.00047)	0.00162*** (0.00022)	-0.00017*** (0.00004)
Mean Dep. Var. Minority	0.03648	0.02482	0.00977	0.00189
N	64800	64800	64800	64800
Panel C: Fixed-term				
(Minority*Period*Shock)	-0.00063** (0.00032)	-0.00149*** (0.00026)	0.00080*** (0.00015)	0.00005** (0.00002)
Mean Dep. Var. Minority	0.02464	0.01530	0.00865	0.00069
N	64800	64800	64800	64800
Panel D: Temp				
(Minority*Period*Shock)	0.00031 (0.00022)	-0.00035* (0.00019)	0.00046*** (0.00008)	0.00020*** (0.00004)
Mean Dep. Var. Minority	0.01242	0.00771	0.00330	0.00141
N	64800	64800	64800	64800
Panel E: Seasonal				
(Minority*Period*Shock)	-0.00002 (0.00002)	-0.00002 (0.00001)	-0.00001 (0.00001)	0.00000 (0.00000)
Mean Dep. Var. Minority	0.00013	0.00008	0.00004	0.00000
N	64800	64800	64800	64800

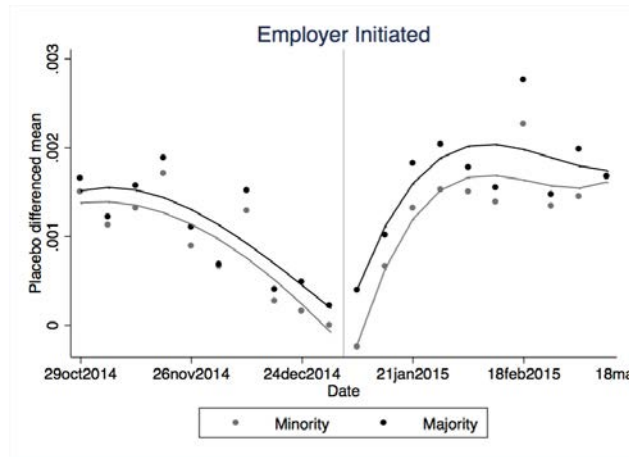
Note: This table presents impacts using weighted least squares on equation 2. The dependent variables are agency averages of potential matches (the number of potential matches divided by the number of registered jobseekers by group) separated by channel as denoted in the column titles. Regression equations are weighted by $\sqrt{n_{imtT}}$ where n_{imtT} is the number of jobseekers contributing to the observation's average. The mean of the dependent variable for minorities is the weekly mean of the dependent variable during the 10 weeks preceding the shock in the year of the shock ($t = 0, T = 1$). Panel A presents results of the estimation for all coefficients in equation 2 while panels B-E display only the estimate for β_1 . Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Figure 3: De-trended potential matches by channel



(a) Candidate

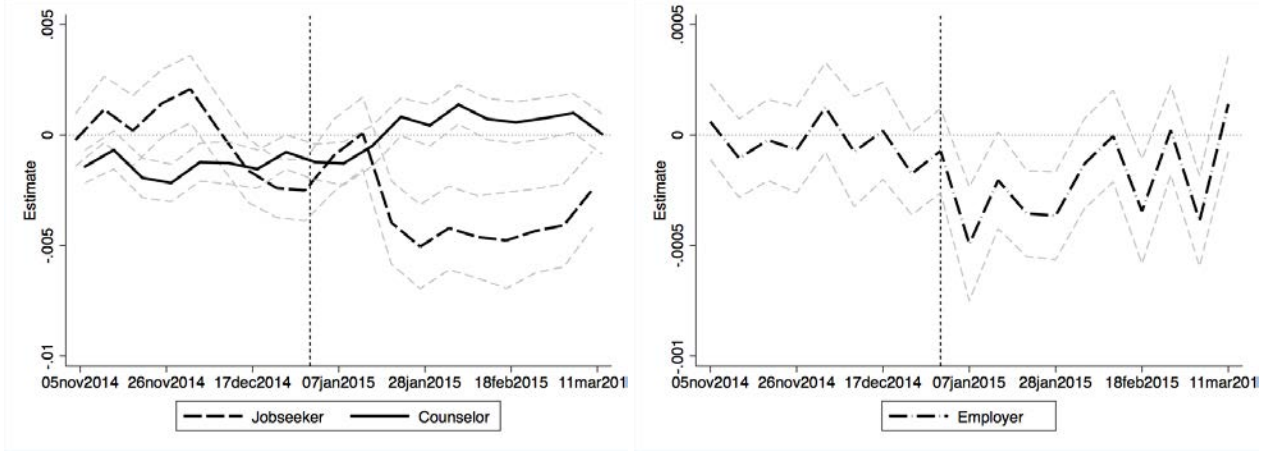
(b) Counselor



(c) Employer

Note: Graphs show the de-trended evolution of average potential matching rates binned by week for the shock year $T = 1$ for minority and majority jobseekers. Observations are de-trended by differencing out the equivalent placebo year ($T = 0$) weekly mean. The weighted means are fitted using an OLS regression with a polynomial time trend of order 3. The vertical line indicates the week of the discontinuity date of the attack. Potential matches are shown for the three matching channels, jobseeker (a), counselors (b) and employer (c).

Figure 4: Evolution of impact on search behavior on minorities



Note: Graphs plot estimates of the β_s coefficients from equation 4 over the weeks of the observation period for permanent contracts. The first graph on the left shows the evolution of the impact for jobseeker and counselor initiated matches while the graph on the right shows this evolution for employer search. Coefficients are connected by lines with grey dashed lines indicating 95% confidence intervals.

Table 3: Impact on contract flows

	(1)	(2)	(3)	(4)
	All Contracts	Permanent	Fixed	Temp
(Minority*Period*Shock)	0.00095** (0.00037)	0.00001 (0.00005)	0.00086*** (0.00020)	0.00009 (0.00030)
Mean Dep. Var. Minority	0.07261	0.00295	0.02091	0.04875
N	64800	64800	64800	64800

Note: This table presents impacts using weighted least squares on equation 2 where the dependent variables are agency averages of contract flows (the number of contracts divided by the number of registered jobseekers by group). The type of contract is denoted in the column titles. Regression equations are weighted by $\sqrt{n_{imtT}}$ where n_{imtT} is the number of jobseekers contributing to the observation's average. The mean of the dependent variable for minorities is the weekly mean of the dependent variable during the 10 weeks preceding the shock in the year of the shock ($t = 0, T = 1$). Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 4: Fixed-term contract flows

	(1) 1 day	(2) ≤ 7 days	(3) ≤ 1 month	(4) ≤ 3 months	(5) ≤ 6 months	(6) > 6 months
(Minority*Period*Shock)	0.00048*** (0.00014)	0.00025*** (0.00008)	0.00011* (0.00006)	0.00001 (0.00003)	0.00006** (0.00003)	-0.00006* (0.00003)
Mean Dep. Var. Minority	0.00748	0.00509	0.00438	0.00179	0.00109	0.00108
N	64800	64800	64800	64800	64800	64800

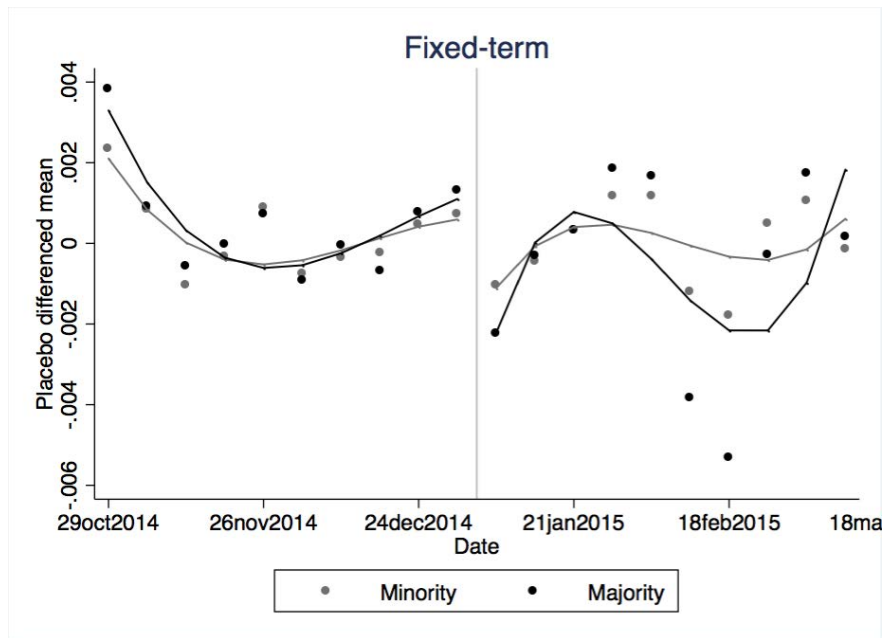
Note: This table replicates results from Table 3, but only for fixed-term contracts. The type of fixed-term contact is labeled in the column header and calculated using the start and end dates in the hiring declaration data. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 5: Impact on workdays

	(1) Permanent contract workdays	(2) Fixed-term workdays	(3) Permanent and fixed workdays
(Minority*Period*Shock)	0.00005 (0.00035)	0.05511 (4.91444)	0.05516 (4.91444)
Mean Dep. Var. Minority	0.021	12.659	12.679
N	64800	64800	64800

Note: Workdays are calculated using the start and end dates in the hiring declaration data. For permanent contracts we censor the date at the 20th week of the observation period within each year. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Figure 5: Detrended evolution of effect on fixed-term contract flows.



Note: Graphs show the de-trended evolution of the average number of fixed-term hires by week for the shock year $T = 1$ for each group of jobseeker. Observations are de-trended by differencing out the equivalent placebo year ($T = 0$) weekly mean.

Table 6: Name placebo tests for potential matches by channel

	(1)	(2)	(3)	(4)
	All potential matches	Jobseeker	Counselor	Employer
Panel A: Minority=Maghreb/Mashiq names				
(Minority*Period*Shock)	-0.00182***	-0.00327***	0.00162***	-0.00017***
	(0.00050)	(0.00047)	(0.00022)	(0.00004)
N	64800	64800	64800	64800
Panel B: Minority=British names				
(Minority*Period*Shock)	-0.00017	-0.00044	0.00030**	-0.00003
	(0.00043)	(0.00036)	(0.00015)	(0.00005)
N	64800	64800	64800	64800
p-value Equality of Coefs.	0.055	0.000	0.000	0.068
Panel C: Minority=Southern European names				
(Minority*Period*Shock)	0.00040	-0.00021	0.00058***	0.00003
	(0.00051)	(0.00045)	(0.00018)	(0.00006)
N	64800	64800	64800	64800
p-value Equality of Coefs.	0.014	0.000	0.001	0.025

Note: The dependent variable is the mean potential match rate by channel as denoted by the column titles. Panel A replicates results from Panel B of Table 2 using our baseline specification, Panels B and C define minority status using British and Southern European first names, respectively. p-values come from a test in the equality of coefficients of the *DDD* parameter between the baseline definition of minority status, Maghreb/Mashriq, and British and S. European. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 7: Name placebo tests for contract flows

	(1)	(2)	(3)	(4)
	All Contracts	Permanent	Fixed	Temp
Panel A: Minority=Maghreb/Mashiq names				
(Minority*Period*Shock)	0.00095** (0.00037)	0.00001 (0.00005)	0.00086*** (0.00020)	0.00009 (0.00030)
N	64800	64800	64800	64800
Panel B: Minority=British names				
(Minority*Period*Shock)	0.00031 (0.00044)	0.00004 (0.00007)	-0.00025 (0.00027)	0.00052 (0.00035)
N	64800	64800	64800	64800
p-value Equality of Coefs.	0.375	0.729	0.006	0.460
Panel C: Minority=Southern European names				
(Minority*Period*Shock)	0.00034 (0.00043)	-0.00010 (0.00007)	0.00036 (0.00028)	0.00008 (0.00031)
N	64800	64800	64800	64800
p-value Equality of Coefs.	0.375	0.321	0.230	0.985

Note: The dependent variable is the mean hiring rate by contract type as denoted by the column titles. Panel A replicates results from Table 3 using our baseline specification, Panels B and C define minority status using British and Southern European first names, respectively. p-values come from a test in the equality of coefficients of the *DDD* parameter between the baseline definition of minority status, Maghreb/Mashriq, and British and S. European. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 8: Robustness to controls

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: All potential matches						
(Minority*Period*Shock)	-0.00182*** (0.00050)	-0.00168*** (0.00051)	-0.00172*** (0.00051)	-0.00157*** (0.00051)	-0.00154*** (0.00050)	-0.00145*** (0.00051)
N	64800	64800	64800	64800	64800	64800
Panel B: Jobseeker						
(Minority*Period*Shock)	-0.00327*** (0.00047)	-0.00284*** (0.00046)	-0.00293*** (0.00045)	-0.00295*** (0.00044)	-0.00291*** (0.00044)	-0.00284*** (0.00044)
N	64800	64800	64800	64800	64800	64800
Panel C: Counselor						
(Minority*Period*Shock)	0.00162*** (0.00022)	0.00133*** (0.00019)	0.00139*** (0.00019)	0.00158*** (0.00019)	0.00153*** (0.00020)	0.00158*** (0.00019)
N	64800	64800	64800	64800	64800	64800
Panel D: Employer						
(Minority*Period*Shock)	-0.00017*** (0.00004)	-0.00018*** (0.00004)	-0.00018*** (0.00005)	-0.00020*** (0.00005)	-0.00016*** (0.00005)	-0.00018*** (0.00005)
N	64800	64800	64800	64800	64800	64800
Nationality	No	Yes	Yes	Yes	Yes	Yes
Type of jobseeker	No	No	Yes	Yes	Yes	Yes
Demographics	No	No	No	Yes	Yes	Yes
Sector	No	No	No	No	Yes	Yes
Regional FE's	No	No	No	No	No	Yes


Note: This table replicates results from Panel B of Table 2 (impact on potential matches to permanent contracts) using our baseline specification while progressively adding covariates. Each column shows the *DDD* estimates while adding the weekly mean of covariates in the pre-shock period from Table 1, interacted with the period t , year T and $t * T$ indicators as well their pairwise interactions with minority status, $m * t$, $m * T$ and $m * t * T$. Interactions are centered at the mean-group level. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 9: Impacts on search intensity controlling for compositional changes

	(1)	(2)	(3)	(4)
	All potential matches	Jobseeker	Counselor	Employer
(Minority*Period*Shock)	-0.00193*** (0.00051)	-0.00337*** (0.00048)	0.00161*** (0.00022)	-0.00018*** (0.00004)
Mean Dep. Var. Minority	0.036	0.025	0.010	0.002
N	64800	64800	64800	64800

Note: This table replicates results from Panel B of Table 2 (impact on potential matches to permanent contracts) including endogenous controls for potential compositional changes in the average number of highly qualified jobseekers registered at the agency, young jobseekers, those with Maghreb nationality and registered jobseekers searching for jobs in construction, the trades, or in the theater and film industry. Each of the controls is centered at the group-mean level and interacted with the period, year and minority indicators as well as their pairwise interactions. The endogenous controls are instrumented with their analogous counterparts using the pre-shock means. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Figure 6: Correlation between FN vote and discrimination sentiment pre-shock



Note: Graphs represent the correlation between the relative regional search interest in the term with the FN vote share in the region in the year preceding the attack. Lines are fitted values from estimating equation 11 where we control for the proportion of minority jobseekers in the region.

Table 10: Heterogeneity in impacts on potential matches across discrimination proxy

	(1)	(2)	(3)	(4)
	All potential matches	Jobseeker	Counselor	Employer
<i>Low FN</i>				
(Minority*Period*Shock)	-0.00230*** (0.00076)	-0.00460*** (0.00069)	0.00244*** (0.00033)	-0.00013** (0.00006)
N	32400	32400	32400	32400
<i>High FN</i>				
(Minority*Period*Shock)	-0.00141** (0.00058)	-0.00146** (0.00056)	0.00028 (0.00019)	-0.00024*** (0.00006)
N	32400	32400	32400	32400
p-value Equality of Coefs.	0.353	0.000	0.000	0.230


Note: This table replicates results from Panel B of Table 2 (impact on potential matches to permanent contracts) using our baseline specification, equation 2 separately for high and low FN vote share agencies (above or below the median). Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 11: Heterogeneous impacts on contract flow

	(1)	(2)	(3)	(4)
	All contracts	Permanent	Fixed-term	Temp
<i>Low FN</i>				
(Minority*Period*Shock)	0.00169*** (0.00051)	-0.00006 (0.00007)	0.00144*** (0.00028)	0.00032 (0.00039)
N	32400	32400	32400	32400
<i>High FN</i>				
(Minority*Period*Shock)	0.00013 (0.00056)	0.00008 (0.00007)	0.00017 (0.00028)	-0.00011 (0.00047)
N	32400	32400	32400	32400
p-value Equality of Coefs.	0.038	0.156	0.001	0.483

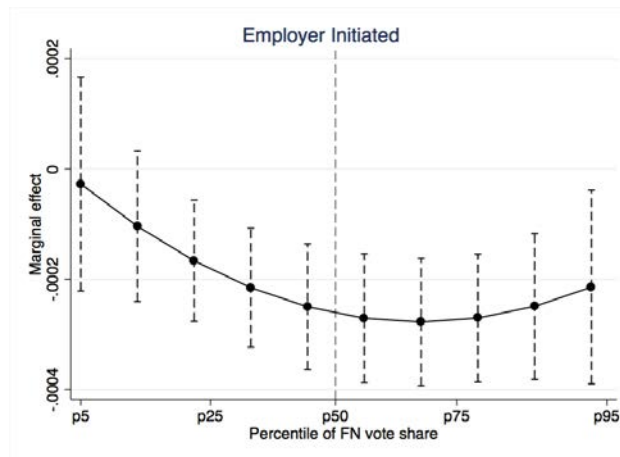
Note: This table replicates results from Table 3 (impact on contract flows contracts) using our baseline specification separately for high and low FN vote share agencies (above or below the median). Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Figure 7: Marginal effects on potential matches over support of existing discrimination



(a) Candidate

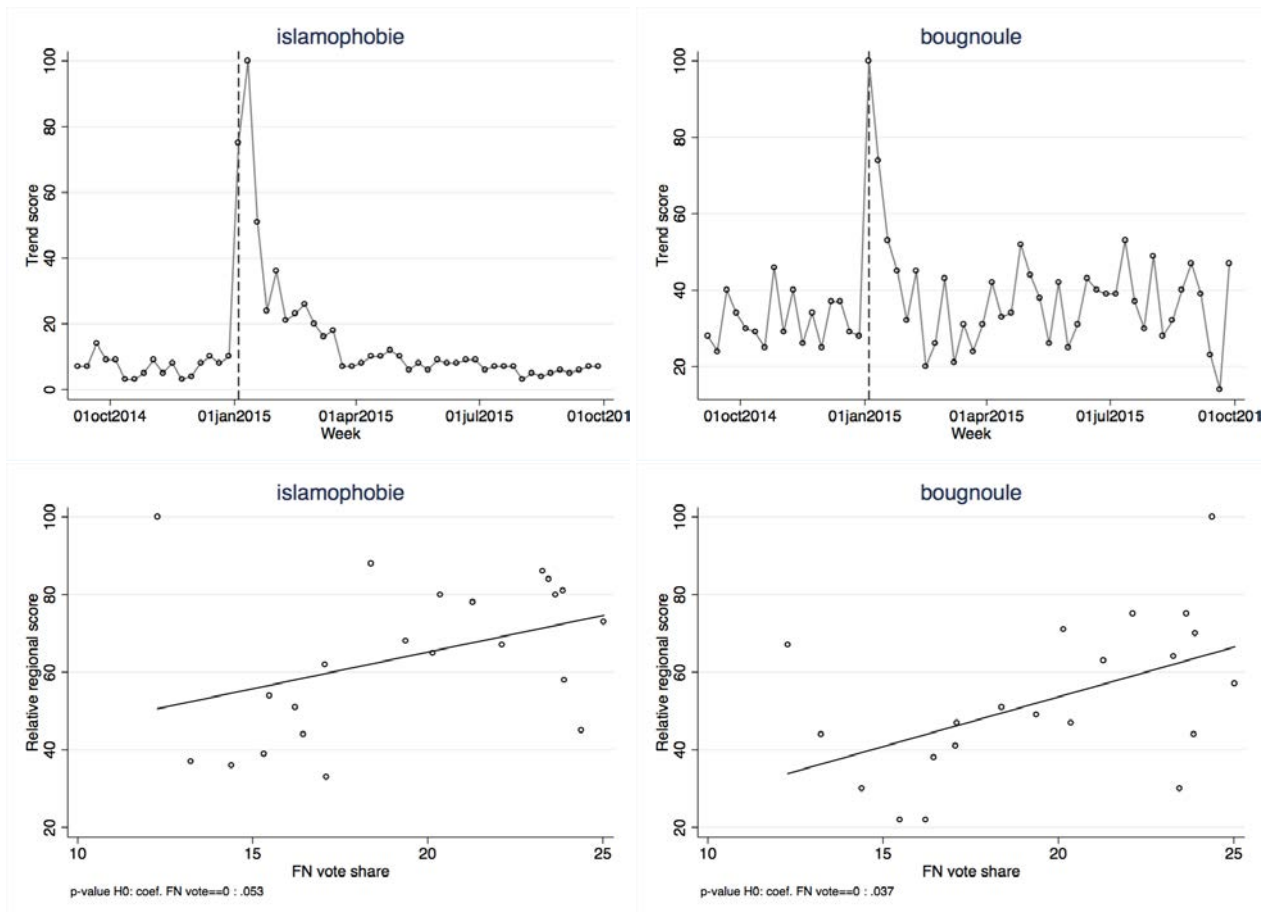
(b) Counselor



(c) Employer

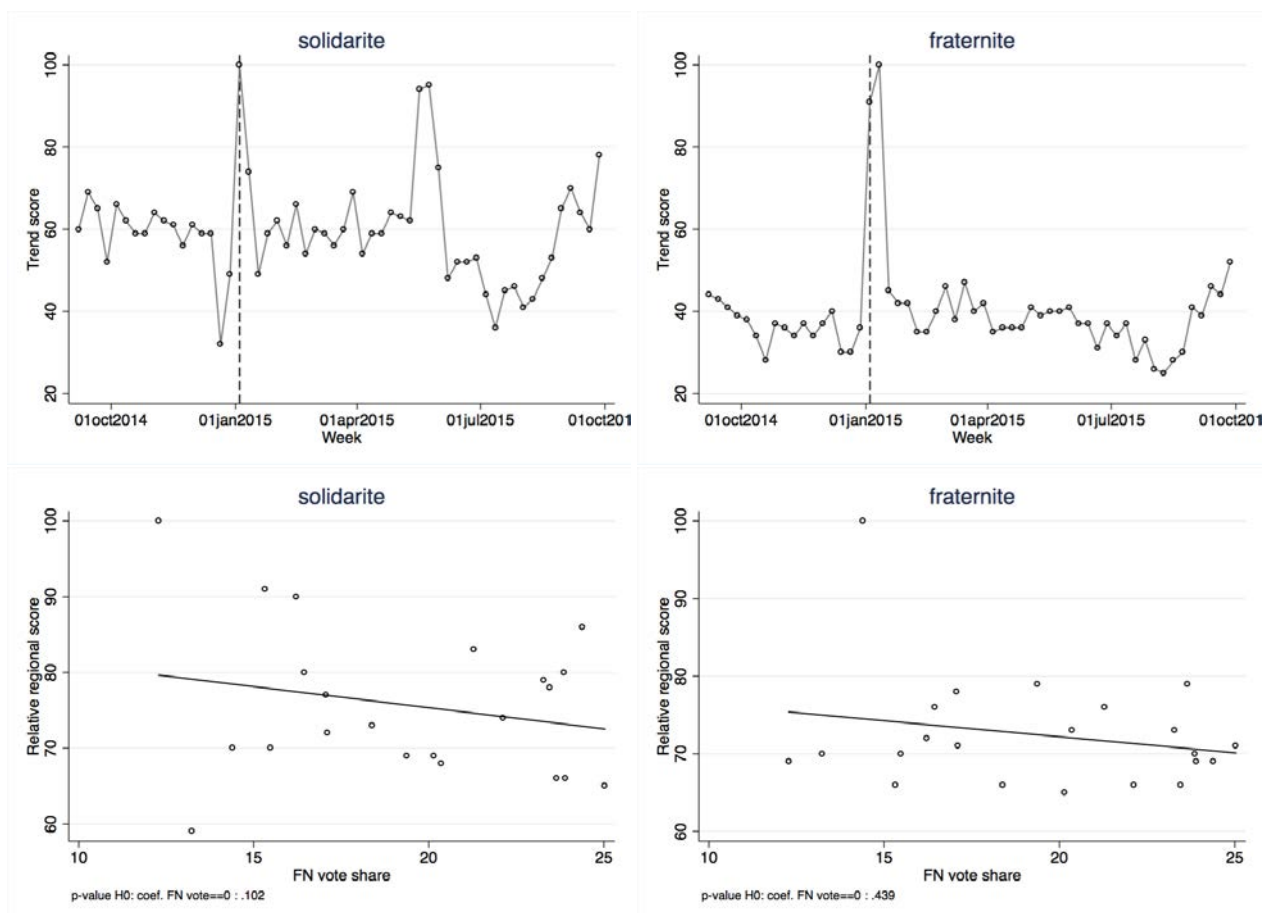
Note: These figures plot the marginal effect of the *DDD* parameter on potential matches made by our three labor market actors over the support of the agency-level FN vote share. Results come from a WLS regression of equation 2 where each term on the right-hand side are interacted with the continuous measure of the vote share and its square. 95% confidence intervals are shown using dashed lines.

Figure 8: Shock impact on search terms and correlation with FN vote



Note: Graphs in the top row are weekly series for the search interest for the term in the graph title in France. See the introduction for the interpretation of the search score. The vertical dashed line indicates the week of the shock. Graphs on the bottom row represent the correlation between the relative regional search interest in the term with the FN vote share during the same period.

Figure 9: Shock impact on positive search terms and correlation with FN vote



Note: Graphs in the top row are weekly series for the search interest for the term in the graph title in France. See the introduction for the interpretation of the search score. The vertical dashed line indicates the week of the shock. Graphs on the bottom row represent the correlation between the relative regional search interest in the term with the FN vote share during the same period.

Table 12: Change in relative regional score due to shock

	(1)	(2)	(3)	(4)
	Islamophobie	Bougnoule	Fraternite	Solidarite
Shock	6.94 (4.77)	13.28*** (2.72)	11.06** (4.56)	5.72* (3.31)
High FN	7.13 (7.59)	20.83*** (7.31)	-4.79 (6.37)	-3.75 (3.62)
(High FN)*(Shock)	3.67 (9.88)	2.71 (5.41)	2.07 (7.04)	-2.26 (3.97)
Prop. Minority	2.24*** (0.39)	0.69** (0.33)	0.54** (0.24)	1.13*** (0.20)
(Prop. Minority)*(Shock)	-0.33 (0.55)	0.06 (0.26)	-0.83*** (0.25)	-0.29 (0.23)
Constant	17.96** (7.86)	14.87** (6.28)	54.63*** (6.05)	55.90*** (3.31)
N	44	44	44	44

Note: This table presents impacts using ordinary least squares on equation 12. The dependent variables are the relative regional Google trend search score for the terms in the column titles. Standard errors in parenthesis are clustered at the regional level. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 13: Counselor matches by counselor minority status

	(1)	(2)	(3)	(4)
	All counselor matches	For Minority Jobseekers	For Majority Jobseekers	Difference: Minority - Majority
(Minority Counselor*Period*Shock)	0.00264*** (0.00089)	0.00398*** (0.00110)	0.00178* (0.00094)	0.00221** (0.00093)
(Period*Shock)	0.01050*** (0.00042)	0.01085*** (0.00049)	0.01028*** (0.00043)	0.00057* (0.00031)
N	635856	635856	635856	635856

Note: This table presents impacts using weighted least squares estimates on equation 13 for parameters λ_1 and λ_2 . The dependent variables are the average number of matches made by counselors for the type of jobseeker as denoted by the column titles. Regression equations are weighted by $\sqrt{n_{jt}^m}$ where n_{jt}^m is the number of minority jobseekers in the counselor's portfolio. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 14: Impacts for intensive versus normal counselors

	(1)	(2)	(3)
	All Counselors	Intensive counselors	Normal counselors
Dependent variable: Difference: (Minority - Majority) matches			
<i>All</i>			
(Minority Counselor*Period*Shock)	0.00221** (0.00093)	0.00303 (0.00499)	0.00229*** (0.00084)
(Period*Shock)	0.00057* (0.00031)	0.00363** (0.00143)	0.00017 (0.00030)
N	635856	113410	522446
<i>Low FN</i>			
(Minority Counselor*Period*Shock)	0.00315*** (0.00119)	0.00636 (0.00631)	0.00304*** (0.00109)
(Period*Shock)	0.00069 (0.00045)	0.00560*** (0.00213)	0.00006 (0.00044)
N	330971	58014	272957
<i>High FN</i>			
(Minority Counselor*Period*Shock)	-0.00001 (0.00143)	-0.00583 (0.00818)	0.00070 (0.00122)
(Period*Shock)	0.00042 (0.00042)	0.00132 (0.00186)	0.00030 (0.00040)
N	304885	55396	249489

Note: This table presents impacts using weighted least squares estimates on equation 13 for parameters λ_1 and λ_2 . The dependent variable is the difference between matches made to their minority jobseekers versus their majority jobseekers. Regression equations are weighted by $\sqrt{n_{jt}^m}$ where n_{jt}^m is the number of minority jobseekers in the counselor's portfolio. The first panel uses all counselor observations while the following two panels split the sample agencies that have below and above the median FN vote shares. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Appendix

Weighted Least Squares regression

Because our dependent variables are averages taken over the number of minority or majority jobseekers within the agency, it is natural to weight observations to recapture the true underlying population parameters and take into account that some means have larger variances than others. This last point is an issue of heteroskedasticity as explored in Solon, Haider, and Wooldridge (2015) who show that weighting for efficiency gains actually depends heavily on the data structure. They point out that estimating $\sqrt{n_i}y_i = \sqrt{n_i}X_i'\beta + \sqrt{n_i}u_i$ can lead to large gains in precision when there is large variation in the underlying group sizes on which the dependent variable is calculated, but will actually inflate standard errors if this is not the case. In section 3 of their paper they provide a test that entails obtaining the residuals from the baseline specification using OLS and then regressing these residuals on $1/n_i$ and then use the t-stat on the estimated coefficient on $1/n_i$ as a test for the presence of significant heteroskedasticity and thus the utility of using WLS. Performing this test using all applications as the dependent variable gives a $|t\text{-stat}| = 3.94$ suggesting that WLS may provide considerable benefit in terms of precision given our data structure. This is most certainly due to the fact that we have very few minority jobseekers in a significant proportion of agencies.

A synthetic placebo year

In this section we compare the comparability of two potential placebo years with the year of the shock. We then develop a synthetic placebo as a weighted average of the two years.³⁹ Unfortunately, PES data collection in their current form on potential matches started in 2013, thus we cannot look at another placebo year for these outcomes, but we can look at contract flows for the 2012-2013 period because the PES began collecting this data in March 2012.

Here, we develop a simple test that entails comparing $t = 0$ (i.e. pre-shock outcomes) differences between mean outcomes for $T = 0$ (previous year) and $T = -1$ (2 years prior) and outcomes in $T = 1$. In comparing the average weekly agency level difference in the pre-shock period in the current year with previous years, we can provide evidence about which placebo year is the most appropriate. The choice of appropriate year to use could simply be given by year that exhibits the smallest absolute $t = 0$ differences with the year

³⁹This technique could of course be generalizable to multiple placebo years.

of the shock:

$$\min\{(\bar{D}_0 = |\bar{y}_{t=0,T=1} - \bar{y}_{t=0,T=0}|), (\bar{D}_{-1} = |\bar{y}_{t=0,T=1} - \bar{y}_{t=0,T=-1}|)\}$$

Isolating the $t = 0$, period we can rearrange the data so that we have observations for D for each population in the local employment agency for each placebo year. Results from a WLS estimation of the following regression equation will show differences in the comparability of the two placebo years,

$$D_{i,m,t=0,T} = \delta_0 + \delta_1 Year_{i,m,t=0,T} + e_{i,m,t=0,T} \quad (14)$$

In this equation $Year = 1$ for the previous year to shock ($T = 0$) and $Year = 0$ for two years before the shock ($T = -1$). Hence, $\delta_0 + \delta_1$ gives the average comparability in outcomes of our baseline placebo year with the year of interest and the constant, δ_0 , the comparability of the placebo year taken two years before the shock. Table A.4 presents results from the estimation of equation 14 with contract flows as our dependent variables. For each type of contract, we separate the results by minority status and display the $t = 0$, $T = 1$ mean of the variable to gauge the size of the differentials. Ideally, we would like the estimate of $\delta_0 + \delta_1$ to be close to zero, meaning that on average there is very little difference in hiring rates between our baseline placebo year and the current year in the $t = 0$ period.

Looking at the results in panels A and B we see that the $T = 0$ placebo year appears to be much more comparable than the $T = -1$ year in terms of hiring rates of both minorities and majorities. Indeed we see that $|\hat{\delta}_0 + \hat{\delta}_1|$ is smaller across all contract types and groups than $|\hat{\delta}_0|$. In terms of the proportional difference off the mean of the original variable we see that for permanent contracts for minorities, the $T = 0$ difference is 2.8% while the $T = -1$ difference is 8.7%, almost 3 times as big. Looking at panel B we see the same story: $T = 0$ appears to be a better counterfactual than $T = -1$ in terms of average hiring rates for majorities as well.

We can visualize these average differences with the year of the shock in Figure A.3. Here we plot the binned de-trended data for both groups for the 10 weeks in the $t = 0$ time period. $T = 0$ data are solid dots and lines and the $T = -1$ data are dashed. We see that $T = 0$ data are much closer to zero, on average and, importantly, the differential between the two groups appears to be much more constant (this is especially true for permanent contracts).⁴⁰

Our estimates of δ_0 and δ_1 in equation 14 allow us to gauge the comparability of years

⁴⁰It should be clear that using the average distance between outcomes in the year of interest and the placebo years is arbitrary. For instance, one could develop weights based on the variance of the difference between minority and majority outcomes, i.e. the quality of the parallel trends pre-shock.

$T = 0$ and $T = -1$ with the year of interest, $T = 1$ in terms of average hiring rates in the pre-shock $t = 0$ period. Using these estimates at the panel entity level, we devise a simple weighting scheme to create a synthetic placebo year to use in the *DDD* estimation of the shock’s effect. Within agency i for population m , combined total absolute deviation over the two years is given by $|\delta_{im0}| + |\delta_{im1} + \delta_{im0}|$. Thus, the proportion of total deviation gives weights for each placebo year within agency i for population m .

$$w_{imt_0} = 1 - \frac{|\delta_{im1} + \delta_{im0}|}{|\delta_{im0}| + |\delta_{im1} + \delta_{im0}|} \quad (15)$$

and

$$w_{imt_{-1}} = 1 - \frac{|\delta_{im0}|}{|\delta_{im0}| + |\delta_{im1} + \delta_{im0}|} = 1 - w_{imt_0} \quad (16)$$

Hence the weights are panel entity-level scalars representing the proportion of each year’s deviation over the total deviation observed over all previous years. The “synthetic” placebo year is then generated as the weighted sum of the dependent variable over the two years for each population m in agency i :

$$y_{imt_s} = w_{imt_0} * y_{imt_0} + (1 - w_{imt_0}) * y_{imt_{-1}} \quad (17)$$

Hence in its general form, we are able to “blend” as many placebo years as we have data for using the weights generated from the $t = 0$ data. Yet, given we have shown that our $T = 0$ year is clearly more comparable than the $T = -1$ year, we should probably consider results from the *DDD* estimation using this synthetic data as a conservative estimate. In essence, we are allowing ourselves to include data of potentially lower “counterfactual quality” in the synthetic placebo, albeit with lower importance.⁴¹ Therefore, we believe it is useful to present results using the synthetic placebo as an “informed robustness check” to our main results. Table A.5 compares results from our baseline specification on contract flows with results obtained from using the synthetic placebo year. We find that results are similar using the two methods. While the coefficients are indeed more conservative using the synthetic control we cannot reject the null hypothesis in the equality of the *DDD* parameter for any contract type.

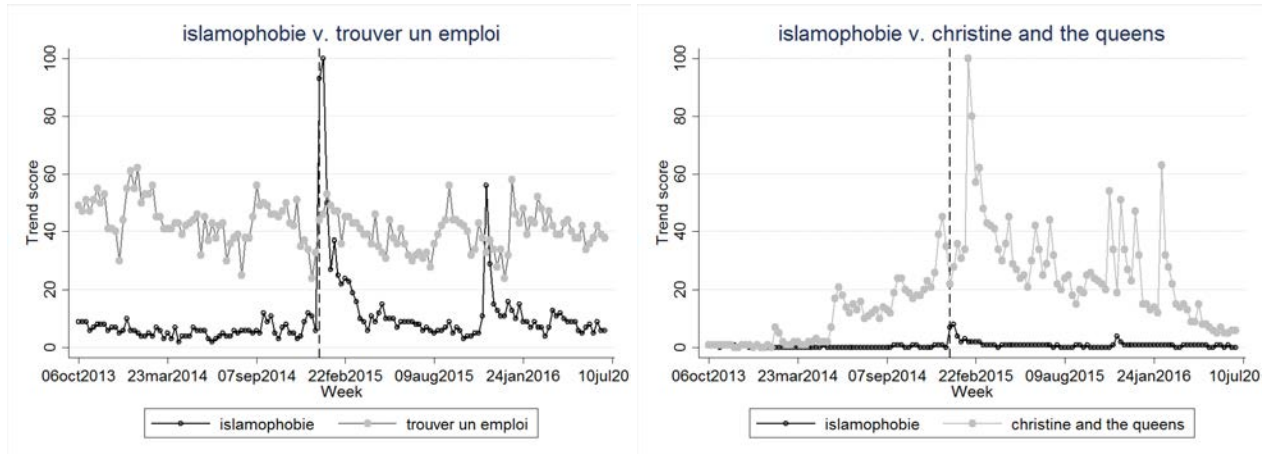
In summary, this simple method introduces a way to be selective about placebo years used in a *DDD* framework and provides a way to include data from multiple potential placebo years in a non arbitrary way in order to test the robustness of results. The test and method developed here are very simple and could be a possible avenue for more serious econometric

⁴¹Indeed, we have no rule-of-thumb about at which point one should or should not use a placebo year. Do the differentials shown above have to be below 5%, 10%, etc.?

work in the future.

Appendix Figures and Tables

Figure A.1: Google searches for “islamophobic” compared to other search terms



Note: Graphs are weekly series for the search interest for the terms in the graph title in France. See the introduction for the interpretation of the search score. The vertical dashed line indicates the week of the January 2015 terrorist attacks.

Table A.1: Potential matches for all contract types (unweighted)

	(1)	(2)	(3)	(4)
	All	Jobseeker	Counselor	Employer
Panel A: All contracts				
(Minority*Period*Shock)	-0.00303*** (0.00098)	-0.00311*** (0.00083)	0.00095*** (0.00034)	-0.00001 (0.00011)
(Period*Shock)	-0.00082** (0.00037)	-0.00203*** (0.00027)	0.00857*** (0.00021)	0.00104*** (0.00004)
(Minority*Period)	0.00383*** (0.00066)	0.00429*** (0.00056)	-0.00094*** (0.00026)	-0.00014** (0.00005)
(Minority*Shock)	0.00568*** (0.00089)	0.00759*** (0.00076)	-0.00195*** (0.00031)	-0.00009 (0.00010)
Period	0.01838*** (0.00033)	0.01131*** (0.00024)	-0.00367*** (0.00018)	0.00064*** (0.00002)
Minority	0.01150*** (0.00073)	0.00583*** (0.00054)	0.00494*** (0.00030)	0.00028*** (0.00006)
Shock	0.00531*** (0.00036)	0.00939*** (0.00025)	-0.01228*** (0.00024)	0.00199*** (0.00004)
Constant	0.05240*** (0.00056)	0.01802*** (0.00032)	0.03040*** (0.00031)	0.00185*** (0.00003)
Mean Dep. Var. Minority	0.07488	0.04082	0.02111	0.00402
N	64800	64800	64800	64800
Panel B: Permanent contracts				
(Minority*Period*Shock)	-0.00164*** (0.00058)	-0.00156*** (0.00050)	0.00006 (0.00017)	-0.00021*** (0.00007)
Mean Dep. Var. Minority	0.03421	0.02014	0.00853	0.00184
N	64800	64800	64800	64800
Panel C: Fixed-term				
(Minority*Period*Shock)	-0.00097** (0.00043)	-0.00102*** (0.00034)	0.00060*** (0.00021)	-0.00002 (0.00005)
Mean Dep. Var. Minority	0.02738	0.01371	0.00914	0.00069
N	64800	64800	64800	64800
Panel D: Temp				
(Minority*Period*Shock)	-0.00037 (0.00027)	-0.00050** (0.00023)	0.00030*** (0.00011)	0.00022*** (0.00006)
Mean Dep. Var. Minority	0.01310	0.00687	0.00337	0.00149
N	64800	64800	64800	64800
Panel E: Seasonal				
(Minority*Period*Shock)	-0.00006 (0.00004)	-0.00003 (0.00002)	-0.00001 (0.00002)	0.00000 (0.00000)
Mean Dep. Var. Minority	0.00020	0.00010	0.00007	0.00000
N	64800	64800	64800	64800

Note: This table replicates results in Table 2 using OLS. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Table A.2: Potential matches for all contract types (Count data)

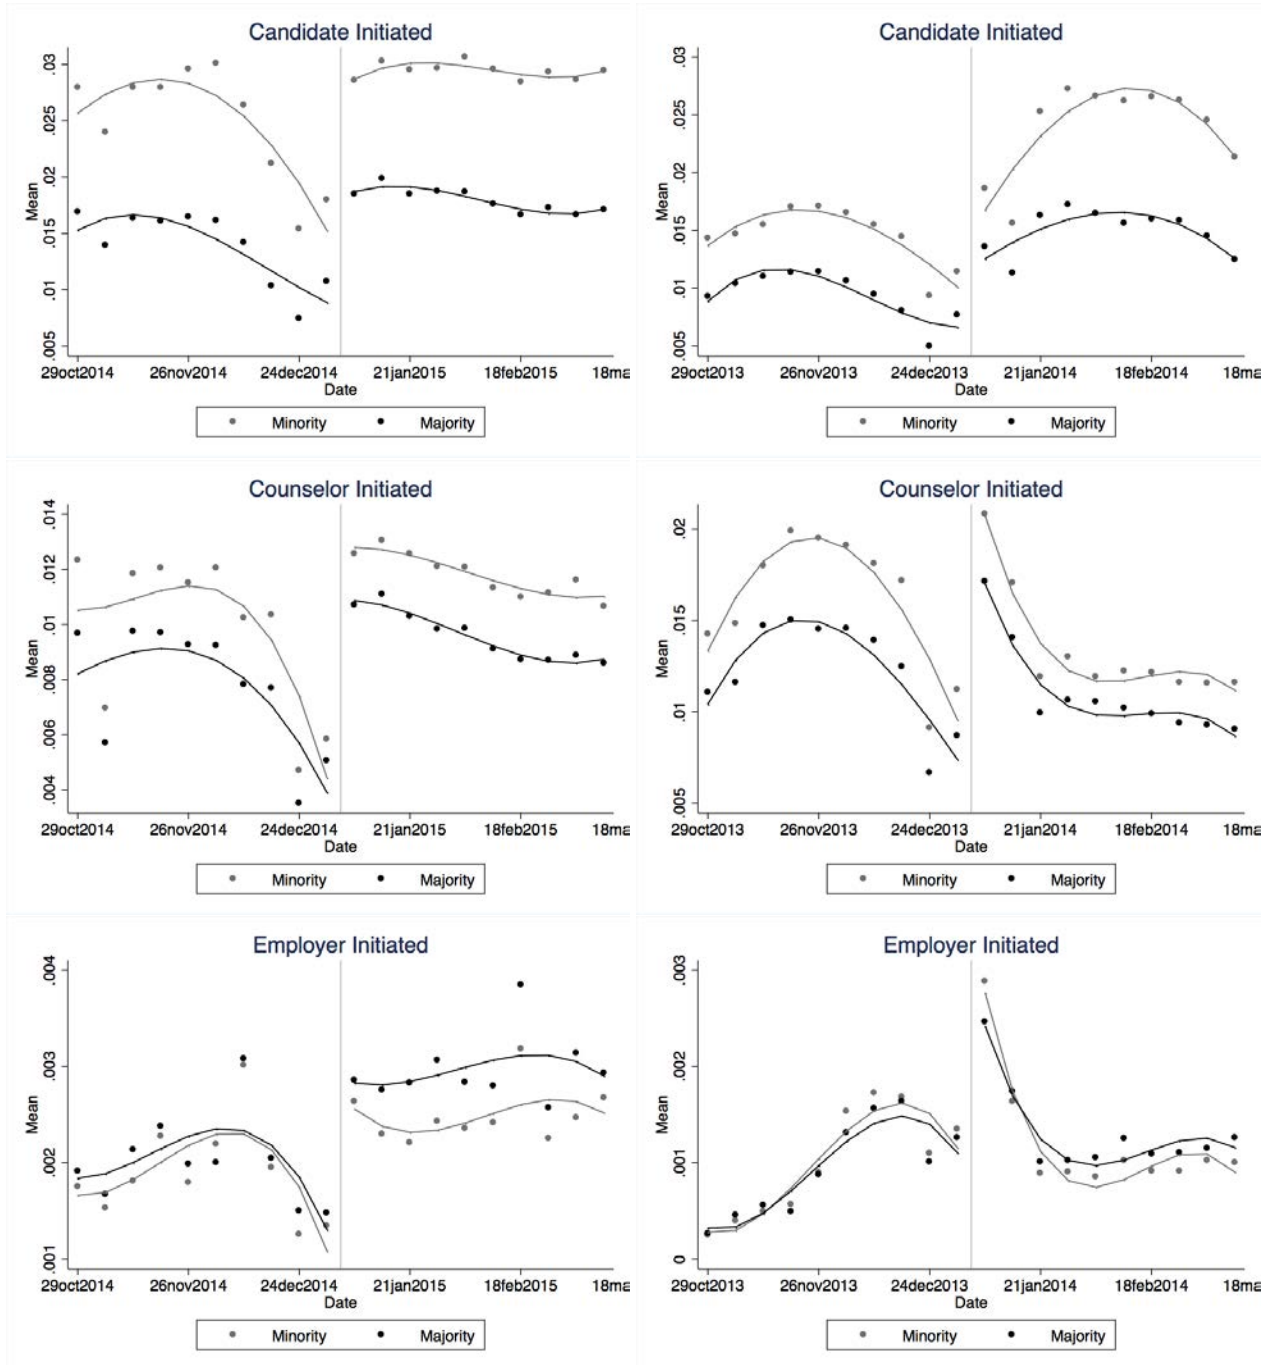
	(1)	(2)	(3)	(4)
	All	Jobseeker	Counselor	Employer
Panel A: All contracts				
(Minority*Period*Shock)	-0.05387*** (0.01015)	-0.10367*** (0.01593)	0.02543** (0.01120)	0.07547*** (0.01982)
(Minority*Period)	0.00950 (0.00711)	0.03281*** (0.01182)	-0.04400*** (0.00900)	-0.11515*** (0.01553)
(Minority*Shock)	0.10203*** (0.00997)	0.15496*** (0.01598)	-0.00406 (0.01165)	-0.06923*** (0.01884)
(Period*Shock)	-0.03898*** (0.00524)	-0.19827*** (0.00737)	0.35393*** (0.00749)	0.06103*** (0.00948)
Shock	0.14573*** (0.00608)	0.46954*** (0.00761)	-0.45630*** (0.00879)	0.77397*** (0.00998)
Period	0.30925*** (0.00398)	0.49366*** (0.00543)	-0.11018*** (0.00612)	0.31099*** (0.00759)
Minority	-1.23479*** (0.03634)	-1.08228*** (0.03866)	-1.32086*** (0.03621)	-1.36558*** (0.03660)
Constant	5.37198*** (0.01504)	4.29869*** (0.01990)	4.83339*** (0.01520)	2.01899*** (0.01713)
Mean Dep. Var. Minority	80.2	46.6	21.2	3.9
N	64800	64800	64800	64800
Panel B: Permanent contracts				
(Minority*Period*Shock)	-0.06031*** (0.01240)	-0.12199*** (0.01869)	0.04951*** (0.01453)	0.03963 (0.02733)
Mean Dep. Var. Minority	39.5	24.1	9.5	1.8
N	64800	64800	64800	64800
Panel C: Fixed-term				
(Minority*Period*Shock)	-0.05375*** (0.01103)	-0.09837*** (0.01789)	0.01980 (0.01249)	0.04062 (0.04332)
Mean Dep. Var. Minority	27.3	14.9	8.4	0.7
N	64800	64800	64800	64800
Panel D: Temp				
(Minority*Period*Shock)	-0.01920 (0.01645)	-0.04137 (0.02555)	-0.01150 (0.02062)	0.10406*** (0.03164)
Mean Dep. Var. Minority	13.3	7.5	3.2	1.4
N	64800	64800	64800	64800
Panel E: Seasonal				
(Minority*Period*Shock)	-0.09028 (0.10558)	-0.17887 (0.18452)	0.00630 (0.13925)	0.59955 (0.51951)
Mean Dep. Var. Minority	0.1	0.1	0.0	0.0
N	64800	64800	64800	64800

Note: This table replicates results in Table 2 using Poisson regression. The dependent variables are count data as opposed to averages. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Figure A.2: Evolution of potential matches by channel, by year

(a) Shock year $T = 1$

(b) Placebo year $T = 0$



Note: Outcomes are the average number of potential matches made by jobseekers, counselors and employers. Graphs on the left display binned averages for the shock year $T = 1$ while graphs on the right display results for the placebo year $T = 0$. Observations are bins of the weighted average at the weekly level for majority and minority populations. The points are then fitted using an OLS regression with a polynomial time trend of order 3. The vertical line indicates the week of the discontinuity date regardless of the year.

Table A.3: Changes in jobseeker composition within agencies

	(1)	(2)
	Minority mean for $t = 0, T = 1$	\overline{DDD}
Unemp. looking for full-time work in permanent contract	0.77217	-0.00007 (0.00041)
Unemp. looking for part-time work in permanent contract	0.09969	-0.00011 (0.00023)
Unemp. looking for work in fixed-term, temp or seasonal contract	0.05479	-0.00003 (0.00018)
Unemp. but not immediately available for work	0.03162	0.00037 (0.00023)
Emp. looking for other work	0.04174	-0.00016 (0.00023)
High qualification	0.47459	0.00095*** (0.00037)
Lives in Sensitive Urban Zone	0.23354	0.00044 (0.00040)
Male	0.58398	0.00032 (0.00034)
< 35 years	0.43258	-0.00110*** (0.00035)
College degree	0.18838	0.00012 (0.00029)
French	0.62076	0.00069** (0.00034)
Maghreb	0.30309	-0.00088*** (0.00031)
Western Europe	0.02276	-0.00011 (0.00011)
Sub-Saharan Africa	0.04137	0.00020 (0.00013)
Other	0.01203	0.00010 (0.00008)
Agriculture	0.02556	-0.00021* (0.00012)
Arts	0.00397	-0.00008* (0.00004)
Banking, insurance and real estate	0.00997	0.00011 (0.00007)
Commercial and Sales	0.11748	0.00033 (0.00026)
Communications, marketing and media	0.00743	-0.00000 (0.00007)
Construction	0.14761	-0.00089*** (0.00026)
Hotel, restaurants and tourism	0.08134	-0.00013 (0.00020)
Manufacturing industry	0.08018	0.00014 (0.00019)
Trades	0.03805	0.00028** (0.00014)
Health	0.02764	0.00005 (0.00013)
Personal services	0.23138	-0.00028 (0.00030)
Theater and film	0.00631	0.00023*** (0.00007)
IT, secretarial, accounting and RH	0.08963	0.00024 (0.00020)
Transport	0.13278	0.00028 (0.00023)
N=64800		

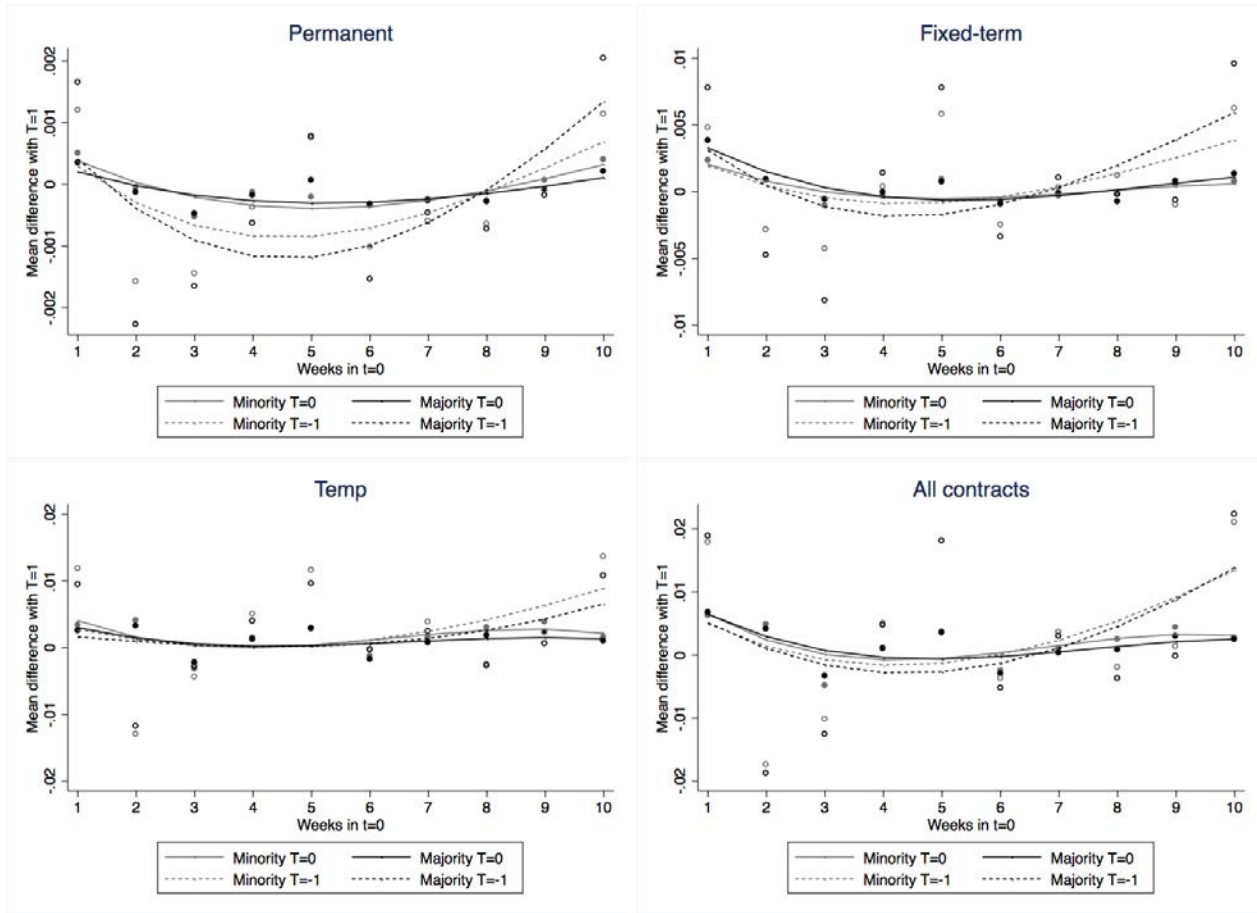
Note: Each row displays results from a separate regression using our *DDD* specification. The dependent variables are the average proportion for each compositional variable as denoted in the first column. We also display the pre-shock mean to gauge effect sizes. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Table A.4: Validity test for placebo year

	(1)	(2)	(3)	(4)
	All contracts	Permanent contract	Fixed-term	Interim
Panel A: Minority				
Year	-0.00162*** (0.00037)	0.00017*** (0.00004)	-0.00060*** (0.00018)	-0.00118*** (0.00032)
Constant	0.00348*** (0.00044)	-0.00026*** (0.00004)	0.00085*** (0.00020)	0.00289*** (0.00039)
Mean original var.	0.07261	0.00295	0.02091	0.04875
Proportional difference T=0	0.02564	-0.02825	0.01162	0.03492
Proportional difference T=-1	0.04791	-0.08653	0.04051	0.05923
N	16200	16200	16200	16200
Panel B: Majority				
Year	-0.00118*** (0.00021)	0.00018*** (0.00003)	-0.00055*** (0.00015)	-0.00081*** (0.00016)
Constant	0.00272*** (0.00025)	-0.00029*** (0.00003)	0.00106*** (0.00018)	0.00195*** (0.00018)
Mean original var.	0.07394	0.00312	0.03406	0.03675
Proportional difference T=0	0.02075	-0.03694	0.01488	0.03109
Proportional difference T=-1	0.03676	-0.09308	0.03099	0.05315
N	16200	16200	16200	16200

Note: This table presents results from estimating equation 14 with results separated by group status. The mean of the original variable is the weekly mean in the $t = 0, T = 1$ period. The proportional difference is calculated as $\frac{\hat{\delta}_1 + \hat{\delta}_0}{\text{Mean of orig. var}}$ for $T = 0$ and $\frac{\hat{\delta}_0}{\text{Mean of orig. var}}$ for $T = -1$. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Figure A.3: Placebo year and trend comparability with shock year



Note: Graphs show the binned de-trended data for the 10 weeks in the $t = 0$ time period. Data that are de-trended using $T = 0$ are solid dots and lines while $T = -1$ are dashed. The binned means are fitted using an OLS regression with a polynomial time trend of order 3.

Table A.5: Comparison with synthetic placebo year

	(1)	(2)	(3)	(4)
	All contracts	Permanent contract	Fixed-term	Interim
Panel A: $T = 0$ placebo				
(Minority*Period*Shock)	0.00095** (0.00037)	0.00001 (0.00005)	0.00086*** (0.00020)	0.00009 (0.00030)
N	64800	64800	64800	64800
Panel B: Synthetic placebo				
(Minority*Period*Shock)	0.00055* (0.00033)	0.00005 (0.00004)	0.00051*** (0.00018)	-0.00001 (0.00027)
N	64800	64800	64800	64800
p-value Equality of Coefs.	0.413	0.528	0.181	0.813

Note: Panel A duplicates results from Table 4 while panel B reproduces the results using a synthetic placebo year i.e. a weighted blend of outcomes from multiple years preceding the shock. See the second section in the appendix for details. Standard errors in parenthesis are clustered at the agency level. * $p < .1$, ** $p < .05$, *** $p < .01$

Bibliography

- ABADIE, A., A. DIAMOND, AND J. HAINMUELLER (2010): “Synthetic control methods for comparative case studies: Estimating the effect of California’s tobacco control program,” *Journal of the American statistical Association*, 105, 493–505.
- ADIDA, C. L., D. D. LAITIN, AND M.-A. VALFORT (2014): “Muslims in France: Identifying a discriminatory equilibrium,” *Journal of Population Economics*, 27, 1039–1086.
- (2016): ““One Muslim is enough!” Evidence from a field experiment in France,” *Annals of Economics and Statistics/Annales d’Économie et de Statistique*, 121–160.
- AEBERHARDT, R., D. FOUGÈRE, J. POUGET, AND R. RATHELOT (2010): “L’emploi et les salaires des enfants d’immigrés,” *Économie et statistique*, 433, 31–46.
- AGERSTRÖM, J., R. CARLSSON, AND D.-O. ROTH (2007): “Ethnicity and obesity: Evidence of implicit work performance stereotypes in Sweden,” *Working Paper, IFAU-Institute for Labour Market Policy Evaluation*.
- AIGNER, D. J. AND G. G. CAIN (1977): “Statistical theories of discrimination in labor markets,” *Industrial and Labor relations review*, 175–187.
- ALGAN, Y., T. MAYER, M. THOENIG, ET AL. (2013): “The Economic Incentives of Cultural Transmission: Spatial Evidence from Naming Patterns across France,” Tech. rep., CEPR Discussion Papers.
- AMODIO, D. M. AND P. G. DEVINE (2006): “Stereotyping and evaluation in implicit race bias: Evidence for independent constructs and unique effects on behavior,” *Journal of Personality and Social Psychology*, 91, 652.
- ANGRIST, J. D. AND J.-S. PISCHKE (2008): *Mostly harmless econometrics: An empiricist’s companion*, Princeton university press.
- ARROW, K. (1973): “The theory of discrimination,” *Discrimination in Labor Markets*, 3, 3–33.
- ÅSLUND, O. AND D.-O. ROTH (2005): “Shifts in attitudes and labor market discrimination: Swedish experiences after 9-11,” *Journal of Population Economics*, 18, 603–629.
- BANSE, R., J. SEISE, AND N. ZERBES (2001): “Implicit attitudes towards homosexuality: Reliability, validity, and controllability of the IAT,” *Zeitschrift für Experimentelle Psychologie*, 48, 145–160.

- BEAMAN, L., R. CHATTOPADHYAY, E. DUFLO, R. PANDE, AND P. TOPALOVA (2009): “Powerful women: Does exposure reduce prejudice?” *Quarterly Journal of Economics*, 124, 1497–1540.
- BECKER, G. S. (1957): *The Economics of Discrimination*, University of Chicago Press, Chicago.
- BEHAGHEL, L., B. CRÉPON, AND T. LE BARBANCHON (2015): “Unintended effects of anonymous résumés,” *American Economic Journal: Applied Economics*, 7, 1–27.
- BEHNCKE, S., M. FRÖLICH, AND M. LECHNER (2010): “A caseworker like me—does the similarity between the unemployed and their caseworkers increase job placements?” *The Economic Journal*, 120, 1430–1459.
- BELL, B., R. BLUNDELL, AND J. VAN REENEN (1999): “Getting the unemployed back to work: the role of targeted wage subsidies,” *International tax and public finance*, 6, 339–360.
- BERSON, C. (2012): “Does Competition Induce Hiring Equity?” Tech. rep., HAL.
- BERTRAND, M., D. CHUGH, AND S. MULLAINATHAN (2005): “Implicit discrimination,” *American Economic Review: Papers and Proceedings*, 95, 94–98.
- BERTRAND, M. AND E. DUFLO (2017): “Field experiments on discrimination,” *Handbook of Economic Field Experiments*, 1, 309–393.
- BERTRAND, M., E. DUFLO, AND S. MULLAINATHAN (2004): “How much should we trust differences-in-differences estimates?” *The Quarterly Journal of Economics*, 119, 249–275.
- 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, 991–1013.
- BLUNDELL, R. AND M. C. DIAS (2009): “Alternative approaches to evaluation in empirical microeconomics,” *Journal of Human Resources*, 44, 565–640.
- BOISJOLY, J., G. J. DUNCAN, M. KREMER, D. M. LEVY, AND J. ECCLES (2006): “Empathy or antipathy? The impact of diversity,” *American Economic Review*, 96, 1890–1905.
- BORDALO, P., K. COFFMAN, N. GENNAIOLI, AND A. SHLEIFER (forthcoming): “Stereotypes,” *Quarterly Journal of Economics*.

- CAMERON, A. C., J. B. GELBACH, AND D. L. MILLER (2008): “Bootstrap-based improvements for inference with clustered errors,” *Review of Economics and Statistics*, 90, 414–427.
- CARD, D., J. KLUVE, AND A. WEBER (2015): “What works? A meta analysis of recent active labor market program evaluations,” Tech. rep., National Bureau of Economic Research.
- COATE, S. AND G. C. LOURY (1993): “Will affirmative-action policies eliminate negative stereotypes?” *American Economic Review*, 83, 1220–1240.
- COFFMAN, K. B. (2014): “Evidence on self-stereotyping and the contribution of ideas,” *Quarterly Journal of Economics*, 129, 1625–1660.
- CRÉPON, B., E. DUFLO, M. GURGAND, R. RATHELOT, AND P. ZAMORA (2013): “Do labor market policies have displacement effects? Evidence from a clustered randomized experiment,” *The quarterly journal of economics*, 128, 531–580.
- DEE, T. S. (2004): “Teachers, race, and student achievement in a randomized experiment,” *Review of Economics and Statistics*, 86, 195–210.
- (2005): “A teacher like me: Does race, ethnicity, or gender matter?” *American Economic Review*, 95, 158–165.
- DICKENS, R., D. BROWN, P. GREGG, S. MACHIN, AND A. MANNING (2001): “Everything under a fiver: recruitment and retention in lower paying labour markets,” .
- DOVIDIO, J. F. AND S. L. GAERTNER (2008): “New directions in aversive racism research: Persistence and pervasiveness,” in *Motivational Aspects of Prejudice and Racism*, Springer, 43–67.
- DOVIDIO, J. F., K. KAWAKAMI, AND S. L. GAERTNER (2002): “Implicit and explicit prejudice and interracial interaction,” *Journal of Personality and Social Psychology*, 82, 62.
- EGLOFF, B. AND S. C. SCHMUKLE (2002): “Predictive validity of an Implicit Association Test for assessing anxiety,” *Journal of Personality and Social Psychology*, 83, 1441.
- FRIESE, M., M. BLUEMKE, AND M. WÄNKE (2007): “Predicting voting behavior with implicit attitude measures: The 2002 German parliamentary election,” *Experimental Psychology*, 54, 247.

- FRÖLICH, M. (2008): “Parametric and nonparametric regression in the presence of endogenous control variables,” *International Statistical Review*, 76, 214–227.
- GARTHWAITE, C., T. GROSS, AND M. J. NOTOWIDIGDO (2014): “Public health insurance, labor supply, and employment lock,” *Quarterly Journal of Economics*, 129, 653–696.
- GAUTIER, P. A., A. SIEGMANN, AND A. VAN VUUREN (2009): “Terrorism and attitudes towards minorities: The effect of the Theo van Gogh murder on house prices in Amsterdam,” *Journal of Urban Economics*, 65, 113–126.
- GIULIANO, L., D. I. LEVINE, AND J. LEONARD (2009): “Manager race and the race of new hires,” *Journal of Labor Economics*, 27, 589–631.
- (2011): “Racial bias in the manager-employee relationship: An analysis of quits, dismissals, and promotions at a large retail firm,” *Journal of Human Resources*, 46, 26–52.
- GLOVER, D., A. PALLAIS, AND W. PARIENTE (2017): “Discrimination as a Self-Fulfilling Prophecy: Evidence from French Grocery Stores,” *The Quarterly Journal of Economics*, 132, 1219–1260.
- 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, 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, 1464.
- GREENWALD, A. G., B. A. NOSEK, AND M. R. BANAJI (2003): “Understanding and using the Implicit Association Test: I. An improved scoring algorithm,” *Journal of Personality and Social Psychology*, 85, 197–216.
- GREENWALD, A. G., T. A. POEHLMAN, E. L. UHLMANN, AND M. R. BANAJI (2009): “Understanding and using the Implicit Association Test: III. Meta-analysis of predictive validity,” *Journal of Personality and Social Psychology*, 97, 17.
- HAHN, A., C. M. JUDD, H. K. HIRSH, AND I. V. BLAIR (2014): “Awareness of implicit attitudes,” *Journal of Experimental Psychology: General*, 143, 1369–1392.

- HAIDER, A. H., J. SEXTON, N. SRIRAM, L. A. COOPER, D. T. EFRON, S. SWOBODA, C. V. VILLEGAS, E. R. HAUT, M. BONDS, P. J. PRONOVOST, ET AL. (2011): “Association of unconscious race and social class bias with vignette-based clinical assessments by medical students,” *Journal of the American Medical Association*, 306, 942–951.
- HEBL, M. R., J. B. FOSTER, L. M. MANNIX, AND J. F. DOVIDIO (2002): “Formal and interpersonal discrimination: A field study of bias toward homosexual applicants,” *Personality and Social Psychology Bulletin*, 28, 815–825.
- HECKMAN, J. J. (1998): “Detecting discrimination,” *The Journal of Economic Perspectives*, 12, 101–116.
- HIRANO, K., G. W. IMBENS, AND G. RIDDER (2003): “Efficient estimation of average treatment effects using the estimated propensity score,” *Econometrica*, 71, 1161–1189.
- HJORT, J. (2014): “Ethnic divisions and production in firms,” *Quarterly Journal of Economics*, 129, 1899–1946.
- HU, X., J. P. ROSENFELD, AND G. V. BODENHAUSEN (2012): “Combating automatic autobiographical associations: The effect of instruction and training in strategically concealing information in the Autobiographical Implicit Association Test,” *Psychological Science*, 23, 1079–1085.
- IMBENS, G. W. AND D. B. RUBIN (2015): *Causal inference in statistics, social, and biomedical sciences*, Cambridge University Press.
- KAUSHAL, N., R. KAESTNER, AND C. REIMERS (2007): “Labor market effects of September 11th on Arab and Muslim residents of the United States,” *Journal of Human Resources*, 42, 275–308.
- KIM, D.-Y. (2003): “Voluntary controllability of the Implicit Association Test (IAT),” *Social Psychology Quarterly*, 66, 83–96.
- LANE, K. A., M. R. BANAJI, B. A. NOSEK, AND A. G. GREENWALD (2007): “Understanding and using the Implicit Association Test: IV. What we know (so far) about the method,” in *Implicit Measures of Attitudes*, ed. by B. Wittenbrink and N. Schwarz, The Guilford Press, 59–102.
- LANG, K. AND J.-Y. K. LEHMANN (2012): “Racial Discrimination in the Labor Market: Theory and Empirics,” *Journal of Economic Literature*, 50, 959–1006.

- LANG, K., M. MANOVE, AND W. T. DICKENS (2005): “Racial discrimination in labor markets with posted wage offers,” *The American Economic Review*, 95, 1327–1340.
- LEE, M.-J. (2016): “Generalized Difference in Differences With Panel Data and Least Squares Estimator,” *Sociological Methods & Research*, 45, 134–157.
- LUNDBERG, S. J. AND R. STARTZ (1983): “Private discrimination and social intervention in competitive labor market,” *American Economic Review*, 73, 340–347.
- MANNING, A. (2011): “Imperfect Competition in the Labor Market,” *Handbook of Labor Economics*, 4, 973 – 1041.
- MAS, A. AND E. MORETTI (2009): “Peers at work,” *American Economic Review*, 99, 112–145.
- MAYER, N. AND P. PERRINEAU (1996): *Le Front national à découvert*, Presses de Sciences Po.
- MCCONNELL, A. R. AND J. M. LEIBOLD (2001): “Relations among the Implicit Association Test, discriminatory behavior, and explicit measures of racial attitudes,” *Journal of Experimental Social Psychology*, 37, 435–442.
- MICHAILLAT, P. (2012): “Do matching frictions explain unemployment? Not in bad times,” *The American Economic Review*, 102, 1721–1750.
- MOONEY, C. (2014): “Across America, whites are biased and they don’t even know it,” *The Washington Post: Wonkblog*.
- MORTENSEN, D. T. AND C. A. PISSARIDES (1994): “Job creation and job destruction in the theory of unemployment,” *The review of economic studies*, 61, 397–415.
- NOSEK, B. A., A. G. GREENWALD, AND M. R. BANAJI (2007): “The Implicit Association Test at age 7: A methodological and conceptual review,” in *Automatic Processes in Social Thinking and Behavior*, ed. by J. Bargh, Psychology Press, 265–292.
- ONZUS (2011): “Observatoire National des Zones Urbaines Sensibles,” Tech. rep., Secrétariat général du Comité interministériel des villes.
- OYER, P., S. SCHAEFER, ET AL. (2011): “Personnel Economics: Hiring and Incentives,” *Handbook of Labor Economics*, 4, 1769–1823.

- PETIT, P., E. DUGUET, L. DU PARQUET, F. SARI, ET AL. (2011): “Discriminations à l’embauche des jeunes franciliens et intersectionnalité du sexe et de l’origine: Les Résultats d’un testing,” Tech. rep., Centre d’Études des Politiques Économiques (EPEE), Université d’Evry Val d’Essonne.
- PHELPS, E. S. (1972): “The Statistical Theory of Racism and Sexism,” *American Economic Review*, 62, 659–661.
- PISSARIDES, C. A. (2000): *Equilibrium unemployment theory*, MIT press.
- PRENDERGAST, C. (2007): “The motivation and bias of bureaucrats,” *The American Economic Review*, 180–196.
- PRICE, J. AND J. WOLFERS (2010): “Racial discrimination among NBA referees,” *Quarterly Journal of Economics*, 125, 1859–1887.
- RAO, G. (2014): “Familiarity does not breed contempt: Diversity, discrimination and generosity in Delhi schools,” *Working Paper*.
- RATCLIFFE, A. AND S. VON HINKE KESSLER SCHOLDER (2015): “The london bombings and racial prejudice: Evidence from the housing and labor market,” *Economic Inquiry*, 53, 276–293.
- RIACH, P. A. AND J. RICH (2002): “Field experiments of discrimination in the market place,” *The economic journal*, 112.
- ROOTH, D.-O. (2010): “Automatic associations and discrimination in hiring: Real world evidence,” *Labour Economics*, 17, 523–534.
- ROSENTHAL, R. AND L. JACOBSON (1968): “Pygmalion in the classroom,” *Urban Review*, 3, 16–20.
- SKANDALIS, D. AND A. PHILIPPE (2016): “The Discouraged-Worker Effect: How do job-seekers respond to information about local labor market tightness?” .
- SMITH-McLALLEN, A., B. T. JOHNSON, J. F. DOVIDIO, AND A. R. PEARSON (2006): “Black and white: The role of color bias in implicit race bias,” *Social Cognition*, 24, 46–73.
- SOLON, G., S. J. HAIDER, AND J. M. WOOLDRIDGE (2015): “What are we weighting for?” *Journal of Human resources*, 50, 301–316.
- STAUFFER, J. M. AND M. R. BUCKLEY (2005): “The existence and nature of racial bias in supervisory ratings,” *Journal of Applied Psychology*, 90, 586–591.

- STEPHENS-DAVIDOWITZ, S. (2014): “The cost of racial animus on a black candidate: Evidence using Google search data,” *Journal of Public Economics*, 118, 26 – 40.
- STOLL, M. A., S. RAPHAEL, AND H. J. HOLZER (2004): “Black job applicants and the hiring officer’s race,” *Industrial & Labor Relations Review*, 57, 267–287.
- TODD, E. (2015): *Qui est Charlie?. Sociologie d’une crise religieuse*, Le Seuil.