AN EXPERIMENT IN HIRING DISCRIMINATION
VIA ONLINE SOCIAL NETWORKS

Alessandro Acquisti and Christina Fong

Carnegie Mellon University

Abstract. We investigate whether personal information posted by job candidates on social media sites is sought and used by prospective employers. We create profiles for job candidates on popular social networks, manipulating information protected under U.S. laws, and submit job applications on their behalf to over 4,000 employers. We find evidence of employers searching online for the candidates. After comparing interview invitations for a Muslim versus a Christian candidate, and a gay versus a straight candidate, we find no difference in callback rates for the gay candidate compared to the straight candidate, but a 13% lower callback rate for the Muslim candidate compared to the Christian candidate. While the difference is not significant at the national level, it exhibits significant and robust heterogeneity in bias at the local level, compatible with existing theories of discrimination. In particular, employers in Republican areas exhibit significant bias both against the Muslim candidate, and in favor of the Christian candidate. This bias is significantly larger than the bias in Democratic areas. The results are robust to using state- and county-level data, to controlling for firm, job, and geographical characteristics, and to several model specifications. The results suggest that 1) the online disclosure of certain personal traits can influence the hiring decisions of U.S. firms and 2) the likelihood of hiring discrimination via online searches varies across employers. The findings also highlight the surprisingly lasting behavioral influence of traditional, offline networks in processes and scenarios where online interactions are becoming increasingly common.

1. Introduction

The rise of Internet and social media services like online social networks has created new channels through which employers and job candidates can find information about each other. Those channels can facilitate and improve the matching between firms and workers. However, job seekers also reveal, online, information that would not be easily discovered during the interview process, and which may be even illegal for employers to request or use in the hiring process. Thus, although new online tools can facilitate labor market matching, they may also create a new arena for labor market discrimination.

To date, no field data has demonstrated how online information affects the hiring behavior of U.S. firms. In surveys, employers admit to using various online services to research job candidates.[1] However, the Equal Employment Opportunity Commission (EEOC) has cautioned firms about risks associated with searching online for protected characteristics,[2] and thus may have dissuaded some firms from using social media in the hiring process. Some states have even drafted bills limiting employers' ability to access candidates' online information.[3] Thus, whether hiring bias results from personal information posted online remains an open question.

Furthermore, in surveys, employers claim to use social media merely to seek job-relevant information about candidates.[4] However, much more private information can be gleaned from the online presences of prospective hires. On a social media profile, a status update can reveal a candidate's place of worship, a comment can suggest sexual orientation, and a personal photo can reveal ethnic origins. In a country known for its cultural assimilation of immigrants (Vigdor 2008), private differences that have traditionally been scrubbed out for work or education might now be more visible online. Whether U.S. employers react to such online personal information, rather than merely to the online professional information they may seek, is not known.

We present a randomized field experiment testing the joint hypothesis that (i) firms search online for information about job applicants and (ii) change their hiring behavior according to manipulated online

---

[1] Various surveys suggest that U.S. employers search job candidates online, but the reported frequency of searches varies considerably across surveys. See "Survey Shows 48% of Employers Conduct Social Media Background Checks." *EmployeeScreenIQ*. Accessed February 26, 2016. http://www.employeescreen.com/iqblog/social-media-2/48-of-employers-conduct-social-media-background-checks/; "Employers Using Social Networks for Screening Applicants." *Wikibin.* Accessed February 26, 2016. http://wikibin.org/articles/employers-using-social-networks-for-screening-applicants.html; "Ponemon Institute/Littler Mendelson Study." *International Association of Privacy Professionals.* Accessed February 26, 2016. https://www.privacyassociation.org/publications/2007_12_ponemon_institute_littler_mendelson_study; Johnston, Stuart J. 2010. "Microsoft Survey: Online 'Reputation' Counts." *Internet News*, January 27. Accessed February 26, 2016. http://www.internetnews.com/webcontent/article.php/3861241/Microsoft+Survey+Online+Reputation+Counts.htm.

[2] See Theodore Claypoole, "EEOC Regulations Spotlight Social Media," Womble Carlyle, last modified May 24, 2011, accessed February 26, 2016, http://www.wcsr.com/Insights/Alerts/2011/May/EEOC-Regulations-Spotlight-Social-Media. On the potential risks of using social media in hiring for firms, see Harpe (2009).

[3] For instance, Texas S.B. 118 aims to prohibit an employer from "requiring or requesting access to the personal accounts of employees and job applicants through electronic communication devices; establishing an unlawful employment practice." See S.B. 118, 83rd Leg., Reg. Sess. (Tex. 2013).

[4] See surveys listed in footnote 1.

personal information. The experiment relies on a methodology consisting of the creation and careful design of online presences of fictional individuals. We design social media profiles for four job candidates – a Muslim versus a Christian candidate, and a gay versus a straight candidate – to manipulate personal information that may be hard to discern from résumés and interviews, and that may be protected either under federal or some state laws (henceforth referred to as "protected information").[5] We manipulate the candidates' personal information exclusively via their online profiles, using material revealed online by actual members of popular social networking sites and job seeking sites. Candidates' professional background and résumés are kept constant across conditions.

After vetting the realism and quality of candidates' online profiles in a randomized online pilot experiment (henceforth the "online pilot"), we submit résumés and cover letters on behalf of those four candidates to over 4,000 U.S. job openings, with a single application sent to each employer (henceforth the "field experiment"). The résumés and letters contain no references or links to the candidates' manipulated personal information: To be treated by our manipulation of online profiles, the employer must independently choose to search online for information about the candidate using the name indicated on the submitted résumé. The main dependent variable in the field experiment is the number of interview invitations each candidate receives (i.e., callbacks). We compare the callback rate for the Christian candidate to that for the Muslim candidate, and the callback rate for the straight candidate to that for the gay candidate. We control for independent variables used in related literature (e.g., Bertrand and Mullainathan 2004; Tilcsik 2011), including firm characteristics, job characteristics, and geographical characteristics of the job's location. We also use Google and LinkedIn data to estimate the frequency with which employers search the candidates online. We test for a stronger callback bias in states and counties that have a higher prevalence of demographic traits associated with bias in prior research. More negative attitudes to both Muslims and gay people have been found among survey respondents who are Republican, older, and who do not personally know Muslims or gay people (Arab American Institute 2012; Pew Research Center 2012; Pew Research Center 2013). Thus, we test for stronger bias in firms located in states and counties that have a high fraction of Republican voters, a high median age, and low fractions of Muslims (in the Muslim-Christian manipulation) or gay people (in the gay-straight manipulation).

Nationwide, we detect no difference in callback rates for the gay candidate compared to the straight candidate, but we find 13% fewer callbacks for the Muslim candidate compared to the Christian candidate. While the difference is not significant at the national level, it exhibits significant and robust heterogeneity in bias at the local level. We find that discrimination a) against the Muslim candidate and b) in favor of the

---

[5] Different types of personal information enjoy different levels of protection across U.S. states. Some personal traits cannot even be inquired about in interviews, while others cannot be used in the hiring decision. Some are protected across all states, and others only in some states. For simplicity, we refer to information about these traits collectively as "protected" information, but investigate state-level differences in the degree of their protection through our empirical analysis.

Christian candidate varies significantly and robustly with employer characteristics in manners predicted by both theoretical and empirical previous work. In counties with a high fraction of Republican voters, the callback rates for the Muslim and Christian candidates are, respectively, 6.25% and 22.58%. In contrast, in the Democratic counties, the callback rates for the Muslim and Christian candidates are, respectively, 12.37% and 12.13%. The results are robust at the state level.

Using the callback rates for the gay and straight candidates as a benchmark, we find that our results are driven by both a *negative* bias against the Muslim candidate, and a *positive* bias toward the Christian candidate. Furthermore, we find that the findings are robust to the inclusion of firm characteristics (including firm size and ownership status) and a host of additional controls, as well as to different categorizations of Republican, mixed, and Democratic states or counties based on electoral results and Gallup Organization surveys. These results are consistent with those of the online pilot, where we find significant bias against the Muslim candidate, relative to the Christian candidate, among subjects with hiring experience who self-identified as Republican and Christian.

We consider possible explanations for the results – under the caveat that causal identification is limited by non-random assignment of employers to geographical areas. Without allowing ad hoc variation in discriminatory tastes, heterogeneous bias can be explained as group beneficial cooperation. Such cooperation can vary across communities because it requires specific conditions which are not satisfied by all communities (Ostrom 1998, 2000; Bowles and Gintis 2004; Chong 2000). These conditions include smaller community size and less migration into and out of the communities, yielding predictions that are testable in, and in fact compatible with, our data: we find significantly more bias in smaller communities and areas with less geographical migration.

These findings make a number of contributions to the literature. First, an emerging stream of work in information systems research investigates the impact of online information systems on offline behaviors (for instance, Chan and Ghose 2014 investigate the potential role of online platforms such as Craigslist in the transmission of STDs). The findings of our experiment suggest that, while hiring discrimination via online searches of candidates may not yet be widespread, online disclosures of personal traits can significantly influence the hiring decisions of a selected set of employers. At the same time, the results suggest an intriguing phenomenon for scholars of online transactions: to the extent that group-beneficial cooperation plays a role in explaining the findings, the experiment highlights the lasting behavioral influence of traditional offline networks of people physically close to each other. Even as online networks and interactions become more common, they may sometimes facilitate parochial cooperation in local, physical networks. Second, the findings highlight an emerging tension between modern information systems and institutional regulation written for a pre-Internet world. The latter aims to preclude certain information from being used in the hiring process; the former can effectively bypass legislation by allowing

individuals to make their information openly available to others online. A similar tension is being studied in the growing empirical (Miller and Tucker 2009; Goldfarb and Tucker 2011) and theoretical (Acquisti, Taylor, and Wagman 2015) information systems literature on privacy and its economic aspects (Stigler 1980; Posner 1981). Third, the paper introduces a methodology – the manipulation of Internet profiles – for field experiments on digital discrimination. One advantage of this method, which we exploit here, is its suitability for investigating discrimination associated with traits that may be protected under a country's laws or that may be difficult to realistically manipulate in résumé-only or audit studies (as most job candidates may refrain from revealing certain types of personal information on résumés). In fact, this methodology (creating online presences to test dynamics in the offline world) may be used not only in studies of discrimination in the job market, but also in comparable studies of bias in access to credit, housing, or educational opportunities. Fourth and finally, this paper illustrates a wrinkle in the literature investigating the role of information in economic outcomes. Economists have long been interested in the role of information (Stigler 1962) and signaling (Spence 1973) in job market matching. Recent work has highlighted how hiring mechanisms that *reduce* candidates' information available to employers (for example, by making the candidate anonymous) may *increase* interview opportunities for certain categories of applicants (Goldin and Rouse 2000; Aslund and Skans 2012) and may raise social welfare (Taylor and Yildirim 2011). In this manuscript, conversely, we find that Internet and social media platforms can affect interview opportunities for some categories of applicants at the expense or benefit of others, by making more information about candidates available to employers.[6] The extent to which this new information channel will improve labor search efficiency (Kroft and Pope 2014) and reduce labor market frictions (by allowing better matching of employers and employees), or will in fact lead to more discrimination, is likely to be an issue of increasing public policy relevance.

2. Background

2.1 Public Disclosures, Social Media, and Social Networking Sites

The rise of social media has both fueled and been fueled by an arguably unprecedented amount of public sharing of personal information. The shared information frequently includes disclosures and revelations of a personal, and sometimes surprisingly candid, nature. In certain cases, personal information is revealed through fully identified profiles.[7] Other times, users provide personal information under pseudonyms that

---

[6] In competitive markets, firms may overinvest and collect an "excessive" amount of information in equilibrium, because of a contrast between social incentives and firms' data-gathering incentives (see Burke, Taylor, and Wagman 2012). In a labor market search model, Wagman (2014) finds that firms may search for negative news about applicants and end up collecting too much information, resulting in applicants inefficiently matched with firms of lower productivity. Seabright and Sen (2014) examine how reductions in the cost of applying for jobs may increase the number of applications to a degree that adversely affects firms.

[7] For instance, an overwhelming majority of Facebook users in a sample of North-American college students surveyed by Acquisti, Gross, and Stutzman (2014) used real first and last names on their profiles.

may still be identified.[8] Some social media users take advantage of privacy settings to manage and restrict their online audiences (Stutzman, Gross, and Acquisti 2012). Others think they do, but actually fail to protect their information.[9]

Employers can access shared personal information in many ways. Some job candidates make their online profiles (on social media or blogging platforms) openly accessible to strangers.[10] Others are more selective, but sensitive information such as religious affiliation, sexual orientation, or family status may still be indirectly inferable from seemingly more mundane data.[11] Finally, some employers engage in social engineering (such as using friends of friends connections to view a candidate's profile), or even ask for the candidates' passwords to access their profiles.[12] Although the legal consensus seems to suggest that seeking information online about job candidates (or employees) may not violate U.S. law, such searches do raise privacy and legal issues, such as the discriminatory behavior that may result from discovering protected information (Sanders 2011; Sprague 2011; Sprague 2012).

Above-cited surveys suggest that U.S. employers have been using social media to screen prospective job candidates. However, no prior controlled experiment has measured the frequency of firms' usage of online profiles in hiring decisions, and how profile information actually affects those decisions. The study that comes closest to our experiment is Manant, Pajak, and Soulié (2015), who investigate the role of social media in the French labor market. Other than differences in the country of study (USA vs. France) and sample size (about 4,000 versus about 800 employers), Manant et al (2015) focus on a different research question from our study: this manuscript focuses on testing the joint hypothesis that firms search online for information about job applicants, and change their hiring activities based on the personal information they find; Manant et al (2015) focus on investigating the impact of different search costs for finding candidates' information via social media. In addition, Bohnert and Ross (2010) use a survey-based experiment to investigate how the content of social networking profiles can influence evaluations of job candidates. Garg and Telang (2012) analyze how job applicants can use LinkedIn to find connections that

---

[8] Numerous examples exist of techniques through which seemingly pseudonymous online profiles can be re-identified across a variety of platforms and scenarios. See, for instance, Narayanan and Shmatikov (2009).

[9] Consider the gap between stated and actual privacy settings of online social network users reported by Acquisti and Gross (2006). Similar results have been subsequently found by Madejski, Johnson, and Bellovin (2012).

[10] For instance, using data collected for this study (see Section 4.2), we estimate that, in 2011, 42% of all profiles members of the Facebook network of a major North-American college shared "likes" publicly (where a like could be an interest, book, movie, music, or a TV program). Data reported in Acquisti, Gross, and Stutzman (2014) also shows that a majority of Facebook members in a North American city used facial images in their primary profile photos (which are public by default). Facebook data analyzed by Johnson, Egelman, and Bellovin (2012) indicates that around 54% of Facebook members made available to strangers at least some of their profile information.

[11] For instance, Jernigan and Mistree (2009) show that a Facebook member's sexual orientation may be inferable from knowledge of her friends' network.

[12] Cases in which employers asked job candidates to provide passwords to their online profiles have been reported in the media. See Stern, Joanna. 2012. "Demanding Facebook Passwords May Break Law, Say Senators." *ABC News*. Last modified March 26. Accessed February 26, 2016. http://abcnews.go.com/Technology/facebook-passwords-employers-schools-demand-access-facebook-senators/story?id=16005565#.T7-fVMWQMQo.

may lead to a job offer. Kluemper, Rosen, and Mossholder (2012) assess the relationship between the content of a person's Facebook profile and her future job performance.

2.2 Discrimination and Résumé Studies

Following Bertrand and Mullainathan (2004), experiments using written applications to real employers have found evidence of discrimination against people with various traits. Pertinent prior literature has found discrimination against job candidates with Muslim rather than Swedish names in Sweden (Carlsson and Rooth 2007), candidates who openly signal Muslim beliefs (relative to other religious beliefs) on their résumés in New England (Wright et al 2013), and candidates who explicitly or implicitly signal same-sex sexual orientation (Weichselbaumer 2003; Ahmed and Hammarstedt 2009; Tilcsik 2011),[13] but no evidence of discrimination against Muslims (and little systematic discrimination by caste in new economy sectors) in India (Banerjee et al 2009).

One crucial difference between existing résumé studies and our approach is that we focus on employers' volitional choices to search online for candidates' information. A second important difference is that we focus on candidates revealing information in personal online social profiles, rather than volunteering personal traits in a professional context. This approach facilitates the investigation of discrimination based on protected information that candidates do not frequently provide in their résumés or during interviews.

3. Specification

Prior literature on job market search strategies has highlighted employers' use of formal or informal information networks (Rees 1966), and their reliance on intensive or extensive searches (Barron, Bishop, and Dunkelberg 1985). A central privacy concern in online job market search is that Internet information exchanges may bundle job relevant information and personal information to an extent not seen in traditional labor market practices. This raises the question we investigate in this manuscript: Is there evidence that employers search for and discriminate on the basis of publicly posted yet personal information? We test for the existence of hiring discrimination that stems jointly from two, possibly dependent, actions: first, each employer decides whether to search online for a candidate's information; and, second, each employer who finds the candidate's profile is unknowingly treated to the experimental manipulation and chooses whether or not to interview the candidate.

It is important to note the difference between our question and more frequently asked questions in the economics literature on discrimination: Does discrimination exist, to what extent, and in what form? That literature has developed sophisticated methods for testing whether or not discrimination is actually present and for empirically separating taste-based from statistical discrimination (Mobius and Rosenblat

---

[13] See also Drydakis (2009) and Hebl et al (2002).

2006; Neumark 2012; Ewens, Tomlin, and Wang 2014). In this paper, we do not attempt to separate bias from search, let alone separate taste-based discrimination from statistical discrimination or other forms of bias. We instead focus on the more basic question of whether or not social media profiles can impact hiring via the combined effects of online search and bias.

Using a between-subjects design, we test for an effect of random assignment to treatment conditions on callbacks:

$$Callback_i = ß_0 + ß_1 * A_i + x_i' \gamma + z_i' \delta + \varepsilon_i \qquad (1)$$

where A_i is an indicator of random assignment of employer i to either the Muslim condition ($A_i = 1$) compared to the Christian condition ($A_i = 0$) or to the gay condition ($A_i = 1$) compared to the straight condition ($A_i = 0$), $ß_0$ and $ß_1$ are unknown parameters, x_i and z_i are vectors of, respectively, observed and unobserved regressors capturing employer i's traits, γ and δ are vectors of unknown parameters, and ε_i is an error term. *Callback_i* equals one if employer i contacts the candidate for an interview and zero if there is no response or an explicit rejection. Assuming successful randomization of A_i, the estimate of $ß_1$ will be an unbiased estimate of the effect of manipulated online profiles on callbacks. It is analogous to an intent-to-treat effect, but we refer to it as an assignment effect to emphasize that the main goal of our design is testing the joint effect of (i) choosing to search (and thus unknowingly self-selecting into our experiment's treatment) and (ii) being treated, whereas researchers in the intent-to-treat literature are primarily interested in isolating the treatment effects.

We do not expect search rates to differ across condition assignments, because employers are treated by our manipulated profiles after they choose to search. However, it is likely that both search and discrimination probabilities will vary with employer characteristics, leading to heterogeneous assignment effects.[14] We test for heterogeneous assignment effects in our regression analysis by including interactions between the random assignment and regressors that prior literature predicts may interact significantly with our manipulations.

3.1 Dependent and Independent Variables
Our experimental design manipulates personal information that is unlikely to be volunteered by candidates in résumés and that may be risky under federal or state law for employers to inquire about during interviews,

---

[14] For instance, screening and hiring procedures differ significantly across firms, with some firms doing all steps of the hiring process in-house and others outsourcing screening to external firms (including emerging services specializing in online screenings, such as HireRight.com) or mathematical algorithms. Hiring procedures used will influence a given employer's experimental treatment probability. We are agnostic regarding the specific hiring procedures used within a given firm. However, we assume this heterogeneity in screening procedures to be similarly distributed across experimental conditions.

but which may be obtained online. Building on résumé studies by Weichselbaumer (2003), Carlsson and Rooth (2007), and Tilcsik (2011), we focus on sexual orientation and religious affiliation.

The dependent measure is interview invitations that the candidate receives from actual employers, either by email, phone call, or letter. Based on prior evidence, we expect callbacks to be more biased in certain areas than in others. Independent surveys consistently find three variables that are associated with more negative attitudes to both Muslims and gay people. Surveys on attitudes to Muslims and Arabs find more bias against Muslims among respondents who are Republican, older, and who know fewer Muslims (Arab American Institute 2012). Separately, surveys on attitudes to gay people find more bias against gay people among respondents who are Republican, older, and who know fewer gay people (Pew Research Center 2013). In communities with a high prevalence of types of people who show bias in these surveys, employers themselves may be more likely to share the attitudes expressed by community members, or they may face incentives to hire candidates whose traits match those of existing employees and community members. Thus, we test for stronger bias in firms located in states and counties that have: (i) a higher median age than the U.S. median age, (ii) a lower fraction of the population that is Muslim than the U.S. fraction of Muslims (in the Muslim-Christian manipulation), or a lower fraction of households that are male-partner households than the U.S. fraction of male-partner households (in the gay-straight manipulation), and (iii) a high fraction of voters in the 2012 Presidential election who voted for the Republican candidate (see Section 5.2.2).

In addition, we control for a host of other variables used in comparable résumé studies, particularly Bertrand and Mullainathan (2004) and Tilcsik (2011), including employer characteristics (such as industry or number of employees), job characteristics (such as job-specific requirements), and regional characteristics based on location of the job (for instance, variations in state and local policies to protect gay people). Finally, we include controls that are specific to our experimental design's online dimension, such as a measure of Facebook penetration across states (a proxy for the likelihood that an individual in a particular state may be inclined to search online for the candidate's social media profile).

3.2 Interpreting a Null Result

Given that our design tests jointly for online searching and for discrimination on the basis of manipulated online profiles, a null effect in the field experiment may signal low levels of either of these behaviors. In addition, the experiment could fail to reject the null hypothesis in a number of additional ways.

First, employers may search the candidates online but fail to find our profiles. Our design addresses this concern by using name-selection criteria that produce high-ranked search results, and by checking, for the duration of the experiment and across a variety of platforms, that candidate name searches produce the desired results (see Section 4.2). Another possibility is that the manipulations may fail to signal the traits we intended to manipulate. To address this possibility, prior to conducting the field experiment, we use an

online pilot to test whether we successfully manipulated beliefs about the candidates' personal traits. A further concern is that the candidates may be either under- or over-qualified, creating floor or ceiling effects in callback rates. We address this possibility by designing the résumés and professional backgrounds of the candidates in concert with human resources professionals, and testing their quality during the online pilot. An additional possibility is that employers may not pursue the candidates because of suspicions that they are fictional. Consequently, we design résumés and profiles using information that real individuals (including members of job search sites and members of the same social network the candidates were purported to belong to) had also made publicly available on their profiles or on job search sites. Furthermore, we choose privacy and visibility settings that were common among the profiles of real members. The online pilot tests the success of our efforts at creating realistic profiles by asking questions aimed at discerning whether online subjects had doubts about the veracity of the candidates (see Section 5.1). Finally, a null effect may be consistent with scenarios in which employers may search, but at a later stage in the hiring process (for instance, after interviews with the candidates). Our design would not capture this behavior, and would not be able to distinguish this scenario from one with little or no discrimination.

4. Design

We implemented a between-subjects design with four treatment conditions consisting of social media profiles for a single job candidate. We manipulated religious affiliation (a Christian versus a Muslim male) and sexual orientation (a gay versus a straight male).[15] Thus, the experimental conditions represent a range of traits that include federally protected information (religious affiliation)[16] and information protected only in certain states (sexual orientation).[17] While these traits enjoy different levels of protection under U.S. law, job candidates can signal this information online both in an explicit manner (e.g., self-descriptions on a profile indicating one's sexual orientation) and an implicit one (e.g., a primary profile photo suggesting the individual's religious affiliation).

For each of the four candidates, we designed: (i) a résumé; (ii) a profile on LinkedIn (a popular online professional network commonly used by human resource professionals and job seekers, henceforth referred to as "PN," as in "Professional Network"); and (iii) a profile on Facebook (a popular social networking site commonly used for socializing and communicating, henceforth referred to as "SN," as in

---

[15] Both the Christian and Muslim profiles are presented as straight. Neither the gay or straight profile provides information about the candidate's religious affiliation.

[16] Under Title VII of the Civil Rights Act of 1964, companies with 15 or more employees may not inquire about a candidate's religious beliefs.

[17] In several, but not all, U.S. states, employers cannot ask job candidates questions regarding their sexual orientations. See U.S. Equal Employment Opportunity Commission. "Laws Enforced by EEOC." Accessed February 26, 2016. http://www.eeoc.gov/laws/statutes/.

"Social Network").[18] We designed candidates' résumés and PN profiles to show identical information across conditions, except for the candidates' names. In contrast, the candidates' SN profiles manipulated information about the candidates' personal traits using various profile fields, including those that some employers admit to using during the hiring process in semi-structured interviews (Llama et al 2012). Each applicant's name corresponds to an experimental condition; the names link together a candidate's résumé, PN profile, and SN profile.[19] Because a single name links all of these materials, we can submit just one application to each employer in our sample.[20] Submitting multiple candidates' applications to the same employer would increase the risk of an employer detecting social media profiles which have identical photos and generic information, yet come from candidates with different names and condition-specific information.

4.1 Design Approach and Priorities

We designed profiles that (i) were realistic and representative of the population of SN users with the traits we were manipulating, and (ii) held constant direct signals of productivity outside of beliefs stemming from the trait itself. As described in the rest of this section, we populated the résumés and online profiles using existing information posted online by actual SN members demographically similar to the candidates, and by individuals who listed their résumés on job searching sites. The SN profiles are the vehicles for the manipulations, and responses to these profiles comprise our study's core. The remainder of this section discusses our design in more detail.

4.2 Candidates' Information

Names. We designed first and last names representing a U.S. male for each of the conditions used in the experiments. We chose first names common among U.S. males in the same age group as our candidates, and assigned identical first names to candidates in matching pairwise conditions (the gay and straight candidates were both named Mike, and the Muslim and Christian candidates were both named Adam).[21] We then designed last names by altering letters of existing but similarly low-frequency U.S. last names. The last names were chosen to have the same number of syllables, identical lexical stress, and similar sounds across matching pairwise conditions (see Appendix Table A1 for more detail). A potential concern

---

[18] In addition, and for completeness, we created candidates' accounts on Google+ and Google Sites. These profiles were identical across conditions and did not include professional and candidate-specific information.

[19] We took down the various online materials we had created only after the experiment's completion, which included candidates' online profiles and the sites of companies that had been created as part of the design.

[20] Many traditional résumé studies use multiple applications to each employer, but for an exception see Ahmed, Andersson, and Hammarstedt (2013), who sent one application per employer.

[21] Our candidates were purported to be born in 1982. Mike was the most popular name for male newborns in the 1980s according to U.S. Social Security data (see Social Security Administration. 2015. "Top Names of the 1980s." Last modified February. Accessed February 26, 2016. http://www.ssa.gov/oact/babynames/decades/names1980s.html) and was assigned to the gay/straight conditions. Adam was the 22nd most popular name, and was assigned to the Christian/Muslim conditions, it being both a Christian and a Muslim name.

is that Christian and Muslim names often differ. However, roughly 35% of the about 2.35 million American Muslims are, like our candidate, US-born, and many have Anglo-Saxon names.[22]

We iteratively tested several combinations of first and last names until we found a subset that satisfied three criteria. Criterion 1 was that the exact first and last name combination would be unique on SN and PN: no other profile with that name should exist on the network. Criterion 2 was that SN and PN profiles designed under each name would appear among those names' top search results conducted with the most popular search engines (Google, Bing, and Yahoo), as well as searches conducted from within the SN and the PN social networks. We continuously monitored the fulfillment of these criteria for the experiment's duration.[23]

Criterion 3 was the most critical one: names alone should not elicit statistically significant differences in perceptions of the manipulated traits. We conducted two checks of this criterion. First, we tested that names, by themselves, did not influence perceptions of the traits manipulated in the profiles. We recruited 496 subjects from Amazon Mechanical Turk (or, MTurk) – a popular platform for behavioral experiments[24] – and presented each with one first and last name combination, randomly chosen from the list of possible names. Each subject then responded to a questionnaire on associations between that name and the traits manipulated in field experiment. We selected names for the field experiment that did not elicit statistically significant differences in perceptions of traits (see Appendix Table A5), and then randomly assigned them to candidates. The second check of Criterion 3 combined names, résumés, and professional online profiles, and checked that names and professional information did not elicit differential propensity to invite a candidate to an interview (see Section 5.1).

Email addresses and telephone numbers. For inclusion in the résumé and cover letters sent to companies, we designed email addresses using a consistent format across the profiles and registered a telephone number for the applicants. We used automated message recordings for the number's voice messaging service. This message was identical across conditions and recorded with the voice messaging service's standard greeting. We also used the same postal address for all candidates, corresponding to a residential area of a mid-size North-American city.

Résumés. The résumés contained information depicting a professional and currently employed candidate. Each résumé contained the candidate's contact information, educational background, and work experience,

---

[22] See Pew Research Center. 2007. "Muslim Americans: Middle Class and Mostly Mainstream." Last modified July 2011. Accessed February 26, 2016. http://pewresearch.org/pubs/483/muslim-americans.

[23] The experiment was not started until search engines had indexed the social media profiles. After the experiment began, and on a weekly basis throughout the experiment's duration, we monitored the search results for the candidates' names. We conducted searches through proxy servers so that they appeared as coming from foreign IP addresses, which enabled us to avoid contaminating the actual employer search data that we hoped to gather from services such as Google Keyword Tool trends (see Section 5.2.1 and Appendix B).

[24] See Appendix A.

as well as technical skills, certifications, and activities. The résumés were held constant across conditions, except for the names of the applicants and their email addresses.[25] Hence, professional and educational backgrounds did not vary across experimental conditions, thus holding constant the candidates' job market competitiveness. The information included in the résumé was modeled after résumés found on websites such as Monster.com and Career.com for job seekers demographically similar (in terms of age and educational background) to the candidates. The résumé represented a candidate with a bachelor's degree in computer science and a master's degree in information systems.[26] Creating a single SN and PN profile per candidate constrained the types of jobs to which we could apply. Hence, all résumés needed to be consistent with the constant information provided in the candidates' social media profiles (for instance, all candidates' profiles exhibited the same master's degree in information systems). Therefore, the openings (and respective résumés) tended to be technical, managerial, or analytic in nature. Two human resource recruiters vetted the résumés for realism, professionalism, and competitiveness before they were tested in the online pilot. A design objective (also tested during the online pilot) was to create a sufficiently competitive candidate for stimulating the level of interest necessary to generate an online search, but not so competitive as to outweigh any potential effect arising from an employer's perusal of the candidate's online profile.

The résumés did not include links to the candidates' personal profiles or references to the personal traits we manipulated on the SN. The experimental design relied entirely on the possibility that employers would autonomously decide to seek information about applicants online, searching for their names either on popular search engines or directly on popular social networks.

PN ("Professional Network") profiles. We designed PN profiles for each name, maintaining identical profile information across conditions. The content of the profiles (see Appendix Table A3) reflected the information provided in the résumés. To increase realism, we also designed additional PN profiles for other fictional individuals and connected them to the candidates' profiles, so that they would become "contacts" for our candidates. We took great care to avoid any possible linkage between the actual candidates' profiles.[27] The number of PN connections was identical across all the profiles and chosen to be as close as

---

[25] Research assistants blind to the experimental conditions were allowed to add up to two technical skills (for instance, Java) or certifications (for instance, CISSP certification) to a résumé if the job description required those skills or certifications as pre-conditions. This fine-tuning always occurred *before* a candidate's name was randomly added to the résumé, and therefore before the candidate was assigned to a job application.

[26] We prepared 10 versions of the same résumés, focusing on slightly different sets of expertise: web development, software development, quality assurance, project or product management, medical/healthcare information, information systems, information security, business intelligence, business development, and analytics. A sample résumé is presented in Appendix Table A2.

[27] We also created websites, email accounts, and PN presence for some of the companies reported in the candidates' résumés, as well as other workers at those companies and potential letter-writers, in order to have a complete, believable background should anyone search. The candidates were also registered as alumni from the institution that granted their degrees. As noted, we took down these materials after the completion of the experiment.

possible to the median number of connections of a US-based profile on the PN professional networking site at the time of the experiment. [28]

SN ("Social Network") profiles. SN profiles served as the principal manipulation vehicle in the field experiment. We paired each candidate with a SN profile created under the candidate's name. As detailed in the rest of this section, we used a number of strategies to create realistic and balanced profiles, and capture the phenomenon of online social networking in an ecologically valid way: we populated our candidates' profiles with data extracted from actual profiles of social network users who were demographically similar to the candidates; we made sure that the overall amount of public self-disclosure in our profiles would be equivalent to the amount of self-disclosure in actual social network profiles of users demographically similar to our candidates; we designed profiles to include a combination of information that changed based on the experimental condition and information that was held constant across all conditions; and finally, we posted the same amount of manipulated and constant information for each candidate (therefore, the Christian and Muslim candidates' profiles, and the gay and straight candidates' profiles, presented equivalent information about the strength of their religion, or sexual orientation).

First, we downloaded the public profiles of 15,065 members of the same Facebook college network from which the candidates purportedly graduated. The vast majority of those profiles belonged to individuals with a similar age range (their 20s), current location (a North-American mid-size city), and educational background as our candidates. As further detailed below, among that set of profiles we then focused on those with the same gender (male) as our candidates, and who self-disclosed traits matching the ones manipulated in the experimental conditions (that is, Christian, Muslim, straight, or gay). Then, we designed several components of each SN profile based on the information we had mined from real online profiles: personal information (such as current location, location of birth, education, and employment history); closed-ended text fields (such as interests, activities, books, movies, or TV shows); open-ended text fields (such as those providing a summary self-description of the individual and his favorite quotations); friends list; primary profile photo and secondary photos; and background image. Based on literature that emerged during the design phase of the experiments,[29] however, we decided not to make all this information publicly available and therefore visible to a human resources professional. To avoid potential confounding

---

[28] We created 41 connected accounts. In 2012, 40% of LinkedIn users had between 51 and 200 connections. See Breitbarth, Wayne. 2012. "LinkedIn Infographic: Want To Know What Others Are Doing?" *Power Formula*. Last modified March 4. Accessed February 26, 2016. http://www.powerformula.net/2347/linkedin-infographic-want-to-know-what-others-are-doing/. Our ability to create a large number of additional connections for our candidates faced technical barriers, such as the need to provide mobile phone numbers to the social network provider in order to validate newly created profiles.

[29] Johnson, Egelman, and Bellovin (2012) report that only 14.2% of the users of a popular social network whose profile information they mined had a public wall (that is, visible status updates and comments). That noted, we reiterate that the profiles used in the experiment contained equivalent amounts of information across conditions, so that any difference in callback rates could not be merely imputed to a higher amount of disclosure by one type of candidate over another. Furthermore, we note that several fields (such as name, gender, primary photo, profile photo, and networks) are mandatorily public on the network we used for the experiment.

effects from over-disclosure, we refrained from showing certain fields that only a minority of users of the network actually publicly show, such as posts, status updates, and friends list.

Appendix Table A3 presents the resulting information included in each of the four SN profiles. By design, some of these profile components (the non-condition-specific traits) were constant across the different profiles and thus across the different experimental conditions. Other components (the condition specific traits) were manipulated across conditions. We chose a mix of condition-specific and non-condition-specific traits that replicated the combination and balance we observed in the actual profiles we mined and analyzed (as described above). The manipulation of fields representing different types of personal information was meant to increase the realism and ecological validity of our experimental design – which we tested in our pilot experiment. We discuss each type of component in the following sub-sections.

Information kept constant across SN profiles. By design, we made the candidates' primary profile photo, secondary photos, current location, hometown, age, education, employment history, and friend list constant across conditions. Basic personal information (such as the SN member's current location and city of birth, educational background, employment history, hometown, and age) was made publicly visible on the profile and kept constant across the conditions. That information was made consistent with the content of the résumé. The candidate was presented as U.S. born and English speaking.[30]

The photo on the candidates' social media profiles was of a Caucasian male with dark hair and brown eyes. We picked the photo after recruiting non-professional models on Craiglist and asking them to submit a portfolio of personal photos. About forty subjects (recruited via MTurk) rated all models who submitted their photos along two 7-point Likert scales, indicating their perceived attractiveness and professionalism. We selected one male model that received median perceived attractiveness and professionalism ratings.

We designed numerous "friends" profiles to connect to the candidates' profiles. Again, the number and identities of friends were identical across conditions.[31] The set of friends showcased a mix of names, backgrounds, and profile visibility settings. We set the friends list to "private" – that is, not visible to an employer – because of extant research suggesting that the overwhelming majority of members of the network do not publicly disclose their friend lists.[32]

---

[30] The Muslim candidate's social network profile also presented him as speaking both Arabic.

[31] Also in the case of the SN, we attempted to create a number of friends that would match the median number of connections of U.S.-based SN profiles at the time the experiment started (around 100). However, we were only able to create 49 friends, due to technical barriers such as (as noted earlier) the need to provide unique mobile phone numbers to create new profiles on the network. As noted, we prevented cross-linkages of the candidates via their common networks of friends by appropriately selecting privacy and visibility settings for the social network profiles.

[32] See Johnson, Egelman, and Bellovin (2012). Although we made the list of friends private, we did not delete the friend profiles because 1) some of those friends' comments appeared on the candidates' profile photos, adding realism to the profiles; 2) the presence of a network of friends decreased the probability that the candidates profiles could be identified as fake and deactivated by the network (none of the candidates profiles was deactivated during the duration of the experiment).

Information manipulated across SN profiles. Some fields in the profiles were manipulated across conditions: close-ended text fields (e.g., interests), open-ended text fields (e.g., quotations), and the background image. Furthermore, we manipulated the candidates' sexual orientation by filling out the field "interested in" (either male interested in females or interested in males), and we manipulated religious affiliation through the "religion" field (either Christian or Muslim, with no specific denomination).

We abided by a number of principles in designing the manipulated information. First, we constructed "baseline" profile data, which was constant across all conditions, including information such as specific interests and activities statistically common among the SN profiles we had previously mined ("baseline information"). We then augmented the profiles with additional data (such as additional interests or activities) specific to the traits that we wanted to manipulate ("treatment information"). Both "baseline" and "treatment" information were extracted from real, existing SN profiles. The profiles therefore represented realistic pastiches of *actual* information retrieved, combined, and remixed from existing social media accounts.[33] We took care to avoid confounding the trait we were manipulating with signals of worker quality beyond signals inherent in the manipulated trait itself. Furthermore, the different candidates' profiles were designed to disclose the same amount of personal information, including the same quantity and types of data revealing their sexual orientation or religious affiliation, so as not to create profiles with unbalanced condition-specific disclosures. For similar reasons, we took great care to ensure that the information revealed by the candidates would not be construed as over-sharing, or as a forced caricature of what a Muslim (or Christian, or gay, or straight) profile should look like. The combination of baseline and treatment information of true existing profiles was one of the strategies we used to meet this goal; using only information existing in other SN profiles, and making certain fields publicly inaccessible, were two others. The online pilot experiment, with its open-ended questions about the profiles, tested whether our goal was attained (see Section 5.1. and Appendix A).

Close-ended text fields such as interests and activities were extracted using a combination of manual and statistical analysis from real SN profiles of people who shared demographic characteristics with the candidate (for the "baseline information") and who exhibited the same characteristics we manipulated (for "treatment information"). For example, when text involved countable objects such as one's favorite books or movies, we calculated the most popular items listed by the overall population of SN profiles and used that as the "baseline information" for the candidate. Then, we repeated the operation, focusing on the subset of the public disclosures of over 15,000 SN profiles that displayed the same traits we manipulated, in order to create the "treatment information." For instance, we constructed a portion of the profile text

---

[33] Technical limitations inherent in designing an experiment via online social networks precluded us from the possibility of randomizing the entire content of the candidates' profiles by creating thousands of different profiles for each candidate. On the other hand, our approach achieves ecological realism by relying on existing profile data.

indicating activities and interests for the Christian male profile using statistical analysis of the entire sample of over 15,000 profiles that also included non-Christian ones; the remaining portion of the profile text concerning his activities and interests was constructed using statistical analysis of the sub-sample of profiles of Christian males at his university. If our sample did not provide enough information (for instance, movies) for individuals with a given trait, we complemented the statistical analysis with a manual analysis of the same SN profiles. We also extracted open-ended fields (such as personal self-descriptions, personal quotations, or non-textual fields such as profile background images) through manual analysis of existing profiles, since the extreme variety of styles and contents across open-ended texts made a statistical approach unfeasible.

5. Results

5.1 Online Pilot Experiment

Before submitting résumés to actual job openings at U.S. firms, we conducted a pilot experiment, consisting of a randomly assigned questionnaire with online participants and a between-subjects design. The online pilot was designed to test whether the treatment conditions successfully manipulate relevant beliefs, such as the religious affiliation or sexual orientations of the candidates. Furthermore, its open-ended questions were designed to test the perceived realism of the profiles. Finally, the pilot experiment was designed to test whether, in absence of the manipulated personal profiles (SN), the candidates' names, résumés, and professional profiles elicited differential propensities to invite the candidate for an interview. Details of the design, analysis, and findings of the online pilot are presented in Appendix A and are summarized here.

We recruited over 1,750 U.S. residents as participants using Amazon Mechanical Turk. Participants were presented with the candidates' social profiles and résumés prepared for the actual field experiment. Participants included individuals with previous hiring experience – something we exploited in the analysis of the results. Participants were randomly assigned to one of four conditions (gay, straight, Christian, or Muslim candidate), and were provided links to one candidate's résumé, PN profile, and SN profile. The survey instrument had four elements: (i) introduction and randomized manipulation; (ii) measurement of perceived employability; (iii) measurement of beliefs for the purpose of manipulation check; and (iv) open-ended questions and demographic characteristics. Participants were asked to evaluate the candidate. The main dependent variables were hypothetical willingness to call the candidate for an interview and perceptions of the candidate's suitability for the job.

The online pilot found, first, that the treatment conditions successfully manipulated relevant beliefs, such as the religious affiliation and sexual orientation of the candidates. Second, open-ended questions checked for perceived realism of the profiles; they provided no evidence of doubt, among the participants, that the candidates were real. Third, the online pilot tested whether, in absence of links to the manipulated

personal profiles (SN), the candidates' names, résumés, and professional profiles elicited differential propensities to invite the candidate for an interview; we found no evidence that the candidates' names and professional materials elicited different responses in the absence of the manipulated online social media profiles. Finally, responses of hypothetical hiring behavior and judgments of employability provided evidence to complement the findings of our field experiment. Consistent with our field findings (presented further below in this section), manipulated profiles in the online pilot elicited no bias against the gay candidate, relative to the straight candidate. However, among subjects with hiring experience, we found highly significant bias against the Muslim candidate relative to the Christian candidate, especially among those who self-identify as Republican and Christian.

5.2 Field Experiment

The field experiment consisted of a between-subjects design in which each employer was randomly assigned to receive one job application from either the Muslim, Christian, gay, or straight candidate. The job application period lasted from early 2013 through the summer of 2013. Samples used in résumé studies (see Section 2.2) have ranged from a few hundred to a few thousand employers. We aimed for the higher end of the spectrum, ultimately applying to every U.S. job opening that fit our pre-defined criteria of a reasonably suitable position for our candidates. This amounted to 4,183 U.S. job applications (or roughly 1,045 employers per experimental condition). Ten applications failed to meet criteria we had defined in our procedures for acceptable and complete applications, leaving us with 4,173 usable applications.

We extracted job openings from Indeed.com, an online job search site that aggregates jobs from several other sites.[34] We selected positions that fit the candidates' backgrounds – namely, positions that required either a graduate degree or some years of work experience. For each position, we sent the most appropriate version of our résumé. (Recall from Section 4.2 that we designed different versions of the résumé to fit ten different job types covering a combination of technical, managerial, and analytic positions.)

We defined several criteria that jobs and companies had to pass for us to apply to them. We focused on private sector firms. Primarily, the job had to be related to the candidates' background and level of experience, although we also included (and controlled for) positions for which the candidates could be considered slightly over- or under-qualified. In addition, we carefully avoided sending two applications to the same company, or to companies that were likely to share HR resources such as databases of applicants (for instance, parent companies of firms to which we already applied).[35] We also excluded staffing

---

[34] Appendix Table A4 lists the search terms used to find different type of jobs.
[35] As noted, we could not tolerate multiple candidates' applications being sent to the same firm, as employers may then have been able to find applications and social media profiles with identical photos and overlapping information from different candidates. We also took great care to avoid cross-linkage of the candidates by the same employer through other means. For instance, employers could not navigate from one of the candidates to the other using their (shared) friends, and profile images were subtly modified

companies, companies located in the same geographic region as the candidates' reported current location, and companies with 15 or fewer employees (to limit the costs imposed on them by the process of searching and vetting fictional job candidates).[36] All applications (résumés and cover letters) were submitted online, either by email or through employer-provided web forms.

We recorded the city and state listed on job postings, and, when possible, we recorded the city and state where the job would be located. When job location was not provided, we used the location of the company's headquarters. We obtained this measure for all observations and employed it for our state level analysis. From the company name, we were able find the street address of the company headquarters for all but a few hundred observations. We used ArcGIS[37] to match this street address to its county. We then merged our data with county level data from the American Community Survey[38] based on the county where the company headquarters is located.

5.2.1 Search Trends

While the focus of the experimental design was capturing employers' callback behavior, tracking and identifying employers' searches and visits to the manipulated profiles would be desirable. Without access to the social network' proprietary traffic data,[39] however, it is not possible to observe how many employers visited the candidates' profiles and thus got treated to the experimental manipulation. Hence, the experimental design does not allow us to capture directly the number of times the candidates' profiles were searched for or perused. In fact, in principle, *all* employers who received the candidates' applications could have searched for the candidates directly via the SN (thus remaining undetectable to us). In practice, that appears unlikely. It is possible, however, to obtain rough estimates of frequency of employers' online searches of job candidates using various indirect approaches. We summarize two approaches and their associated findings in this section, and present detailed assumptions, calculations, and additional results in Appendix B.

We used a combination of two data sources to estimate the frequency with which employers searched for the candidates online. One data source consisted of Google AdWords "Keyword Tool" statistics. These publicly accessible statistics capture the number of times a certain term is searched on Google from various locations. We used this tool to estimate the number of times the exact names of the candidates were searched from U.S. IP addresses. The second data source consisted of statistics provided by the PN network (LinkedIn) via so-called "Premium" accounts. If a user subscribes to a Premium account,

---

with random-noise (invisible to the human eye) in order to avoid being listed as similar images by Google's reverse image search functions.

[36] Small firms were excluded from the experiment per agreement with our institution's IRB. Note that these firms may not be constrained by EEOC laws, so their exclusion actually may yield more conservative estimates.

[37] See website for ArcGIS; https://www.arcgis.com/features/.

[38] See website for the United States Census Bureau, American Community Survey; https://www.census.gov/programs-surveys/acs/.

[39] Facebook, the SN we used in the experiment, does not allow its members to track or count visitors to their profiles.

that user will be able to get information such as the count of visits to its PN profiles, and in some cases the actual identity of the visitors. We subscribed to Premium accounts for each of our candidates' profiles in order to track visits to those profiles.

Each of these sources of data is imperfect. Google Keyword Tool statistics provide aggregate monthly means for searches of a given term, rather than raw data. Furthermore, if the mean is higher than zero but below 10 searches per month, no exact count is provided. Similarly, "Premium" PN accounts do not actually track all visits to that account (for instance, in our tests, visits from subjects that had not logged onto the network went undetected and uncounted; visitors who only viewed the summary profile of the candidate, rather than his full profile, also were not detected or counted; in addition, certain Premium accounts may allow "invisible" views of other LinkedIn profiles). Nevertheless, considered together, these data sources do offer a rough estimate of the proportions of employers searching the candidates online.

We tracked Google Keyword Tool statistics and LinkedIn data over a period of several months. Based on the documented searches we detected on Google and on the PN, we can estimate a minimum lower threshold of employers who searched for the profiles at 10.33%, and the likely proportion of employers who searched at 28.82% (see Appendix B for details). These estimates appear consistent with the results of a 2013 CareerBuilder survey of 2,291 hiring managers and HR managers, according to which 24% claimed to "occasionally" use social media sites to search for candidates' information; 8% answered "frequently;" and only 5% answered "always."[40]

The search rates highlighted in this section are sample-wide, aggregate estimates, because (as noted) in most cases we could not identify the specific employers visiting the profiles. The sub-sample of cases where we could directly capture the identity of an employer searching for our profiles is small (N = 121), but, unsurprisingly, represents employers who are strongly interested in our candidates. The overall callback rate for the four candidates in this sub-sample was 39.67%. The callback rates for the straight and gay candidates were 31.25% and 40.00%, respectively, and were not significantly different from each other. However, the callback rates for the Christian and Muslim candidates were 54.84% and 32.14%, respectively and were significantly different at the ten-percent level (using a chi-squared test). This is consistent with the main results presented in the remainder of this section, where we do find evidence of callback bias in favor of the Christian candidate compared to the Muslim candidate among certain types of employers, but no evidence of bias against the gay candidate compared to the straight candidate.

5.2.2 Employer Callbacks

---

[40] See Erwin, Michael. 2013. "More Employers Finding Reasons Not to Hire Candidates on Social Media, Finds CareerBuilder Survey," *CareerBuilder*. Last modified June 27. Accessed February 26, 2016. http://www.careerbuilder.com/share /aboutus/pressreleasesdetail.aspx?sd=6%2f26%2f2013&id=pr766&ed=12%2f31%2f2013.

Our primary dependent variable is a callback dummy that equals one if the candidate was contacted for an interview and zero if the candidate received a rejection, no response, or a scripted request for more information. Of the 4,173 observations from both manipulations, 11.20% were callbacks, 15.86% were rejections, 69.29% had no response, and 3.32% were scripted requests for more information.

Timing of Callbacks. Figure A1 (in the Appendix) shows the geographical distribution of applications across the United States. Figure A2 (in the Appendix) provides a distribution of the timing of responses we received across all candidates, where a response is an interview invitation, a rejection, or a request for more information. This figure does not include the 70% of employers who did not respond to the application. Most responses came within the first 14 days.[41] Roughly 77% of responses arrived within one month from the application, and roughly 90% within two months. The distribution of timing of responses follows the same patterns across the four candidates (unreported).

Summary of Callback Rates. Overall callback rates are presented in Table 1, together with other summary statistics for various variables captured in the experiment. The callback rates for the straight and gay candidates are nearly identical: 10.63% and 10.69%, respectively. The callback rates for the Christian and Muslim candidates are 12.64% and 10.96%, respectively. This difference is not statistically significant, but masks strong heterogeneity in hiring bias: as discussed in Section 3.1, in surveys, certain traits predict stronger bias against Muslims or gay people (Republican Party identification, older age, and knowing fewer Muslims or gay people). We thus test for differences in bias according to geographical differences in the prevalence of these traits in the population.

Although we cannot identify the individual making the callback decision, and thus cannot observe his or her traits, following the literature we use demographic measures in the firm's geographical area. Specifically, we indicate states and counties with a: (i) higher median age than the U.S. median age, (ii) lower fraction of the population that is Muslim (or gay) than the U.S. fraction, and (iii) a high fraction of voters who voted for Mitt Romney in the 2012 Presidential election. Our measure of political party support in (iii) follows Tilcsik (2011), who used Presidential election data in his audit study, and the Gallup Organization, which regularly produces lists of the ten most Republican and ten most Democratic states based on survey data. Voting results are available at the county level and are based on real behavior. We thus indicate, respectively, the ten states with the highest fractions of 2012 Romney voters and the ten states with the lowest. We refer to the remaining states as politically mixed.[42] We indicate Republican and

---

[41] We did not apply to job openings that had been posted for over 30 days. It is possible that some job openings may have been filled or were otherwise inactive. We expect these cases to be uncorrelated with the condition assignments.

[42] We use Presidential election data from Leip (2013), which is compiled, where possible, from final election results from official sources (e.g., certified election results posted by state election boards). We found that election data from news organizations is often not updated to include final election results. In Section 5.2.3 we show that the results are robust with respect to using the Gallup Organization's classification of Republican and Democratic states for the year of 2012, according to daily tracking of self-reported political party identification.

Democratic counties as those with the same cutoff values for Romney vote fractions that were used to indicate the Republican and Democratic states. Columns (1)-(3) of Table 2 present summary statistics on callbacks for the Muslim and Christian candidates by age, fraction Muslim, and political party support. The most striking and robust result is that callback bias against the Muslim candidate relative to the Christian candidate is substantially higher in Republican states and counties (Panel A) than in Democratic states or counties (Panel B). In Republican states, 17.31% of applications by the Christian candidate received interview invitations compared to only 2.27% for the Muslim candidate. Though striking, this difference is consistent with our estimates of search rates.[43] At the county level, the callback rates for the Christian and Muslim candidates in Republican counties are 22.58% and 6.25%, respectively. Column (3) presents test statistics for differences between the callback rates in the Muslim and Christian conditions: these biases are significant at the five- and one-percent levels at the state and county levels, respectively. In contrast, Panel B shows no significant callback biases in Democratic states or counties. In Democratic states, the Christian and Muslim candidates have callback rates of 11.61% and 11.74%, respectively. In Democratic counties, the Christian and Muslim candidates have callback rates of 12.13% and 12.37%, respectively. These results are robust with respect to multiple specifications (see *Analysis* section, below), robustness checks (Section 5.2.3), and using alternative definitions of Republican states (Section 5.2.3).

Table 2 also shows, for the Muslim-Christian manipulation, that employers in older counties are significantly less likely to call back the Muslim candidate compared to the Christian candidate. The age interaction is less robust than the political party support interaction: it is not present at the state level (we show in Section 5.4 that theories of underlying mechanisms can fully explain the age interaction, so it may be a statistical artifact). Finally, we find no significant differences in bias between high and low Muslim counties or states. In the gay-straight manipulation, we find no differences in bias across any of the geographical areas. For instance, in Republican states, 15.38% of the straight candidate applications and 14.29% of the gay candidate applications received callbacks (see Table 2, Columns (4) and (5)). In

---

[43] Consider an employer search rate of 30% (see our search rate estimates in Section 5.2.1). Suppose there are three types of employers: Type 1 rejects the application and does not search the candidate online; Type 2 considers the candidate further and searches online; and Type 3 considers the candidate further but does not search online. Suppose that 65% of the employers are Type 1, 30% are Type 2, and 5% are Type 3. In this example, the overall search rate is just the prevalence of Type 1 employers, i.e., 30%. If we assume that the callback rate by Type 3 employers (who do not search) equals the average of the callback rates for the two candidates by Type 2 employers (who do search), then we have two equations with two unknowns: P1*0+P2*M+P3(M+C)/2=.02 and P1*0+P2C+P3(M+C)/2=.17, where P1, P2, and P3 are the probabilities of employers being types 1, 2, and 3; and M and C are the callback rates for the Muslim and Christian candidates among employers who search. Sample-wide callback rates of 2% for the Muslim candidate and 17% for the Christian candidate can be explained by callback rates among employers who search of 2% for the Muslim candidate and 52% for the Christian candidate, and a callback rate when employers consider the candidate but do not search of 27%. It is worth noting that if, in Section 5.2.1, we underestimated actual employers' search rates, or if employers in Republican states have a higher than average search rate, then this example implies a smaller callback bias for employers who search.

Democratic states the callback percentages for the straight and gay candidates were 11.24 and 11.71, respectively.[44]

We also test whether the biases found in Table 2 stem from both the Muslim and Christian conditions or only from one. We compare callbacks in each of these conditions to callbacks in the pooled gay-straight conditions – using the latter as benchmarks. Consider columns (6), (7), and (8) of Table 2: both the Muslim and the Christian conditions contribute to the biases shown in Panel A (i.e., in the bias-prevalent areas). Columns (7) and (8), respectively, test for differences in the number of callbacks in (i) the Muslim condition to the pooled gay/straight conditions and (ii) the Christian condition compared to the pooled gay/straight condition. The first row shows that in Republican states, the Muslim callback rate of 2.27% is significantly lower than the pooled gay/straight callback rate of 14.77%. The remaining rows of Panel A show that in Republican counties, older counties, and states and counties with a lower fraction of Muslims, the callback rate in the Christian condition is significantly higher than the callback rate for the pooled gay/straight condition. Thus, both the Muslim and the Christian conditions play significant roles in generating callback bias in areas with a high prevalence of types who show bias in surveys.

The result that both the Muslim and the Christian conditions contribute significantly to the bias in Republican areas is not sensitive to which sample is chosen as the comparison sample. We find similar results when we separate the pooled gay/straight candidate sample into just the gay candidate sample or the straight candidate sample. Regardless of whether we use the sample with the gay candidate or the sample with the straight candidate, we find that the callback rate for the Muslim candidate in Republican states is significantly lower and the callback rate for the Christian candidate in Republican counties is significantly higher.

In the next sub-section, we present regression analysis using both the state-level and county-level measures of Republican, politically mixed, and Democratic areas. We present the state-level analysis first. However, the state-level analysis has two drawbacks. First, taking all four conditions together, there are 1,745 employers in 10 Democratic states or districts (including Washington DC), 2,244 in 31 mixed states, and just 184 in 10 Republican states. Second, we cannot fully control for state-level covariates with state-level definitions of political areas. Thus, we address these issues by also using county-level regressions with state-fixed effects. The county-level measures, in addition to allowing state controls (and thus focusing on within-state variation), have a larger sample size in the Republican areas. The sample sizes in the Democratic and Republican counties are 2,475 and 228, respectively. Finally, in Section 5.2.3, we test for

---

[44] We also test for any effect of the application channel (email application versus web application) or the date of the job opening. We do not find any evidence of impact of those variables on callback rates.

robustness with respect to changes in estimators and types of standard errors, and with respect to additional specifications and measures of political areas.

Regression Analysis. Table 3 presents OLS regressions for the Muslim-Christian manipulation, where the dependent variable is callbacks.[45] Column (1) presents a baseline regression that includes a dummy for assignment to the Muslim condition, dummies for the politically mixed states and Democratic states (Republican states are the omitted category), a dummy for states where the median age is below the U.S. median age, a dummy for states where the fraction Muslim is above the U.S. fraction Muslim, and interactions between the Muslim candidate assignment dummy and each of the other independent variables.

The estimated callback rate for the Christian candidate in the default category – namely, Republican states where the median age is high and the fraction Muslim is low – is 0.176. In the default category, the estimated callback rate in the Muslim condition is significantly lower compared to the Christian condition, by -0.133, yielding an estimated callback fraction for the Muslim candidate of 0.043. In column (1) the interaction between the Muslim assignment effect and the Democratic states dummy is positive, 0.145, and significant at the five-percent level. The interaction between politically mixed states and the Muslim assignment dummy is also positive, 0.120, and significant at the ten-percent level. Column (2) is the same as column (1), except that it includes additional state-level control variables (see Table 3 notes). The Muslim assignment effect and the interaction effects are similar to column (1). Column (3) is the same as Column (2) except that is adds firm- and job-level controls to the regressions (see Table 3 notes).[46] These results are, again, consistent with the results of Columns (1) and (2).

Column (4) is identical to Column (1), except that it uses county-level measures with state fixed-effects.[47] The effect of the Muslim assignment dummy in the default category – namely, Republican counties where the median age is higher and the fraction Muslim lower – is negative and significant at the one-percent level. The interactions between the Muslim assignment dummy and the politically mixed and Democratic counties are positive and significant at the ten- and five-percent levels, respectively. The interaction between a low median age and the Muslim assignment dummy is positive and significant at the five-percent level. Column (5) is the same as Column (4), except that it includes additional county-level control variables analogous to the state-level controls that were included in Column (2) (see Table 3 notes). Column (6) is the same as Column (5), except that it adds the firm level controls that were included in Column (3). The results in Columns (5) and (6) are consistent with the Column (4) results just summarized.

---

[45] Table 3 presents robust standard errors with no clustering. The results are robust with respect to different specifications of the model, as discussed in Section 5.2.3 (see Appendix Table A16 for probit estimates and Appendix Table A17 for OLS estimates with clustered standard errors).

[46] Fewer than 500 employees is a widely used cutoff standard for identifying small firms in U.S. Small Business Administration programs.

[47] The county-level sample is smaller because we were unable to identify counties for all of our observations. We lose a few additional observations in the county-level regressions because of missing county-level demographic data.

Furthermore, looking across Columns (4)-(6), one can see, in rows 2 and 3, that the callback rate for the Christian candidate is slightly lower in politically mixed and Democratic counties than it is in Republican counties.

Table 4 presents regression results for our sexual orientation manipulation. In contrast to the results of the religion manipulation, there are no significant interactions between the treatment assignment and interacted geographical areas.

In short, the regression analysis highlights results that are robust to a variety of specifications. The results in Tables 2, 3, and 4 are consistent with prior results that were suggestive of similar patterns but not definitive. First, they are consistent with our online pilot experiment (see Section 5.1): the online pilot showed significant bias against the Muslim candidate relative to the Christian candidates among subjects with hiring experience, but no bias against the gay candidate relative to the straight candidate (see Appendix A). Second, in a Republican area, Gift and Gift (2014) find evidence (significant at the one-percent level in a one-tailed test) of discrimination against job candidates who signal Democratic Party affiliation on résumés compared to those who signal Republican Party affiliation. In a Democratic area, there is a directionally consistent bias against job candidates who signal Republican Party affiliation, but it is not significant at the five-percent level in a one-tailed test. This is consistent with our finding of stronger discrimination in Republican areas. In Section 5.2.4, we consider a theory that explains why bias is stronger in some places, and derive testable predictions about where to expect stronger bias.

The lack of evidence of bias against the gay candidate differs from the results reported by Tilcsik (2011), who found a significant effect of sexual orientation on callbacks. Rather, our finding is consistent with more recent results presented by Bailey, Wallace, and Wright (2013), who do not find evidence of discrimination against gay men or lesbians. Evolving attitudes toward gay people may explain the differences between our results and Tilcsik's (2011), whose experiment was conducted in 2005. According to our analysis of General Social Survey data,[48] acceptance of gay marriage increased among self-identified Republican, Independent, and Democratic respondents from 2006 to 2012.

5.2.3 Robustness checks

Despite having a binary dependent variable, we used OLS in our analysis above because of concerns about interaction effects in probit regressions (Ai and Norton 2003). Appendix Table A16 is identical to Table 3, except that it uses probit, and the results are robust to this change. Appendix Table A17 is identical to Table

---

[48] Smith, Tom W., Peter V. Marsden, Michael Hout, and Jibum Kim. "General Social Surveys, 1972-2012." [machine-readable data file]. Principal Investigator, Tom W. Smith; Co-Principal Investigators, Peter V. Marsden and Michael Hout, NORC ed. Chicago: National Opinion Research Center, producer, 2005; Storrs, CT: The Roper Center for Public Opinion Research, University of Connecticut, distributor. 1 data file (55,087 logical records) and 1 codebook (3,610 pp). Results are available upon request.

3, except that it presents clustered standard errors (at the state or county level, as appropriate) instead of unclustered errors. The results are robust to this change as well.[49]

Table 3 presents county-level regressions with state fixed effects to address potential weaknesses associated with state-level data. Appendix Table A19 presents additional *state*-level robustness checks. We have so far defined the Republican, Democratic, and politically mixed states according to fractions of voters who voted for Romney in the 2012 U.S. Presidential election. We refer to this categorization of states as the *Romney vote list*. As a robustness check, we use two additional definitions of Republican, Democratic, and politically-mixed states. The first is the ten most Republican and Democratic states in 2012, according to Gallup survey measures of self-reported political party identification (Saad 2013). We categorize the remaining states as politically mixed. We refer to this categorization as the *Gallup list*.[50] The second additional definition is a combination of the Romney vote list and the Gallup list (hereafter the *Combined list)*. The Combined list of Republican states includes any state on either the Romney vote list or the Gallup list of Republican states. Similarly, the combined list of Democratic states includes any state on either the Romney vote list or Gallup list of Democratic states. There is no overlap between the combined Republican and Democratic lists. The remaining states are politically mixed (see Appendix Table A18 for a list of states according to these three definitions).

Appendix Table A19 presents regression results using the Gallup list (Column 1) and Combined list (Column 2) definitions of Republican, Democratic, and politically mixed states. Columns (1) and (2) show that the results are robust to the changed definitions. The rest of the table uses the Combined list for additional robustness checks. Columns (3) and (4) show the smallest and largest interactions with the Democratic states that we could find by deleting each of the 50 states (including Washington, DC), one regression at a time. Column (3) removes California from the sample and column (4) removes Idaho. The results remain robust with respect to dropping one state at a time. Even the smallest interaction with the Democratic states dummy is significant at the five-percent level. Columns (5) and (6) are the same as column (2) except that they add, respectively, the additional controls included in columns (2) and (3) of Table 3. Again, the results are robust with respect to the inclusion of these additional controls.

5.2.4 Explaining bias

This section considers two possible explanations for the finding that hiring bias in favor of the Christian candidate and against the Muslim candidate in Republican and older areas is both significant and significantly stronger than in non-Republican and younger areas.

---

[49] We conducted these same robustness checks for Table 4 (the sexual orientation manipulation). There were no noteworthy changes.

[50] Gallup produced this list of states from a sample of 321,233 respondents surveyed by Gallup Daily tracking over the course of the year at the rate of 1,000 respondents per day.

The first is a taste for discrimination based on ethnic group loyalty. Employers may seek candidates who will be a closer ethnic match with existing employees or themselves. Thus, we might expect stronger bias in areas with a higher prevalence of traits similar to our Christian candidate and different from our Muslim candidate.[51] We do not find strong support for this explanation. Our analysis (see Table 3) already failed to find a significant interaction between the fraction of the population that is Muslim and hiring bias. In Appendix Table A20, we test for additional interactions between the Muslim condition dummy and the fraction of the population that is non-white and the fraction that is not Evangelical Christian. We do not find robust interactions. This does not imply that a taste for group loyalty is unimportant in other forms or settings. However, as an explanation for our main findings, we turn to a second explanation.

A second possibility is that hiring bias in our experiment may be an act of group-beneficial cooperation among employers who benefit from hiring their own types. Graham, Haidt and Nosek (2009) find evidence of stronger group cooperation among political conservatives. They interpret this as political conservatives and liberals having different sets of "moral foundations." A different approach models discrimination as "optimal parochialism" – namely, group-beneficial cooperation in relatively small social or professional networks, such as "close knit residential neighborhoods and ethnically linked businesses" (Bowles and Gintis 2004, p. 1). Unlike taste-based group loyalty, cooperative parochial discrimination is group-beneficial. In particular, when contracts are incomplete, group-beneficial cooperation in markets can increase efficiency by increasing the flow of information. Insiders to the network share a common set of traits. Individuals who possess some but not all of these traits may be excluded. This may help explain why our tests for ethnic group loyalty above failed: a single trait may be insufficient to define an in-group. This particular model predicts that networks are stable for an intermediate group size, and that in equilibrium there is no incentive to migrate into or out of a network. More broadly, the model follows insights from Ostrom (1998, 2000), who emphasized the importance of repeated interactions in relatively small communities with clearly delimited group boundaries that define who is included and who is excluded from the group, to solve common pool resource problems. In closely related work, Chong (2000) has argued that socioeconomic mobility, among other characteristics of communities, can break down norms of trust and reciprocity. Habyarimana et al (2009) have distinguished norm driven group-beneficial cooperation from taste-based discrimination in African communities, showing that subjects failed to cooperate with co-ethnics when anonymous but cooperated with co-ethnics when their identity was known. Thus, a common theme in this literature is that communities that are large or mobile may be less likely to form exclusionary groups. We check the plausibility of these arguments in our data by investigating the effects of (i) community size, measured with a dummy variable for counties in large metropolitan areas with over 1

---

[51] See Brewer and Miller (1996), Luttmer (2001), and Chen and Li (2009) on group loyalty. For a paper broadly related to both explanations considered in this section, see Akerlof and Kranton (2000).

million people and (ii) geographical mobility, measured with Census data on the fraction of households that have not changed homes in the last year.

In Appendix Table A21, columns (1) and (2) are identical to the county-level regressions that were presented in columns (5) and (6) of Table 3, except that (i) the demeaned fraction of the county population that moved in the last year is added to the regression (a metro area dummy was already included in the regressions of Table 3) and (ii) the variables that are interacted with the Muslim condition dummy are now a dummy for large metro counties with over 1 million people, and the demeaned fraction of the county population that moved in the last year. These columns show significant effects in the expected direction of the interactions between the Muslim condition dummy and the metro area dummy and the fraction of movers. Columns (3) and (4) are identical to columns (1) and (2) except that they add back in the interactions from our main analysis of Table 3. These columns show that when all of these interactions are included, the interaction with the metro area dummy and the fraction of movers is no longer significant, and, more important for our purposes, the interaction with low-age counties is no longer significant, and the political area dummies are only significant in one column. In short, a large community size and high levels of geographical mobility are associated with lower hiring bias in Columns (1) and (2) of Table A21, and adding these interactions to our main specifications from Table 3 may partially account for the stronger bias in Republican and older counties.

Bowles and Gintis (2004) argue that information technology may assist cooperation in parochial networks. The interpretation of the results we presented in this section suggests that the phenomenon we reported may actually have been the result of an interaction between "traditional" offline parochial networks and "new" online social networks: online social networks are not necessarily parochial, but they may have nonetheless assisted cooperation in employers' existing parochial networks.

5.3 Limitations

A number of limitations are inherent to the experiment's design. First, a null effect in the Gay-Straight manipulation may not necessarily be interpreted as an absence of discrimination, because a number of potential explanations may exist. Second, heterogeneity in the assignment effect according to firm characteristics cannot be separated into firms' differences in search rates and differences in the treatment effect. Third, the results stem from candidates with certain characteristics, applying to certain types of jobs, with certain types of employers; they may not apply to other types of candidates, positions, or organizations. Fourth, we measure employers' traits with state or county level data; we do not capture the traits of the actual employers or human resource decision makers. Finally, discrimination based on information posted online may stem partly from the signal that the posted information sends about a candidate's judgment in choosing to post that information. However, the disclosures we investigated in this paper took place in personal profiles that employers must actively seek, rather than in résumés in which candidates intentionally

reveal potentially sensitive information to employers, and the consequences of the fact that people use online social networks to disclose this information in this way, is part of what we sought to understand. Our experimental approach – namely, designing profiles to replicate information found in real social media profiles of individuals demographically comparable to our candidates (see Section 4.2) – increased the realism and ecological validity of our experimental design, and decreased the chances that employers may interpret the profiles as "over-sharing" beyond the levels that we observe in real online social network profiles.

6. Conclusion

We set out to answer the question: do hiring decision makers seek online information which should not be used in the hiring process and are they affected by what they find? We focused on the effects of religious affiliation and sexual orientation. We used randomized experiments to answer those questions, capturing discrimination both in behaviors of real employers and attitudes of survey subjects. Our results were broadly consistent across the online pilot and the field experiment. They suggest that while hiring discrimination via Internet searches and social media may not be widespread for the companies and jobs we considered in this experiment, revealing certain traits online may have a significant effect on the hiring behavior of self-selected employers who search for the candidates online.

The findings provide further evidence of a phenomenon increasingly studied in the information systems literature: the impact of online IT services on offline outcomes and behaviors. They also highlight the novel tensions arising between regulatory interventions designed in a non-digital world (that attempt to protect personal information), and technological innovations that bypass those protections (by making said information otherwise available).

This work suggests various directions for future research. There is a tension between a potentially efficiency-enhancing effect of online information on labor market search and a potential for that same information to increase labor market discrimination. To the extent that online markets facilitate not just matching but also discrimination, people in disadvantaged categories may face a dilemma: if they censor their online activity, they may be able to protect their privacy, but a limited online presence might, by itself, signal that there is something to hide, or that something is missing. Imagine a charismatic Muslim candidate with both a strong social network in his religious community and many professional contacts. If he were to reduce his online presence to hide his friendships in his religious community, at least three problems arise. The first is a "lemons" effect: people may assume that those with a restricted online presence belong to less preferred categories. The second is what we may term an "online network exclusion effect": by protecting his privacy, the candidate is deprived of the opportunity to present professionally relevant positive traits about himself. The third effect is a bundled traits effect: candidates may have some traits that employers

find less desirable and some traits that are more desirable, and online mechanisms sometimes bundle together information about undesirable and desirable traits. For the charismatic Muslim candidate described above, information about personal traits that are potentially less desirable to employers may be bundled with information about traits that are more desirable to employers, including the strength of his professionally relevant social networks. This presents additional challenges for disadvantaged candidates who may wish to protect information about traits that are disadvantageous in the labor market while promoting their positively viewed traits. This dilemma may motivate design of mechanisms that facilitate communication only about job-relevant characteristics. Thus, few multi-faceted questions for future research arise, which could be asked descriptively using the methodology in this paper: does the amount of online self-disclosure affect labor market success? (For instance, in an age of self-disclosure, is the absence of social media presence a neutral, positive, or even negative signal? Do employers pay attention to the size of a candidate's personal social network?) Will market mechanisms develop to solve some of these problems? It may also be interesting to ask these questions normatively: how should online self-disclosure affect labor market success? How should markets be structured to address these problems?

Some employers openly admit to using social media in the hiring process (Llama et al 2012). At the same time, legal professionals warn firms about the risks of doing so. Both the EEOC and state legislators have become interested in the role of the Internet and social media in the job search process. However, companies like HireRight.com and Rapleaf.com have offered related services – arguably shielding firms from the risks associated with doing the search themselves. In sum, the rise of the Internet and social media has increased the amount of personal and professional information available to employers about a job candidate. This manuscript offers a first insight as well as a methodological path towards understanding the extent to which this new information channel may reduce labor market frictions, or lead to more discrimination.

REFERENCES

Acquisti, Alessandro, Curtis R. Taylor, and Liad Wagman. 2015. "The Economics of Privacy." *Journal of Economic Literature,* forthcoming.

Acquisti, Alessandro, and Ralph Gross. 2006. "Imagined Communities: Awareness, Information Sharing, and Privacy on the Facebook." *Proceedings of Privacy Enhancing Technologies* 4258: 36-58.

Acquisti, Alessandro, Ralph Gross, and Fred Stutzman. 2014. "Face Recognition and Privacy in the Age of Augmented Reality." *Journal of Privacy and Confidentiality* 6 (2): 1.

Ahmed, Ali M., and Mats Hammarstedt. 2009. "Detecting Discrimination against Homosexuals: Evidence from a Field Experiment on the Internet." *Economica* 76 (303): 588-597.

Ahmed, Ali M., Lina Andersson, and Mats Hammarstedt. 2013. "Are Gay Men and Lesbians Discriminated Against in the Hiring Process?" *Southern Economic Journal* 79 (3): 565-585.

Ai, Chunrong, and Edward Norton. 2003. "Interaction Terms in Logit and Probit Models." *Economics Letters* 80 (1): 123-129.

Akerlof, George, and Rachel Kranton. 2000. "The Economics of Identity." *Quarterly Journal of Economics* 115 (3): 715-753.

Arab American Institute. 2012. "The American Divide: How We View Arabs and Muslims." *Arab American Institute,* August 23. http://b.3cdn.net/aai/4502fc68043380af12_oum6b1i7z.pdf.

Aslund, Olof, and Oskar N. Skans. 2012. "Do Anonymous Job Application Procedures Level the Playing Field?" *Industrial & Labor Relations Review* 65 (1): 82-107.

Bailey, John, Michael Wallace, and Bradley Wright. 2013. "Are Gay Men and Lesbians Discriminated Against When Applying for Jobs? A Four-City, Internet-Based Field Experiment." *Journal of Homosexuality* 60 (6): 873-894.

Banerjee, Abhijit, Marianne Bertrand, Saugato Datta, and Sendhil Mullainathan. 2009. "Labor Market Discrimination in Delhi: Evidence from a Field Experiment." *Journal of Comparative Economics* 37 (1): 14-27.

Barron, John M., John Bishop, and William C. Dunkelberg. 1985. "Employer Search: The Interviewing and Hiring of New Employees." *Review of Economics and Statistics* 67 (1): 43-52.

Bertrand, Marianne, and Sendhil Mullainathan. 2004. "Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination." *American Economic Review* 94 (4): 991-1013.

Bohnert, Daniel, and William H. Ross. 2010. "The Influence of Social Networking Web Sites on the Evaluation of Job Candidates." *Cyberpsychology, Behavior, and Social Networking* 13 (3): 341-347.

Bowles, Samuel, and Herbert Gintis. 2004. "Persistent Parochialism: Trust and Exclusion in Ethnic Networks." *Journal of Economic Behavior and Organization* 55 (1): 1-23.

Brewer, Marilynn B., and Norman Miller. 1996. *Intergroup Relations*. Pacific Grove: Brooks/Cole Publishing Company.

Burke, Jeremy M., Curtis R. Taylor, and Liad Wagman. 2012. "Information Acquisition in Competitive Markets: An Application to the US Mortgage Market." *American Economic Journal: Microeconomics* 4 (4): 65-106.

Carlsson, Magnus, and Dan-Olof Rooth. 2007. "Evidence of ethnic discrimination in the Swedish labor market using experimental data." *Labour Economics* 14 (4): 716-729.

Chan, Jason, and Anindya Ghose. 2014. "Internet's dirty secret: assessing the impact of online intermediaries on HIV transmission." *MIS Quarterly* 38 (4): 955-976.

Chen, Yan, and Sherry X. Li. 2009. "Group Identity and Social Preferences." *American Economic Review*, 99(1): 431-57.

Chong, Dennis. 2000. *Rational Lives: Norms and Values in Politics and Society*. Chicago: The University of Chicago Press.

Drydakis, Nick. 2009. "Sexual Orientation Discrimination in the Labour Market." *Labour Economics* 16 (4): 364-372.

Ewens, Michael, Bryan Tomlin, and Liang C. Wang. 2014. "Statistical Discrimination or Prejudice? A Large Sample Field Experiment." *Review of Economics and Statistics* 96 (1): 119-134.

Garg, Rajiv, and Rahul Telang. 2012. "To Be or Not To Be LinkedIn: Job Search Using Online Social Networks." NBER Summer Institute Working Paper.

Gift, Thomas, and Karen Gift. 2014. "Does Politics Influence Hiring? Evidence from a Randomized Experiment." *Political Behavior*: 1-23.

Goldfarb, Avi, and Catherine E. Tucker. 2011. "Privacy Regulation and Online Advertising." *Management Science* 57 (1): 57-71.

Goldin, Claudia, and Cecilia Rouse. 2000. "Orchestrating Impartiality: The Impact of 'Blind' Auditions on Female Musicians." *American Economic Review* 90 (4): 715-741.

Graham, Jesse, Jonathan Haidt, and Brian A. Nosek. 2009. "Liberals and Conservatives Rely on Different Sets of Moral Foundations." *Journal of Personality and Social Psychology* 96 (5): 1029-1046.

Habyarimana, James, Macartan Humphreys, Daniel N. Posner, Jeremy M. Weinstein. 2009. *Coethnicity: Diversity and the Dilemmas of Collective Action*. New York: Russel Sage Foundation.

Harpe, Lisa D., "Social Networks and Employment Law: Are you Putting your Organization at Risk?," Peopleclick Authoria Research Institute, last modified 2009, accessed June 23, 2014, http://www.peoplefluent.com/sites/default/files/resources/files/Social_Networks_Employment_Law.pdf.

Hebl, Michelle R., Jessica B. Foster, Laura M. Mannix, and John F. Dovidio. 2002. "Formal and Interpersonal Discrimination: A Field Study of Bias Toward Homosexual Applicants." *Personality and Social Psychology Bulletin* 28 (6): 815-825.

Jernigan, Carter, and Behram F.T. Mistree. 2009. "Gaydar: Facebook Friendships Expose Sexual Orientation." *First Monday* 14 (10).

Johnson, Maritza, Serge Egelman, and Steven M. Bellovin. 2012. "Facebook and Privacy: It's Complicated." *Proceedings of the Eighth Symposium on Usable Privacy and Security* 9.

Kluemper, Donald H., Peter A. Rosen, and Kevin W. Mossholder. 2012. "Social Networking Websites, Personality Ratings, and the Organizational Context: More Than Meets the Eye?" *Journal of Applied Social Psychology* 42 (5): 1143-1172.

Kroft, Kory, and Devin G. Pope. 2014. "Does Online Search Crowd Out Traditional Search and Improve Matching Efficiency? Evidence from Craigslist." *Journal of Labor Economics* 32 (2): 259-303.

Leip, David. 2013. "Election 2012." Atlas of U.S. Presidential Elections. http://uselectionatlas.org/2012.php.

Llama, Vanessa A. De La, Isabel Trueba, Carola Voges, Claudia Barreto, and David J. Park. 2012. "At Face (book) value: Uses of Facebook in Hiring Processes and the Role of Identity in Social Networks." *International Journal of Work Innovation* 1 (1): 114-136.

Luttmer, Erzo F. P. 2001. "Group Loyalty and the Taste for Redistribution." *Journal of Political Economy* 109 (3): 500-528.

Madejski, Michelle, Maritza Johnson, and Steven M. Bellovin. 2012. "A Study of Privacy Setting Errors in an Online Social Network." *Proceedings of the 4th IEEE International Workshop on Security and Social Networking*: 340-345.

Manant, Matthieu, Serge Pajak, and Nicolas Soulié. 2015. "Do Recruiters 'Like' It? Online Social Networks and Privacy in Hiring: A Pseudo-Randomized Experiment." Université Paris-Sud Working Paper.

Miller, Amalia R., and Catherine Tucker. 2009. "Privacy Protection and Technology Diffusion: The Case of Electronic Medical Records." *Management Science* 55 (7): 1077-1093.

Mobius, Markus M., and Tanya S. Rosenblat. 2006. "Why Beauty Matters." *American Economic Review* 96 (1): 222-235.

Narayanan, Arvind, and Vitaly Shmatikov. 2009. "De-Anonymizing Social Networks." *Proceedings of 30th IEEE Symposium on Security and Privacy*: 173-187.

Neumark, David. 2012. "Detecting Discrimination in Audit and Correspondence Studies." *Journal of Human Resources* 47 (4): 1128-1157.

Ostrom, Elinor. 1998. "A Behavioral Approach to the Rational Choice Theory of Collective Action: Presidential Address, American Political Science Association, 1997." *American Political Science Review* 92 (1): 1-22.

Ostrom, Elinor. 2000. "Collective action and the evolution of social norms." *Journal of Economic Perspectives* 14 (3):137–158.

Pew Research Center. 2012. "Two-Thirds of Democrats Now Support Gay Marriage." *Pew Research Center U.S. Politics & Policy,* July 31. http://www.pewforum.org/2012/07/31/2012-opinions-on-for-gay-marriage-unchanged-after-obamas-announcement/.

Pew Research Center. 2013. "In Gay Marriage Debate, Both Supporters and Opponents See Legal Recognition as 'Inevitable.'" *Pew Research Center,* June 6. http://www.people-press.org/2013/06/06/in-gay-marriage-debate-both-supporters-and-opponents-see-legal-recognition-as-inevitable/.

Posner, Richard A. 1981. "The Economics of Privacy." *American Economic Review* 71 (2): 405-409.

Rees, Albert. 1966. "Information networks in labor markets." *American Economic Review* 56 (1/2): 559-566.

Saad, Lydia. 2013. "In the U.S., Blue States Outnumber Red States, 20 to 12." *Gallup Politics*, January 30. http://www.gallup.com/poll/160175/blue-states-outnumber-red-states.aspx (accessed February 9, 2016).

Sanders, Sherry D. 2011. "Privacy is Dead: The Birth of Social Media Background Checks." *Southern University Law Review* 39: 243-264.

Seabright, Paul, and Ananya Sen. 2014. "Statistical Externalities and the Labour Market in the Digital Age." IDEI Working Paper.

Spence, Michael. 1973. "Job Market Signaling." Quarterly Journal of Economics 87 (3): 355-374.

Sprague, Robert. 2011. "Invasion of the Social Networks: Blurring the Line Between Personal Life and the Employment Relationship." *University of Louisville Law Review* 50 (1): 1-34.

Sprague, Robert. 2012. "Facebook Meets the NLRB: Employee Online Communications and Unfair Labor Practices." *University of Pennsylvania Journal of Business Law* 14 (4): 957-1011.

Stigler, George J. 1962. "Information in the labor market." *The Journal of Political Economy* 70 (5): 94-105.

Stigler, George J. 1980. "An Introduction to Privacy in Economics and Politics." *Journal of Legal Studies* *9*: 623-644.

Stutzman, Fred, Ralph Gross, and Alessandro Acquisti. 2012. "Silent Listeners: The Evolution of Privacy and Disclosure on Facebook." *Journal of Privacy and Confidentiality* 4 (2): 7-41.

Taylor, Curtis R., and Huseyin Yildirim. 2011. "Subjective Performance and the Value of Blind Evaluation." *The Review of Economic Studies* 78 (2): 762-794.

Tilcsik, Andras. 2011. "Pride and Prejudice: Employment Discrimination Against Openly Gay Men in the United States." *American Journal of Sociology* 117 (2): 586-626.

Vigdor, Jacob L. 2008. "Measuring Immigrant Assimilation in the United States." *Civic Report* 53. Retrieved from http://www.manhattan-institute.org/pdf/cr_53.pdf (accessed February 9, 2016).

Wagman, Liad. 2014. "Good News or Bad News? Information Acquisition and Applicant Screening in Competitive Labor Markets." Northwestern University Working Paper.

Weichselbaumer, Doris. 2003. "Sexual Orientation Discrimination in Hiring." *Labour Economics* 10 (6): 629–642.

Wright, Bradley R.E., Michael Wallace, John Bailey, and Allen Hyde. 2013. "Religious Affiliation and Hiring Discrimination in New England: A Field Experiment." *Research in Social Stratification and Mobility* 34: 111-126.

Table 1. Summary Statistics.

	N	Mean	S.D.	Min.	Max.
Panel A: Christian Candidate					
Callback (1=interview invitation)	1060	0.1264	0.3325	0.0000	1.0000
Fraction Romney vote, state	1060	0.4371	0.0994	0.0728	0.7279
Fraction Romney vote, county	1019	0.3676	0.1504	0.0728	0.8832
Median age, state	1060	37.20	2.20	29.90	43.50
Median age, county	1020	36.51	3.08	24.50	46.60
Fraction Muslim, state	1060	0.0093	0.0081	0.0005	0.0280
Fraction Muslim, county	1020	0.0120	0.0133	0.0000	0.1061
Panel B: Muslim Candidate					
Callback (1=interview invitation)	1022	0.1096	0.3125	0.0000	1.0000
Fraction Romney vote, state	1022	0.4415	0.0913	0.0728	0.7279
Fraction Romney vote, county	998	0.3747	0.1476	0.0728	0.8832
Median age, state	1022	37.28	2.12	29.90	43.50
Median age, county	990	36.57	3.19	24.60	53.10
Fraction Muslim, state	1022	0.0090	0.0078	0.0005	0.0280
Fraction Muslim, county	1000	0.0119	0.0129	0.0000	0.0878
Panel C: Straight Candidate					
Callback (1=interview invitation)	1025	0.1063	0.2969	0.0000	1.0000
Fraction Romney vote, state	1025	0.4376	0.0936	0.0728	0.7279
Fraction Romney vote, county	991	0.3600	0.1454	0.0728	0.8832
Median age, state	1025	37.31	2.07	29.90	43.50
Median age, county	991	36.65	3.12	24.10	53.10
Fraction male&male HH out of total HH, state	1025	0.0042	0.0025	0.0006	0.0126
Fraction male&male HH out of total HH, county	993	0.0044	0.0039	0.0000	0.0173
Panel D: Gay Candidate					
Callback (1=interview invitation)	1066	0.1069	0.3092	0.0000	1.0000
Fraction Romney vote, state	1066	0.4395	0.0886	0.0728	0.7279
Fraction Romney vote, county	1038	0.3604	0.1452	0.0728	0.8832
Median age, state	1066	37.35	2.11	29.90	43.50
Median age, county	1033	36.57	3.02	24.60	50.40
Fraction male&male HH out of total HH, state	1066	0.0042	0.0020	0.0011	0.0126
Fraction male&male HH out of total HH, county	1039	0.0043	0.0040	0.0000	0.0173

See Table A13 for sources. Additional summary statistics are presented in Tables 2 and A13.

Table 2. Callback rates in all conditions by geographical area.

	(1) Muslim callback rate [N]	(2) Christian callback rate [N]	(3) χ² p-value {Fisher's exact 2-sided} for H₀: (1)=(2)	(4) Gay callback rate [N]	(5) Straight callback rate [N]	(6) Pooled gay and straight callback rate [N]	(7) χ² p-value {Fisher's exact 2-sided} for H₀: (1)=(6)	(8) χ² p-value for H₀: (2)=(6)
Panel A: Callback bias in areas with higher prevalence of types of people who report the strongest bias in survey data								
Republican states	**.0227 [44]**	**.1731 [52]**	**0.016 {0.019}**	.1429 [49]	.1538 [39]	**.1477 [88]**	**0.028 {0.034}**	0.690
Republican counties	**.0625 [64]**	**.2258 [62]**	**0.009 {0.011}**	.0816 [49]	.098 [51]	**.0900 [100]**	0.525 {0.768}	**0.016**
States with median age > U.S. median	.1243 [555]	.1284 [553]	0.839	.1134 [582]	.1165 [558]	.1149 [1140]	0.573	0.422
Counties with median age > U.S. median	**.1154 [364]**	**.1722 [389]**	**0.027**	.1290 [372]	.1156 [372]	**.1223 [744]**	0.739	**0.021**
States with fraction Muslim< U.S. fraction	.1078 [705]	**.1293 [727]**	0.209			**.1033 [1433]**	0.748	**0.070**
States with fraction gay HH< U.S. fraction				.1126 [613]	.1115 [601]			
Counties with fraction Muslim < U.S. fraction	.1056 [568]	**.1353 [584]**	0.123			**.1071 [1130]**	0.927	**0.085**
Counties with fraction gayHH < U.S. fraction				.1047 [774]	.1005 [736]			
Panel B: Areas with lower prevalence of types of people who report the strongest bias in survey data								
Politically Mixed states	.1105 [552]	**.1304 [560]**	0.309	.0960 [573]	.0984 [559]	**.0972 [1132]**	0.395	**0.039**
Politically Mixed counties	.1000 [360]	.1239 [347]	0.313	.0938 [352]	.1078 [334]	.1006 [686]	0.976	0.255
Democratic states	.1174 [426]	.1161 [448]	0.952	.1171 [444]	.1124 [427]	.1148 [871]	0.892	0.946
Democratic counties	.1237 [574]	.1213 [610]	0.901	.1165 [635]	.1102 [608]	.1134 [1243]	0.527	0.619
States with median age <U.S. median	.0921 [467]	.1243 [507]	0.107	.0992 [484]	.0942 [467]	.0967 [951]	0.779	0.104
Counties with median age < U.S. median	.1086 [626]	.1030 [631]	0.746	.0968 [661]	.1018 [619]	.0992 [1280]	0.524	0.795
States with fraction Muslim > U.S. fraction	.1136 [317]	.1201 [333]	0.795			.1140 [658]	0.985	0.776
States with fraction gay HH > U.S. fraction				.0993 [453]	.0991 [424]			
Counties with fraction Muslim > U.S. fraction	.1181 [432]	.1216 [436]	0.874			.1086 [902]	0.610	0.484
Counties with fraction gayHH > U.S. fraction				.1170 [265]	.1284 [257]			

There are no significant differences in callback rates shown in Columns (4) and (5).

Table 3. OLS Regressions in the Christian-Muslim Conditions. Dependent Variable is Callbacks.

	(1)	(2)	(3)	(4)	(5)	(6)
	State-level geographical characteristics			County-level geographical characteristics		
Muslim candidate	-0.133**	-0.130**	-0.113*	-0.197***	-0.222***	-0.199***
	(0.060)	(0.060)	(0.061)	(0.071)	(0.072)	(0.071)
Politically mixed area	-0.040	-0.029	-0.006	-0.109*	-0.121*	-0.085
	(0.056)	(0.058)	(0.058)	(0.062)	(0.065)	(0.064)
Democratic area	-0.056	-0.015	0.008	-0.097	-0.130*	-0.100
	(0.055)	(0.071)	(0.072)	(0.063)	(0.071)	(0.070)
Median age < U.S. median age	-0.005	0.018	0.019	-0.065**	-0.065**	-0.076***
	(0.021)	(0.027)	(0.028)	(0.026)	(0.028)	(0.028)
Fraction Muslim > U.S. fraction Muslim	-0.010	0.008	0.006	0.010	0.007	-0.001
	(0.023)	(0.028)	(0.029)	(0.027)	(0.028)	(0.028)
Muslim*Politically mixed area	0.120*	0.115*	0.098	0.137*	0.163**	0.120*
	(0.062)	(0.062)	(0.063)	(0.071)	(0.072)	(0.071)
Muslim*Democratic area	0.145**	0.139**	0.113*	0.156**	0.179**	0.139**
	(0.062)	(0.062)	(0.063)	(0.070)	(0.071)	(0.069)
Muslim*Low median age	-0.025	-0.023	-0.025	0.065**	0.065**	0.086***
	(0.029)	(0.029)	(0.029)	(0.032)	(0.032)	(0.033)
Muslim*High fraction Muslim	0.008	0.007	0.016	0.004	0.008	0.014
	(0.032)	(0.032)	(0.032)	(0.030)	(0.030)	(0.031)
State fixed effects?	na	na	na	Yes	Yes	Yes
Additional geo controls?	No	Yes	Yes	No	Yes	Yes
Firm-level controls?	No	No	Yes	No	No	Yes
Constant	0.176***	0.143**	0.088	0.259***	0.636***	0.547***
	(0.055)	(0.066)	(0.071)	(0.060)	(0.184)	(0.195)
Observations	2082	2082	2009	2008	2006	1934
R^2	0.004	0.007	0.042	0.030	0.038	0.073

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Dep. Var. =1 if the candidate is contacted for an interview, zero otherwise. Numbers reported are OLS coefficients (robust standard errors in parentheses). Additional state-level geographical controls included in column 2 and 3 are: 2012 unemployment rate, fraction foreign-born, natural log of median income, fraction non-white, fraction college educated or more, fraction evangelical Christian, fraction urban, Facebook penetration, and legal protection from religious discrimination. Additional county-level geographical controls in col. 5 and 6 are: 2012 U.E. rate, fraction foreign-born, natural log of median income, fraction non-white, fraction college educated or more, fraction evangelical Christian, and rural-urban continuum code dummies. Firm-level controls included in columns 3 and 6 are: dummies for women and minority owned, public firm, large firm (500 employees or more), federal contractor, and included application-level characteristics are dummies for entry level position, references required, preferred salary required, master's degree required, one or more years of experience required, multiple fields of experience required, and 9 dummies for field of employment. The continuous state and county variables – namely, 2012 unemployment rate, fraction foreign-born, natural log of median income, fraction non-white, fraction college educated or more, fraction evangelical Christian, fraction urban, and Facebook penetration – are centered on their means. The omitted categories for dummies capturing variables with more than two categories are: Republican states or counties; counties with less than 20,000 people; and jobs in information systems. The remaining variables are binary. Full regression results are available in the Appendix.

Table 4. OLS Regressions in the Gay-Straight Manipulation. Dependent Variable is Callbacks.

	(1)	(2)	(3)	(4)	(5)	(6)
	State-level geographical characteristics			County-level geographical characteristics		
Gay candidate	-0.016	-0.019	-0.038	0.014	0.005	-0.034
	(0.078)	(0.078)	(0.082)	(0.066)	(0.066)	(0.067)
Politically mixed area	-0.059	-0.040	-0.060	0.026	0.000	-0.022
	(0.060)	(0.063)	(0.067)	(0.046)	(0.049)	(0.052)
Democratic area	-0.039	-0.001	-0.017	0.025	-0.018	-0.041
	(0.063)	(0.077)	(0.081)	(0.047)	(0.058)	(0.061)
Median age < U.S. median age	-0.023	-0.017	-0.025	-0.020	-0.020	-0.032
	(0.021)	(0.024)	(0.024)	(0.024)	(0.026)	(0.027)
Fraction gay HH > U.S. fraction gay HH	-0.010	-0.015	-0.008	0.055^{*}	0.051	0.044
	(0.023)	(0.026)	(0.027)	(0.029)	(0.032)	(0.033)
Gay*Politically mixed area	0.013	0.015	0.033	-0.018	-0.011	0.020
	(0.079)	(0.079)	(0.084)	(0.067)	(0.068)	(0.068)
Gay*Democratic area	0.024	0.025	0.027	0.022	0.030	0.061
	(0.083)	(0.083)	(0.086)	(0.067)	(0.067)	(0.068)
Gay*Low median age	0.009	0.009	0.025	-0.021	-0.019	-0.011
	(0.029)	(0.029)	(0.030)	(0.030)	(0.030)	(0.031)
Gay*High fraction gay households	-0.009	-0.009	-0.012	-0.028	-0.028	-0.019
	(0.032)	(0.032)	(0.032)	(0.036)	(0.036)	(0.037)
State fixed effects?	na	na	na	Yes	Yes	Yes
Additional geo controls?	No	Yes	Yes	No	Yes	Yes
Firm-level controls?	No	No	Yes	No	No	Yes
Constant	0.170^{***}	0.140^{**}	0.128^{*}	0.082^{*}	0.243^{*}	0.233
	(0.060)	(0.066)	(0.075)	(0.045)	(0.132)	(0.144)
Observations	2091	2091	2011	2021	2019	1942
R^2	0.003	0.006	0.035	0.025	0.032	0.063

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Numbers reported are OLS coefficients (robust standard errors in parentheses). See notes to Table 2. The dependent variable and all additional controls not shown are the same as in Table 2, with the exception that columns 2 and 3 control for legal protection from sexual orientation discrimination, instead of legal protection from religious discrimination. Full regression results are available in the Appendix.