

NBER WORKING PAPER SERIES

INCREASING THE DEMAND FOR WORKERS WITH A CRIMINAL RECORD

Zoe B. Cullen  
Will S. Dobbie  
Mitchell Hoffman

Working Paper 29947  
<http://www.nber.org/papers/w29947>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
April 2022

We thank Amanda Agan, Paul Goldsmith-Pinkham, Sara Heller, Lawrence Katz, Judd Kessler, Eddie Lazear, Melvin Stephens, Crystal Yang, and numerous seminar participants for helpful comments and suggestions. Dylan Balla-Elliott, Dan Ma, and Alexia Olaizola provided excellent research assistance. We thank the Abdul Latif Jameel Poverty Action Lab (J-PAL), the Michael Lee-Chin Family Institute for Corporate Citizenship, and the Social Sciences and Humanities Research Council (SSHRC) for financial support. The views expressed here are those of the authors alone and do not necessarily reflect those of the institutions or funders involved with this work. The RCT and a pre-analysis plan were pre-registered with the AEA RCT registry under ID AEARCTR-0005200. We are deeply grateful to Wonolo Inc. for their support of this work. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Zoe B. Cullen, Will S. Dobbie, and Mitchell Hoffman. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Increasing the Demand for Workers with a Criminal Record  
Zoe B. Cullen, Will S. Dobbie, and Mitchell Hoffman  
NBER Working Paper No. 29947  
April 2022  
JEL No. C93,J23,J24,M51

**ABSTRACT**

State and local policies increasingly restrict employers' access to criminal records, but without addressing the underlying reasons that employers may conduct criminal background checks. Employers may thus still want to ask about a job applicant's criminal record later in the hiring process or make inaccurate judgments based on an applicant's demographic characteristics. In this paper, we use a field experiment conducted in partnership with a nationwide staffing platform to test policies that more directly address the reasons that employers may conduct criminal background checks. The experiment asked hiring managers at nearly a thousand U.S. businesses to make incentive-compatible decisions under different randomized conditions. We find that 39% of businesses in our sample are willing to work with individuals with a criminal record at baseline, which rises to over 50% when businesses are offered crime and safety insurance, a single performance review, or a limited background check covering just the past year. Wage subsidies can achieve similar increases but at substantially higher cost. Based on our findings, the staffing platform relaxed the criminal background check requirement and offered crime and safety insurance to interested businesses.

Zoe B. Cullen  
Rock Center 210  
Harvard Business School  
60 N. Harvard  
Boston, MA 02163  
and NBER  
zcullen@hbs.edu

Mitchell Hoffman  
Rotman School of Management  
University of Toronto  
105 St. George Street  
Toronto, ON M5S 3E6  
CANADA  
and NBER  
mitchell.hoffman@rotman.utoronto.ca

Will S. Dobbie  
Harvard Kennedy School  
79 John F. Kennedy St.  
Cambridge, MA 02138  
and NBER  
will\_dobbie@hks.harvard.edu

A data appendix is available at <http://www.nber.org/data-appendix/w29947>

# INCREASING THE DEMAND FOR WORKERS WITH A CRIMINAL RECORD\*

ZOË CULLEN  
WILL DOBBIE  
MITCHELL HOFFMAN

APRIL 2022

State and local policies increasingly restrict employers’ access to criminal records, but without addressing the underlying reasons that employers may conduct criminal background checks. Employers may thus still want to ask about a job applicant’s criminal record later in the hiring process or make inaccurate judgments based on an applicant’s demographic characteristics. In this paper, we use a field experiment conducted in partnership with a nationwide staffing platform to test policies that more directly address the reasons that employers may conduct criminal background checks. The experiment asked hiring managers at nearly a thousand U.S. businesses to make incentive-compatible decisions under different randomized conditions. We find that 39% of businesses in our sample are willing to work with individuals with a criminal record at baseline, which rises to over 50% when businesses are offered crime and safety insurance, a single performance review, or a limited background check covering just the past year. Wage subsidies can achieve similar increases but at substantially higher cost. Based on our findings, the staffing platform relaxed the criminal background check requirement and offered crime and safety insurance to interested businesses. *JEL* Codes: C93, J23, J24, M51

## I. INTRODUCTION

Employers are significantly less likely to interview or hire workers with a criminal record (WCs) compared to otherwise similar workers without a record (e.g., Pager, 2003; Holzer et al., 2006; Holzer, 2007; Agan and Starr, 2017). In 2008, for example, the average unemployment rate among formerly incarcerated people—27%—was higher than the U.S. unemployment rate for the general population at any point in history, including the Great Depression (Couloute and Kopf, 2018). The limited employment opportunities for WCs exacerbate existing socioeconomic and racial inequalities and likely contribute to the high rates of recidivism among recently released individuals (e.g., Yang, 2017; Schnepel, 2018).

In an attempt to mitigate the scarring effects of a criminal record, 35 states and over 150 cities and counties have adopted “Ban the Box” (BTB) policies that delay questions about a job applicant’s arrest and conviction record. These policies are meant to increase hiring and employment

\*We thank Amanda Agan, Paul Goldsmith-Pinkham, Sara Heller, Lawrence Katz, Judd Kessler, Eddie Lazear, Melvin Stephens, Crystal Yang, and numerous seminar participants for helpful comments and suggestions. Dylan Balla-Elliott, Dan Ma, and Alexia Olaizola provided excellent research assistance. We thank the Abdul Latif Jameel Poverty Action Lab (J-PAL), the Michael Lee-Chin Family Institute for Corporate Citizenship, and the Social Sciences and Humanities Research Council (SSHRC) for financial support. The views expressed here are those of the authors alone and do not necessarily reflect those of the institutions or funders involved with this work. The RCT and a pre-analysis plan were pre-registered with the AEA RCT registry under ID [AEARCTR-0005200](#). We are deeply grateful to Wonolo Inc. for their support of this work.

among WCs by making it more difficult to screen applicants based on their criminal history, helping WCs get “a foot in the door” when seeking employment. However, BTB policies do not address the underlying reasons that employers may conduct criminal background checks, such as the potential for lower productivity or higher downside risk. Employers may therefore still want to ask about an applicant’s criminal record later in the hiring process or make inaccurate judgments about an applicant’s criminal record based on their race or other demographic characteristics (e.g., Bushway, 2004; Holzer et al., 2006; Stoll, 2009; Agan and Starr, 2018; Doleac and Hansen, 2020).<sup>1</sup>

In this paper, we use a field experiment involving incentive-compatible decisions at nearly a thousand U.S. businesses to test several alternative approaches to increasing the demand for WCs. The alternatives we consider are meant to address potential underlying reasons that businesses may conduct criminal background checks. We offer crime and safety insurance to address downside risk concerns, and screening based on past performance reviews and the time since the most recent criminal record to address downside risk and productivity concerns. We also provide objective information on the average productivity of WCs compared to non-WCs to address downside risk and productivity concerns through a different channel, including the possibility that manager beliefs about WC productivity may be inaccurate. We benchmark the effects of each of these alternatives against the effects of a wage subsidy, a natural but potentially costly approach to increasing the demand for WCs.

The partner for our study is a large nationwide staffing platform based in the United States (hereafter, “the Platform”), which third-party businesses use to connect with available workers. Businesses submit job requests to the Platform that include a job description (typically short term), the pay for the job, and qualifying criteria. The Platform sends out the job offer to workers who meet the qualifying criteria and workers can then accept the job on a first-come, first-serve basis. Businesses rarely cancel jobs after they have been accepted due to clear procedures and fee structure. Cancellations of matches occur in less than 1% of cases. Presenting workers with the option to accept a job is therefore equivalent to that business extending a job offer to those workers.

The Platform’s design allows us to ask hiring managers to make incentive-compatible choices over potential hiring decisions, as opposed to callback or interview decisions that have generally been considered in past work.<sup>2</sup> Hiring managers that use the Platform were already familiar with submitting criteria for workers who can accept their jobs. We truthfully informed these managers

<sup>1</sup>Perhaps as a result, existing work shows mixed effects of BTB policies on the employment of WCs and other populations of interest. Jackson and Zhao (2017) find, for example, that a 2010 law in Massachusetts that included a BTB provision had a small negative impact on the employment and earnings of ex-offenders, while Rose (2021) finds negligible impacts of a 2013 Seattle law prohibiting employers from asking about an applicant’s arrest and conviction record until after an initial screening. By contrast, Craigie (2020) finds that BTB policies increase the probability of public-sector employment for WCs. In aggregate data that include both those with and without a criminal record, Doleac and Hansen (2020) find that state and local BTB laws decrease employment rates for young, low-skill Black and Hispanic men, while Shoag and Veuger (2016) find that these state and local BTB laws increase employment rates in high-crime counties. In related work, Sabia et al. (2020) and Sherrard (2021) estimate the impact of BTB laws on crime and recidivism, respectively.

<sup>2</sup>Workers who use the Platform to connect with third-party businesses are independent contractors. Accordingly, many of the terms in this paper (e.g., “employment,” “hire,” “hiring managers,” “wages,” and “staffing”) are used only for convenience and simplicity, and do not legally apply to the Platform or to the workers who use it.

that, in the same way, their responses during the experiment constituted high-stakes decisions that could determine whether WCs would be allowed to accept their jobs in the future. For example, if a hiring manager indicated that they would be willing to work with WCs under a certain insurance policy during the experiment, then WCs would be eligible to accept jobs posted by that manager after such an insurance policy is made available. The high-stakes nature of these choices was not just theoretical—the choices that hiring managers made during the experiment actually affected whether WCs could accept their posted jobs.

In the experiment, the Platform asked hiring managers at nearly 1,000 businesses whether they would allow WCs to accept their jobs given the availability and level of wage subsidies, crime and safety insurance, past performance reviews, and a more limited screening of criminal records. Starting with the baseline level of demand for WCs in our sample, we find that a sizable share of businesses, 39%, are willing to work with WCs without additional incentives or conditions. The level of demand, still without additional incentives or conditions, increases to about 45% for jobs that do not involve customer interactions and 51% for jobs that do not involve high-value inventory, consistent with businesses perceiving risks related to customer safety or inventory theft. We also find that the share of businesses willing to work with WCs increases to 68% if businesses are having a hard time filling a job on the Platform, consistent with businesses being more likely to consider non-traditional workers in tight labor markets.

Turning to our main experimental results, we find that the share of businesses willing to work with WCs further increases by at least 10 percentage points when businesses are offered a modest level of crime and safety insurance, or when WCs must have received a single positive performance review to qualify, or when a more limited background check rules out a conviction or arrest in the past year. Wage subsidies can achieve similar increases but only at relatively high subsidy levels that may be cost prohibitive.<sup>3</sup> For example, providing crime and safety insurance covering damages up to \$5,000 increases the level of demand for WCs by about 12 percentage points, approximately equivalent to the effects of an 80% wage subsidy according to a linear extrapolation of our experimental subsidy estimates. Requiring that WCs have successfully completed one prior job posted on the Platform similarly increases the level of demand by about 11 percentage points, again roughly equivalent to the effect of an 80% wage subsidy. Limiting the pool of WCs to those who have not been convicted or arrested in the past one year increases the level of demand by about 21 percentage points, greater than the effect of a 100% wage subsidy. We note, however, that limiting the pool of WCs mechanically reduces the level of demand to zero for the screened out individuals.

The final option we consider is providing hiring managers with objective information on the performance of WCs. We exploit the fact that some WCs inadvertently accessed the Platform before their background screening results were known, allowing us to objectively compare the customer performance ratings of WCs and non-WCs in their first jobs. Using incentive compatible

<sup>3</sup>We find that the share of businesses willing to work with WCs increases by approximately 2.1% for every 10% increase in the offered wage subsidy. We show in Section VIII that these estimates imply that all of the non-subsidy policies we consider can increase the demand for WCs at  $\frac{1}{2}$  to  $\frac{1}{10}$  the cost of wage subsidies under reasonable assumptions.

elicitation about performance beliefs in our experiment, we find that hiring managers underestimate the performance of WCs in terms of both high- and low-performance ratings. Providing objective information on the true share of high-performance ratings received by WCs leads to more accurate beliefs and increases WC hiring by about 6 percentage points on average, roughly equivalent to the effect of a 40% wage subsidy. Providing objective information on the share of low-performance ratings, typically resulting from no-shows, also leads to more accurate beliefs but does not have a statistically significant impact on the willingness to work with WCs. The muted effect of low-performance ratings compared to high-performance ratings suggests that businesses may be less concerned with no-shows when deciding whether to work with WCs, and more concerned that WCs might not satisfactorily meet their performance standards while on the job.<sup>4</sup>

Based on our findings, the Platform carried out a staged roll-out to relax the criminal background check procedures it coordinates for customers. First, using the responses we collected in our experiment, the Platform extended job offers to WCs from businesses that responded affirmatively to working with WCs under then-current Platform conditions. Second, the Platform created the option for thousands of businesses posting new jobs to allow WCs to accept their jobs with up to \$1 million covered by crime and safety insurance, one of the most promising randomized conditions tested in this study. Eventually, the Platform plans for the inclusion of WCs to be the default option for all businesses, with businesses needing to pay an additional cost to exclude WCs from the pool of qualified workers. Through August 2021, demand from our study participants and the staged roll-out led to over 12,000 jobs being available to WCs on the Platform. This rapid expansion in the number of jobs available to WCs opens new questions for future research, including the evolution of demand as businesses gain experience working with WCs and the opportunities created for WCs that accept jobs posted on the Platform.

Beyond demonstrating that a range of policies can increase the demand for WCs, our paper provides new evidence on why businesses may conduct criminal background checks and hence what types of policies are likely to increase the demand for WCs. Several of our results, including the large demand response to crime and safety insurance and the even larger response to insurance among businesses whose jobs involve customer interaction or high-value inventory, suggest that businesses are sensitive to the downside risk of hiring WCs. Other results, including the large response to objective information on the productivity of WCs, are consistent with the view that some businesses view WCs as less productive on average than otherwise similar non-WCs (a view that is incorrect in our setting). The positive effects of the wage subsidies, performance screening, and screening of the most recent records are consistent with both explanations.<sup>5</sup>

<sup>4</sup>Our high-performance estimates are comparable to other information treatment experiments. We find that our high-performance information treatment increased the share of firms willing to work with WCs by 6 percentage points on average, or about 15%. By comparison, Bursztyn et al. (2020) find that correcting beliefs about male support for female labor market participation in Saudi Arabia increases the likelihood that husbands sign up their wives for job platform by 9 percentage points, or 36%, while Allcott (2011) find that a letter on energy use comparisons reduces consumption by 2%.

<sup>5</sup>These results are related to the literature in personnel and organizational economics on hiring (see Oyer and Schaefer (2011) for a review of recent work in this area). While past work explores the option value of hiring high-variance workers for long-term positions (Lazear, 1998; Bollinger and Hotchkiss, 2003), we explore how to protect



Our paper also provides new evidence on why there may be such low take-up of existing federal programs meant to increase the demand for WCs, such as the Work Opportunity Tax Credit or Federal Bonding Program. For example, the Work Opportunity Tax Credit provides up to a 40% wage subsidy for hiring disadvantaged workers, including individuals convicted of a felony and released from prison in the past year. Our findings suggest that the relatively low take-up of this program may be because employers are less willing to work with recently convicted or released WCs, particularly in the first year. The Federal Bonding Program instead provides insurance to employers hiring hard-to-place job candidates, which our findings suggest should increase the demand for WCs. The relatively low take-up of this program is therefore more likely the result of employers being unaware of the program or having difficulty navigating the program requirements.

Finally, our paper builds on important work by Holzer (2007), Holzer et al. (2007), and Hunt et al. (2018) measuring the demand for WCs using employer surveys. In a survey of 107 businesses, for example, Hunt et al. (2018) find that employers report being more willing to hire WCs if there are wage subsidies, certificates of validated work performance history, or guaranteed replacement workers. These employers also report that “any violent felony conviction” and the “skills to get the job done” are their two most serious concerns with hiring WCs in the absence of these policies in these studies. However, the hypothetical and low-stakes nature of these surveys makes it difficult to know whether employers are expressing their true preferences or just their aspirations. We add to this literature by measuring the demand for WCs using incentive-compatible choices of nearly a thousand U.S. businesses under different counterfactual policies.

The remainder of the paper is organized as follows. Section II describes the experimental context and design. Section III presents our baseline estimates of the labor demand for WCs with and without wage subsidies. Section IV presents results for our primary experimental interventions. Section V presents results from the experiment on providing businesses information on the average productivity of WCs relative to non-WCs. Section VI discusses alternative explanations for our results, and Section VII concludes. The Online Appendix provides additional results and details of the experimental design.

## II. CONTEXT AND EXPERIMENTAL DESIGN

### *II.A. The Platform*

The context for our study is a leading online labor platform that thousands of third-party businesses use to connect with workers seeking short-term jobs. Businesses use the Platform to connect with workers to fill a wide range of entry-level jobs in sectors that report being more willing to hire WCs, such as general labor, hospitality, and transportation, as well as entry-level jobs in customer-facing or administrative sectors that are traditionally more averse to hiring WCs (e.g., Holzer et al., 2004; Raphael, 2011; Yang, 2017; Schnepel, 2018). The Platform is hosted on web and mobile app, but the work they support generally does not involve computers or the internet, nor businesses from the perceived downside potential of hiring disadvantaged workers for short-term positions.

does it require a college degree or significant prior experience. The variety of job types and focus on entry-level jobs provide an ideal setting for estimating the demand for WCs under different policy alternatives.

Three institutional features of the Platform are important for our analysis. First, the Platform’s labor market allows us to ask businesses to make incentive-compatible choices over potential hiring decisions. Businesses submit job requests to the Platform that include a job description, the pay for the job, and qualifying criteria. For example, some requests specify that workers must have experience driving a truck or be comfortable with heavy lifting. Businesses do not decide whether to work with individual workers. The Platform extends the job offer to workers who meet the job qualifications. Workers then have the option to accept or reject these job offers on a first-come, first-serve basis, at their discretion. Upon acceptance, the job is reserved for the worker who accepted it (i.e., other workers no longer have the ability to accept the job). Posted jobs are typically accepted within a few hours of the initial posting, though jobs can sometimes go unaccepted for several days. If a business wishes to cancel a job request that has been accepted within 12 hours of the start time for the job, the business generally must pay a cancellation fee consisting of a set percentage of the anticipated payment for the job (typically 50%). Less than 1% of accepted job requests are canceled in practice. By asking businesses to make decisions about what workers they would allow in their pool, we are therefore asking businesses to make incentive-compatible choices over potential hiring decisions, as opposed to callback or interview decisions that have generally been considered in previous work. While a particular job may be short-term (one day or one week), 54% of workers who accept a job with a given business will later be re-hired by the same business, suggesting potential for long term employment relationships.

Second, the performance data collected by the Platform allows us to provide hiring managers with objective information about the performance of both WCs and non-WCs. At the end of each job, the business’s hiring manager is asked to rate each worker’s overall performance on a scale of one to five. Hiring managers also rate workers on specific attributes such as timeliness, cooperation, and quality of work. In practice, however, hiring managers only complete the more comprehensive review after 8% of jobs compared to 86% for the overall ratings. We therefore focus on the overall ratings throughout the paper. In 2019, roughly 85% of the overall ratings are perfect five-star ratings and 1.3% are one- or two-star ratings. No-shows comprise an additional 4% of the overall job ratings. In our analysis, we refer to information about the share of jobs resulting in five-star ratings as “high-performance” information and information about the share of jobs resulting in a one- or two-star rating or a no-show as “low-performance” information. The intuition for considering the two ends of the performance spectrum separately is that they are only weakly correlated—a worker can perform at a high level conditional on completing the job while also exhibiting high no-show rates. Some businesses might care more about mitigating poor performance and absenteeism than about the ability to perform well.

Finally, like many other labor platforms for independent contract workers (e.g., Uber, Lyft), up to 30% of Platform applicants are currently screened out by a criminal background check, and



the company incurs a cost to run each background check. Beyond the background check, the Platform does not do extensive screening of potential workers. The researchers' collaboration with the Platform grew out of a series of conversations between the researchers and the Platform's Chief Executive Officer, Chief Technology Officer, board members, and other top executives and managers, each party recognizing that the costs and benefits of the criminal background check were largely unknown despite its significant impact on operations.

## *II.B. Experimental Design*

The experiment came about through an intense multi-year collaboration between the research team and the Platform's top executives and managers following the initial conversations discussed above. The goal of the collaboration was to understand the potential barriers to including WCs in the pool of independent contract workers on the Platform so that the company could modify, reduce, or eliminate the criminal background check requirement and provide opportunities to a broader set of workers. The Platform's top executives and managers piloted the randomized conditions, while the Platform's general counsel closely scrutinized and edited the conditions to ensure that the hiring managers' responses could legally determine whether WCs would be allowed to accept their business' jobs in the future (hence, ensuring the high-stakes nature of the responses provided during the experiment). The Platform's general counsel also ensured that the proposed policies were in compliance with the relevant local, state, and federal laws.

A central feature of the experiment is that hiring managers make *ex-ante* incentive-compatible choices under different randomized conditions. As discussed above, the Platform's labor market features a matching process where workers who meet the posted job requirements are matched on a first-come, first-served basis, with no additional screening after the initial matching process. In addition, hiring managers on the Platform were already familiar with submitting criteria for workers who can accept their jobs, making the high-stakes nature of their choices both apparent and natural. These institutional features, as well as the input of the Platform's general counsel, allowed the Platform to truthfully inform hiring managers that their responses during the experiment constituted high-stakes decisions that could determine whether WCs would accept their jobs in the future.<sup>6,7</sup>

<sup>6</sup>Incentive compatibility can be judged by whether the questions are perceived by respondents as having the potential to affect their outcome, with the exact probability that a choice is implemented generally mattering relatively little (Carson and Groves, 2007; Charness et al., 2016). The Platform carefully crafted the initial outreach to businesses to elicit truthful answers: "Please share your truthful and considered views—they matter to us," and the Platform's executives and board members were especially committed to making it possible for those who expressed interest in hiring WCs to do so. As a result, the questions were designed to be both directly implementable and to change the Platform's practices moving forward.

<sup>7</sup>Methodologically, we build on Mas and Pallais (2017), Low (2017), and Kessler et al. (2019) to generate incentive-compatible responses in field experiments. Mas and Pallais (2017) examine the choice of applicants regarding schedule flexibility over jobs within a call center. Low (2017) and Kessler et al. (2019) examine hypothetical candidates (for dating and hiring, respectively) with randomized attributes, where respondents are truthfully informed that their decisions will affect who can accept their position. In our experiment, businesses make multiple decisions about hiring under different randomized conditions, where the Platform truthfully informs them that their decisions may affect whether WCs are included among the workers who can accept their jobs and under what conditions. Our approach is

The *ex-ante* incentive-compatible nature of the hiring managers’ choices was reinforced by two *ex-post* actions taken by the Platform following the experiment. First, the Platform immediately implemented businesses’ choices by allowing WCs to accept jobs posted by the businesses that were willing to work with WCs under current Platform conditions if there was a pool of WCs in the businesses’ location. Second, the Platform used businesses’ choices to set Platform-wide policy. The Platform offered crime and safety insurance up to \$1 million, regardless of client choices or participation in the experiment. The Platform is also planning to change the default option to allow WCs to accept the jobs businesses post, while retaining the crime and safety insurance, in the coming months. The connection of WCs to study participants and these platform-wide changes have already generated over 12,000 jobs (through August 2021) newly open to WCs, with additional jobs being opened to WCs each day.<sup>8</sup>

We leverage this institutional setting to test several approaches to increasing the demand for WCs on the Platform. Each of the alternatives we consider is meant to address the underlying reasons that businesses may screen workers using a criminal background check. We began by asking hiring managers about their willingness to accept WCs under one of several randomly assigned wage subsidy levels (0%, where we made no mention of the subsidy at all, or one of several positive levels) to establish the baseline level of demand and provide a benchmark for the other randomized treatments. We then asked hiring managers about their willingness to accept WCs under different randomized conditions, including different levels of crime and safety insurance and past performance reviews, as well as background checks covering only the most recent criminal records. We then asked a series of descriptive questions about the hiring practices at the business and the types of jobs posted on the Platform to allow for heterogeneity analyses, before concluding with an information experiment motivated by the large dispersion in prior beliefs about the performance of WCs on the Platform. The randomly assigned subsidy level remained in place throughout all of these subsequent questions.

The remainder of this section summarizes the most important details of the experiment, also detailed in Table I and Appendix Table B.1. We begin by describing how the Platform contacted hiring managers, before describing each of the main experimental conditions and subsample comparisons.

*Outreach.* From March 6, 2020 to April 11, 2020, the Platform emailed active hiring managers the following message: “We are considering expanding the pool of [workers] who can perform the jobs that you post, and we want your guidance.” Interested hiring managers were instructed to click on a link that took them to the hiring flow questions that constitute our randomized experiment. The initial outreach emails did not mention WCs, and hiring managers were not aware that they were part of a randomized study at any point during the outreach or experiment. The Platform sent the emails from a Platform-branded account using their own signature (“Sincerely, [Platform]

also similar to the “strategy method” in lab experiments (Brandts and Charness, 2011), where players make multiple conditional decisions and one decision is potentially randomly implemented.

<sup>8</sup>The Platform is not able to automatically track the number of WCs that have been on-boarded and assigned jobs. Based on conversations with the Platform’s team that manually matches WCs to jobs, for jobs made open to WCs, about 4% of jobs were worked by WCs through August 2021.

Management”) and logo. The Platform also offered a \$35 or \$50 cash gift for complete answers to underscore the value of thoughtful and considered responses, as well as to motivate businesses to complete all questions. Such cash transfers are standard practice for the Platform when requesting input from hiring managers to make Platform design decisions.<sup>9</sup>

Following a series of short introduction questions, the Platform explained:

We are considering expanding our pool of [Platform Workers] to include individuals that have a criminal record. We want to learn whether this expanded pool would suit your needs.

If you indicate that you’re interested in connecting with [Platform Workers] with a criminal record, then (and only then) your choice could affect whether these [Platform Workers] are able to accept jobs you post. These individuals would be at most 5% of your pool of possible matches.

The Platform then asked participating hiring managers about their willingness to work with WCs under different randomized conditions, where randomization occurred at the business level to ensure that hiring managers at the same business were not given conflicting options.

In total, 1,095 hiring managers from 913 businesses completed the hiring flow questions. Eighty-six percent of hiring managers completed the hiring flow questions conditional on opening the email, with 91% of managers completing the questions conditional on reaching the first question related to WCs. Eighty percent of the hiring managers in our sample report having the authority to unilaterally allow WCs to perform the jobs they post or to significantly influence this decision. We show in Appendix Tables B.7 and B.8 that our results are qualitatively unchanged if we focus on the subset of hiring managers with unilateral authority. In unreported results, we also find that our results are also qualitatively unchanged if we calculate upper and lower bounds of all treatment effects that account for early attrition.<sup>10</sup> We include responses from all hiring managers completing our survey to allow for separate decisions in multiple-establishment businesses (e.g., a hardware chain with multiple stores). In unreported results, we find very similar results when excluding businesses with responses from multiple managers.

*Baseline Demand.* We measure the baseline demand for WCs with no additional incentives or conditions by simply asking hiring managers whether their business would permit WCs to accept their jobs:

<sup>9</sup>The Platform emailed all hiring managers who had been active in listing positions on the Platform in the last 16 months up to five times during the outreach period. The Platform did not email hiring managers who had joined within the last three months at the request of the company’s account managers.

<sup>10</sup>The Platform emailed 7,450 hiring managers during the outreach, meaning that 14% of hiring managers opened the email and completed all of the hiring flow questions. The Platform generally expects less than 10% of the full pool of past hiring managers to respond when requesting input to make Platform design decisions. Appendix Table B.2 shows that firms responding to the survey and entering our experimental sample are similar in terms of industrial composition to firms that did not respond to the survey, but skew younger (19 years vs. 22 years) and larger (40 employees vs. 21 employees). Respondents are more active than non-respondents in the labor market at the time of the survey, averaging 2,828 posted jobs on the platform over the past two years compared to 561 posted jobs. We are unable to compare firms based on WC hiring policies, as we gather this information during the survey experiment.

Would you permit [Platform Workers] with a criminal background to perform jobs you post?

We asked  $\frac{1}{5}$  of hiring managers this question, meant to measure demand for WCs under the Platform’s current conditions and establish a baseline for a wage subsidy of 0%. Hiring managers were given the option of selecting “Yes,” “Only if it’s hard to fill my jobs,” or “No.” Answering “Yes” to this question immediately extended permission to the Platform to allow WCs to accept the client’s job posting, without any policy changes or conditions being met.

*Wage Subsidies.* We measure demand for WCs under different wage subsidies by asking hiring managers whether their business would permit WCs to accept their jobs under one of several randomly assigned wage subsidy levels:

If the [Platform] gave you a {Wage Subsidy} discount for [Platform Workers] with a criminal record, would you permit such [Workers] to perform jobs you post? This means you would only pay  $(100 - \{\text{Wage Subsidy}\})$  of the wage for those with a criminal record. All [Platform Workers] would still receive the full pay amount after the discount (the [Platform] would pay the difference).

The wage subsidy levels were 5%, 10%, 25%, 50%, and 100%, randomly assigned with probabilities  $\frac{1}{10}$ ,  $\frac{1}{10}$ ,  $\frac{1}{5}$ ,  $\frac{1}{5}$ , and  $\frac{1}{5}$ , respectively. The remaining  $\frac{1}{5}$  of hiring managers were randomly assigned to no subsidy condition described above. The wide range of randomized subsidy levels allows us to trace out a labor demand curve with minimal assumptions and explore whether there are non-linear effects of very small or very large wage subsidies. We cover a range of economically relevant subsidy levels, with the Federal Work Opportunity Tax Credit currently offering a 25% wage subsidy to businesses who employ WCs for at least 120 hours in their first year of employment and a 40% wage subsidy to businesses who employ WCs for at least 400 hours in their first year. For expository purposes, we pool the 5% and 10% subsidy levels, which results in a uniform number of observations across values displayed.

The wage subsidy assigned to each manager is kept constant throughout all subsequent experimental treatments. We agreed with the Platform that it would be most natural to have the wage subsidy remain constant throughout the experiment. We also find nearly identical results when restricting to the subset of hiring managers who were randomized to have no wage subsidy or, equivalently, when fully interacting the other treatments and the wage subsidy. We discuss these results in additional detail below.

*Crime and Safety Insurance.* We measure the effect of crime and safety insurance by asking hiring managers if, at a given subsidy level, their business would permit WCs to accept their jobs under one of several randomly assigned insurance levels:

If the [Platform] could cover damages up to {Crime and Safety Insurance Cap} related to theft or safety incurred by [Platform Workers] with a criminal record, would you

permit such [Workers] to perform jobs you post? The [Platform] would still give you a {Wage Subsidy} discount, but no other supplementary policies would apply.

The randomly assigned insurance levels were \$1,000, \$5,000, \$100,000, and \$5 million, randomly assigned with probabilities of  $\frac{1}{6}$ ,  $\frac{1}{6}$ ,  $\frac{1}{3}$ , and  $\frac{1}{3}$ , respectively. These randomized insurance levels cover a wide range of economically relevant values. The U.S. Federal Bonding Program, for example, offers an insurance bond of \$5,000 to provide insurance against liability for relatively less serious crimes like robbery or theft. The highest level of insurance in our experiment, \$5 million, would also provide liability against much more serious crimes like sexual assault and murder. Crime and safety insurance directly addresses the concern that a WC might act violently towards coworkers or customers. For expository purposes, we pool the \$1,000 and \$5,000 insurance levels, which results in a uniform number of observations across values displayed.

*Screening Based on Performance History.* We measure the effect of having a satisfactory performance history on the Platform by asking hiring managers if, at a given subsidy level, their business would permit WCs to accept their jobs under one of several randomly assigned job histories:

If the [Platform] required [Platform Workers] with a criminal record to have satisfactorily completed {Performance History} job(s), receiving more than 85% five-star reviews, would you permit such [Workers] to perform jobs you post? The [Platform] would still give you a {Wage Subsidy} discount, but no other supplementary policies would apply.

The randomly assigned job histories consisted of one, five, and 25 jobs, randomly assigned with  $\frac{1}{3}$  probability each. These randomized job histories again cover a wide range of economically relevant values. Pallais (2014) shows that workers having one prior job substantially increases the chance of getting hired on oDesk, motivating the inclusion of this job history in our experiment, while the highest value of 25 jobs corresponds to an above the 90th percentile of past performance history on the Platform. Performance screening could potentially address business concerns about both productivity and on-the-job crime.

*Screening Based on Criminal Record History.* We measure the effect of more targeted screening by asking hiring managers if, at a given subsidy level, their business would permit WCs to accept their jobs if the WC had maintained a clean record for at least one, three, or seven years, with these values randomly chosen with  $\frac{1}{3}$  probability each. We chose these randomized values because the probability of criminal re-offending is particularly high in the first two years post-incarceration, while background checks often extend to criminal convictions within the last seven years.

We additionally measure the effect of selectively screening by conviction type by asking hiring managers if, at a given subsidy level, their business would permit WCs to accept their jobs if they were convicted for a distinct category of crimes, including a violent felony, a violent misdemeanor, a property/financial felony, a property/financial misdemeanor, a drug-related felony, and a drug-related misdemeanor. These categories include a wide variety of crimes, but do not encompass all possible

conviction types and do not include arrests that do not result in a conviction but nevertheless are reported on a criminal background check. We therefore do not expect these conviction-specific results to aggregate to our baseline results that include all arrest and conviction types.

*Objective Performance Information.* We measure the effect of providing hiring managers with objective information about the average performance of WCs on the Platform on hiring managers’ beliefs and subsequent hiring choices using an information experiment embedded in the hiring flow questions. Here, we exploit the fact that 5% of WCs who attempted the onboarding process onto the Platform completed their first job on the Platform before the results of their background screening results prohibited them from further activity. Given the quasi-random nature of the relative timing between the start of the job and criminal screening completion time, we consider those WCs who have been screened out after completing their first job to be representative of the pool of WCs excluded from the Platform only due to the crime screen. This allows us to compare the performance ratings of the first job of WC (otherwise qualified) applicants with non-WCs that join the Platform. These performance data reveal that WCs modestly out-perform non-WCs on the Platform, consistent with prior work showing that individuals with criminal records stay in jobs longer and are less likely to voluntarily quit compared to other workers (Minor et al., 2018).

We first measure baseline performance beliefs using an incentive-compatible guessing game about the relative performance of WCs on the Platform that rewards accuracy. Rewards in the guessing game ranged between \$2 and \$10 for an answer within 5% of the truth, where in unreported results we find no difference in respondent accuracy across the reward amounts. For approximately  $\frac{1}{2}$  of the participants, the guessing game asked the following question about the share of high-performance ratings:

In 2019, 86% of jobs on the [Platform] resulted in a five-star rating. What percentage of jobs completed by [Platform Workers] with a criminal record do you think would result in a five-star rating on the [Platform] or a similar platform? If your guess is within 5% of the truth, we will send you an additional {Bonus} reward!

We asked the other  $\frac{1}{2}$  of participants about the share of low-performance ratings in the guessing game:

In 2019, 5% of jobs on the [Platform] resulted in either a no-show or low rating (one or two stars). What percentage of jobs completed by [Platform Workers] with a criminal record do you think would result in a no-show or low rating on the [Platform] or a similar platform? If your guess is within 5% of the truth, we will send you an additional {Bonus} reward!

One-half of participants who initially made guesses about the share of high-performance ratings randomly received objective information about the true share of high-performance ratings:

The truth is that 87% of jobs completed by [Platform Workers] with a criminal record resulted in a five-star rating on the same or a similar platform—actually better than



everyone else. Please take some time to read and understand this information carefully. When you are ready, proceed to the next screen.

Similarly,  $\frac{1}{2}$  of the participants who initially made guesses about the share of low-performance ratings received objective information about the true share of low-performance ratings:

The truth is that only 3% of jobs completed by [Platform Workers] with a criminal record resulted in either a no-show or a low rating (one or two stars) on the same or a similar platform—actually fewer no-shows and low ratings than everyone else. Please take some time to read and understand this information carefully. When you are ready, proceed to the next screen.

We then asked all participants, regardless of whether or not they were shown the new objective information, to report their posterior beliefs about the performance of WCs using the same incentive-compatible guessing game question. Finally, we allowed all participants, again regardless of whether or not they were shown the new objective information, to revise their answer to the very first question about hiring WCs with or without a wage subsidy. By allowing participants to revise their willingness to work with WCs, we can learn how the objective information about WC performance on the Platform impacts hiring decisions.

*Heterogeneity by Labor Market Conditions.* We also explore heterogeneity across labor market conditions, testing whether particular business concerns are more salient under different market conditions. For all of the questions in the experiment, hiring managers were given the option of selecting “Only if it’s hard to fill my jobs,” providing a targeted measure of labor market tightness that is specific to each business’s context. In addition, we explore the effects of local labor market unemployment, a more traditional measure of labor market tightness, by asking whether a hiring manager would want to work with WCs if the local unemployment rate were to be at a certain level, randomized between 2%, 6%, or 10%. Finally, we estimate results separately for businesses located in metropolitan areas with above- and below-median unemployment rates from January to March 2020 measured using all workers, workers in just the businesses’ industry, and workers with just a high school diploma or below.

*Heterogeneity by Job Characteristics.* Finally, we asked hiring managers about the typical jobs they post, including whether there are any customer interactions or access to high-value inventory, thereby allowing us to explore to what extent businesses are concerned about theft as compared to violence when considering downside risk. We also asked hiring managers whether their business has a hiring policy related to WCs.

### *II.C. Motivating Framework*

The experiment is motivated by a stylized theoretical framework that explains why businesses may conduct criminal background checks and decide not to work with WCs. The framework

formalizes the idea that businesses may use a criminal record as a (potentially inaccurate) signal of lower productivity or higher downside risk. The framework helps explain how we can use our results to understand why businesses may want to conduct criminal background checks and what types of policies are likely to increase the demand for WCs.

Consider a single business deciding whether to work with a single WC. The business’s expected profits from working with the WC is a function of expected productivity and the risk of a costly event occurring on the job:

$$\pi = y - w - b \cdot \max\{k - I, 0\}$$

where  $y$  is the expected productivity of the WC (e.g., the rate at which the worker packs boxes),  $w$  is the WC’s wage,  $b$  is the probability of a bad event occurring as a result of the WC’s behavior (e.g., theft),  $k \geq 0$  is the cost of a bad event (e.g., the value of the stolen inventory), and  $I \geq 0$  is the amount of crime and safety insurance provided, if any. The business has an unobserved shadow value,  $\theta$ , of not hiring the WC (e.g., there is some probability that the business can fill the slot with a non-WC). The business chooses to work with the WC,  $H = 1$ , when  $\pi > \theta$ . The first prediction is that wage subsidies increase demand for WCs regardless of expected productivity or downside risk. We are thus able to use the effect of the wage subsidy as a benchmark, comparing its effect to policies that primarily target either expected productivity or downside risk. Our framework yields the following predictions.

*Crime and Safety Insurance.* Crime and safety insurance (i.e., greater  $I$ ) increases the demand for WCs as long as downside risk is a relevant factor for businesses. The effect of insurance on hiring should also be larger for businesses with a greater probability of, or larger costs from, a bad event. This is likely to include businesses with jobs that involve high-valued inventory or frequent customer interactions (although we note that insurance may increase hiring at all firms due to concerns about, e.g., violence towards coworkers). If the primary downside risks involve infrequent but very costly events (e.g., violence towards customers or coworkers), then we expect that an insurance policy with a low cap will not impact hiring demand but a very generous insurance policy will. If, on the other hand, the primary downside risks are minor infractions (e.g., petty theft), then we expect the effect of insurance policies with low and high insurance caps per event to similarly increase hiring demand.

*Screening Based on Performance History and Objective Performance Information.* Requiring that WCs successfully complete a prior job can be viewed as increasing the expectation about productivity,  $y$ , for that worker. While screening could also decrease the perceived probability of a bad event  $b$ , our conversations with the Platform suggest that expected productivity is the primary signal contained in prior ratings. Nevertheless, screening based on performance history should increase the demand for WCs if either productivity or downside risk is a relevant factor for businesses. If businesses have negatively biased beliefs about  $y$  and  $b$  for WCs, providing objective performance information regarding WC performance can also increase the demand for WCs.

*Screening Based on Criminal Record History.* Expected productivity  $y$  may be higher and both

the probability of a bad event  $b$  and the cost of that bad event  $k$  may be lower for WCs with less recent criminal histories or convicted of less serious crimes. This combination of lower  $b$  and  $k$  leads to higher demand for WCs with less recent criminal histories or convicted of less serious crimes compared to WCs with more recent criminal histories or convicted of more serious crimes if either productivity or downside risk is a relevant factor for businesses.

*Heterogeneity by Labor Market Conditions.* When local labor market conditions are such that the business has strong alternative options to hiring WCs, the shadow value of labor in our framework,  $\theta$ , rises. The business chooses to work with the WC,  $H = 1$ , when the value reaped from hiring is greater than the alternative,  $\pi > \theta$ . Hence, we expect that demand for WCs is lower when the businesses face favorable labor market conditions and their jobs are easy to fill.

*Heterogeneity by Job Characteristics.* Businesses with jobs that involve high-valued inventory likely face a higher probability of a bad event occurring (i.e., higher  $b$ ) and a higher cost from such a bad event (i.e., higher  $k$ ). Thus, businesses with jobs involving high-valued inventory should have relatively lower demand for WCs if downside risk is a relevant factor. Similarly, businesses with jobs that require frequent customer interactions likely have more opportunities for costly infractions to occur (i.e., higher  $b$ ) and a higher cost from such an infraction event (i.e., higher  $k$ ), again leading to relatively lower demand for WCs if downside risk is a relevant factor.

#### *II.D. Descriptive Statistics and Randomization Assessment*

Table II presents descriptive statistics for the experimental sample comprised of the 1,095 hiring managers from 913 businesses that completed the experiment, and a broader set of businesses in the United States. Panel A reports information on basic business characteristics from the Infogroup Historical Business Database (Infogroup, 2016), which contains basic profile data for more than a million U.S. businesses. Businesses in our experimental sample are broadly representative of U.S. businesses in terms of industrial composition, but skew older (19 years vs. 16 years) and larger (40 employees vs. 2.5 employees). Businesses in our experimental sample are also somewhat more likely to be in manufacturing (19% vs. 6%), transportation (10% vs. 3%), and public administration (10% vs. 2%), and less likely to be in service (31% vs. 37%), finance (3% vs. 7%), and construction (1% vs. 8%).

Panel B reports information on WC hiring policies, where information for the broader set of U.S. businesses comes from a nationwide survey of over 1,000 HR professionals commissioned by the Society for Human Resource Management (SHRM) (Society for Human Resource Management, 2018). Compared to other U.S. businesses, businesses in our experimental sample are somewhat less likely to have a business-wide WC hiring policy (45% vs. 66%) and to indicate that they want to work with the best candidate for the job regardless of criminal history (46% vs. 53%). Slightly more businesses in our sample indicated that they would want to work with WCs to help give individuals a second chance (50% vs. 38%) or for financial incentives (8% vs. 2%), and a similar number of businesses in both samples are concerned about local or state regulations that make hiring WCs

difficult (26% vs. 22%). Section VI explores how our results change if observations are weighted to match either the distribution of firms in the U.S. economy based on industry and firm size or based on answers to the SHRM questions about attitudes and policies in place with respect to WCs.

Table III shows that the randomization was balanced in our experimental sample. We regress seventeen business characteristics on indicator variables for all levels of the six randomized treatments. Table III reports p-values from an F-test of each of the 90 regressions. Only three of the p-values are statistically significant at the 5% level and only an additional six are significant at the 10% level, which is to be expected given the number of tests. These results indicate that randomization was performed correctly and that our sample is balanced across treatment arms.

### III. THE LABOR DEMAND FOR WORKERS WITH A CRIMINAL RECORD

In this section, we measure the baseline demand for WCs using the randomized wage subsidies. We first analyze the effects of the wage subsidies on the willingness to work with WCs for all jobs, before measuring the level of demand for different types of jobs and local labor market conditions. We present our results graphically, providing regression tables in the Appendix when noted.<sup>11</sup>

#### III.A. Baseline Results

Panel A of Figure I plots the fraction of businesses that are willing to work with WCs at each randomized wage subsidy. We show our baseline results, where we code businesses as willing to work with WCs if they responded “Yes” and unwilling to work with WCs if they responded “No” or “Only if it’s hard to fill my jobs.” We pre-registered our main analyses using this form of the dependent variable since the answer of “Yes” is unambiguous and allows for choices to be legally binding. Below we will also show results where we code businesses as willing to work with WCs if they responded “Yes” or “Only if it’s hard to fill my jobs” to the relevant question and unwilling to work with WCs if they responded “No.” Throughout the paper, standard errors are clustered by business, the level of random assignment.

Panel B of Figure I uses the randomized subsidy results to show the labor demand curve for WCs on the Platform. We also report an estimate of the average demand elasticity,  $\epsilon^D$ . To estimate  $\epsilon^D$ , we first estimate the following linear specification that includes one observation per respondent:

$$\text{Hire}_i = \gamma_0 + \gamma_1 \cdot \text{EffectiveWage}_i + u_i \tag{1}$$

where  $\text{Hire}_i$  is an indicator that represents whether business  $i$ ’s is willing to work with WCs and  $\text{EffectiveWage}_i$  is equal to 100% minus the randomized wage subsidy in the set  $\{0\%; 10\%; 25\%; 50\%; 100\%\}$ , as a pseudo-continuous variable, and  $u_i$  is an error term. We then calculate  $\epsilon^D = \hat{\gamma}_1 \cdot \frac{\bar{w}}{\bar{h}}$ , where  $\bar{w}$  and  $\bar{h}$  are the average effective wage and hiring rates, respectively. The results shown in both panels of Figure I are from a regression model similar to Equation (1), but where we include

<sup>11</sup>Our empirical analysis closely follows our pre-analysis plan, available at [AEARCTR-0005200](https://www.aearctr.org/0005200).

dummies for the different subsidy levels instead of a pseudo-continuous effective wage. Appendix Table B.3 explores robustness to alternate methods of calculating the demand elasticity.

We find that 39% businesses are willing to work with WCs in our baseline case when there is no wage subsidy and businesses must pay 100% of the posted wage. This baseline number of businesses willing to work with WCs is qualitatively similar to other reported estimates, such as the approximately 40% of employers that would “definitely” or “probably” hire WCs in most low-stakes surveys (e.g., Holzer, 2007). The share of businesses willing to work with WCs is also generally increasing in the subsidy level, with 54% of businesses willing to work with WCs when there is a full wage subsidy and no out-of-pocket costs for the business.<sup>12</sup> Estimates of (1) that use information from all of the randomized subsidy levels show that the share of businesses willing to work with WCs increases by approximately 2.1% for every 10% increase in the offered wage subsidy.

Appendix Table B.4 presents descriptive statistics for the businesses that are and are not willing to hire WCs at different subsidy levels to better understand the results from Figure I and the hesitation of many businesses to hire WCs even when there are no out-of-pocket costs. Businesses that are and are not willing to hire WCs are similar in terms of hiring manager experience and business size, but those that are willing to hire WCs are less likely to have a business-wide policy in place regarding WCs and utilize the Platform less frequently compared to business that are not willing to hire WCs. Businesses that are willing to hire WCs are also more likely to say that they want to hire the best candidate regardless of criminal history and that they want to give people a second chance. These businesses are also more confident that WCs will perform well and less concerned that WCs will put others at risk or steal or cause damage while on the job. These patterns generally hold regardless of the subsidy offer, suggesting that the wage subsidies do not substantially change the mix of businesses willing to hire WCs. The remainder of this section will explore how these baseline results change when businesses that use the Platform are having a hard time filling jobs and provide a more formal heterogeneity analysis motivated by our theoretical framework.

### *III.B. Heterogeneity by Labor Market Conditions*

Figure II first explores heterogeneity in the demand for WCs by whether a job is hard to fill on the Platform, a targeted and Platform-specific measure of the labor market conditions for the firm. The first bar plots the fraction of businesses that are willing to work with WCs at baseline from Figure I, where we code businesses as willing to work with WCs only if they responded “Yes.” The second bar instead plots the fraction of businesses that are willing to work with WCs when jobs are hard to fill on the Platform, where we code businesses as willing to work with WCs if they responded “Yes” or “Only if it’s hard to fill my jobs.” Panel A reports these results for the subset of 234 hiring

<sup>12</sup>The point estimates shown in Figure I are slightly non-monotonic over some ranges, with a slightly lower fraction of businesses willing to work with WCs at a wage subsidy of 25% compared to a subsidy of 10%. However, the hiring rates are statistically indistinguishable at the 10% and 25% subsidies, with  $p = 0.60$ . The simplest explanation for these results is that our non-parametric estimates include sampling error due to having only 1/5 of the sample at each subsidy level.

managers from 203 businesses who were randomized to have no wage subsidy, a comparison that is free of possible concerns of treatment effect interactions but may be statistically underpowered. Panel B reports analogous results for the full sample of 1,095 hiring managers from 913 businesses that completed the experiment using a regression specification that measures the mean effect of labor market conditions across all of the wage subsidy levels. We discuss assumptions underlying the full-sample specification in greater detail in Section IV, including potential issues of treatment effect interactions. In both panels, we use a stacked regression model with two observations per person, one for each way of coding the answer of “Only if it’s hard to fill my jobs.”

In our no wage subsidy sample, we find that the share of businesses willing to work with WCs increases by 29 percentage points, to 68%, if the Platform is having a hard time filling a job in our baseline case when there is no wage subsidy. The increase in the share of businesses willing to work with WCs is similar in the full sample, with an average increase of 27 percentage points, to 66%. These results indicate that businesses are more likely to consider non-traditional workers when jobs are hard to fill, consistent with prior work showing that individuals released from prison when local economic conditions are good are less likely to re-offend (e.g., Yang, 2017).

By comparison, we find no economically significant differences in the willingness to work with WCs by the randomized unemployment rate, local unemployment rates for all workers, or local unemployment rates for all workers in the businesses’ industry, as shown in Appendix Figure B.1. There is mixed evidence that businesses are more willing to work with WCs when there is a tight labor market for low-skilled workers, however. We interpret these results as broadly suggesting that measures such as local unemployment rates for all workers do not capture all of the relevant variation in labor market tightness for the businesses in our sample and that more targeted measures are required to accurately understand the importance of labor market conditions on the demand for WCs.

### III.C. Heterogeneity by Job Characteristics

The second and third set of results in Figure II explore heterogeneity in the demand for WCs by whether the job involves high-value inventory and customer interactions, two highly salient characteristics that map to our motivating framework. Panel A again reports results for the subset of hiring managers randomized to have no wage subsidy and Panel B reports results for the full sample. Following our baseline results, we code businesses as willing to work with WCs if they responded “Yes” to the relevant question and unwilling to work with WCs if they responded “No” or “Only if it’s hard to fill my jobs.” We estimate models of the below form:

$$\text{Hire}_i = \alpha_0 \text{No HVI}_i + \alpha_1 \text{HVI}_i + \sum_{k \in K} \lambda_k \cdot \text{Subsidy}_{ik} + e_i \quad (2)$$

where  $\text{No HVI}_i$  is a dummy for not having high-value inventory and  $\text{HVI}_i$  is a dummy for having high-value inventory (and likewise for having or not having customer interactions).  $\text{Subsidy}_{ik}$  is a set of indicator variables for the assigned wage subsidy in the set  $K = \{10\%; 25\%; 50\%; 100\%\}$ .



Note that we do not include an indicator for no wage subsidy and we omit the constant term.

We find in both panels that the share of businesses willing to work with WCs is 17 percentage points higher for jobs that do not involve high-value inventory compared to jobs that do involve high-value inventory. We similarly find that the share of businesses willing to work with WCs is 9 to 12 percentage points higher for jobs that do not involve customer interactions compared to jobs that do. Put another way, this means that the share of businesses willing to work with WCs increases by at least 6 percentage points, to 45%, for jobs that do not involve customer interactions and 12 percentage points, to 51%, for jobs that do not involve high-value inventory.

All of these results are consistent with businesses perceiving greater risks related to customer safety or inventory theft when working with WCs, as suggested by our motivating framework. These results are also consistent with prior work suggesting that employers with jobs that require “trust” are generally less willing to hire WCs (e.g., Holzer, 2007).

#### IV. CRIME AND SAFETY INSURANCE AND TARGETED SCREENING

This section tests several approaches to increasing the demand for WCs that directly address the reasons that businesses may conduct criminal background checks, such as lower average productivity and higher downside risk. We begin by measuring the effects of crime and safety insurance that is meant to address downside risk concerns. We then measure the effects of performance screening and screening based on criminal record history, policies that are meant to address both average productivity and downside risk concerns.

##### IV.A. *Crime and Safety Insurance*

Figure III plots the fraction of businesses that are willing to work with WCs under different counterfactual policies, beginning with a policy where the Platform would provide crime and safety insurance policy that covers damages up to \$5,000. We also plot the baseline level of demand from Figure I as a reference. Panel A reports these results for the subset of 234 hiring managers from 203 businesses who were randomized to have no wage subsidy, again a comparison that is free of possible concerns of treatment effect interactions but may be statistically underpowered. Panel B instead exploits the full sample of 1,095 hiring managers from 913 businesses that completed the experiment to measure the average impact of the insurance policy across the different subsidy levels, using the following stacked regression specification that includes two observations per respondent,  $i$ :

$$\text{Hire}_{ij} = \sum_{l \in L} \lambda_l \cdot \text{PolicyLevel}_{il} + \sum_{k \in K} \lambda_k \cdot \text{Subsidy}_{ik} + e_{ij} \quad (3)$$

where the first observation codes willingness to work with WCs in the baseline case ( $\text{Hire}_{i0}$ ) and the second observation codes willingness to work with WCs under the counterfactual policy ( $\text{Hire}_{i1}$ ).  $\text{PolicyLevel}_{il}$  is a set of indicators that is one for the assigned counterfactual policy level (e.g., insurance caps of \$5,000; \$100,000; \$5m). We include indicators for all possible counterfactual policy

levels including no policy.  $\text{Subsidy}_{ik}$  is a set of indicator variables for the assigned wage subsidy in the set  $K = \{10\%; 25\%; 50\%; 100\%\}$ , and  $e_{ij}$  is an error term. We do not include an indicator for no subsidy and we omit the constant term. We use regressions of this form to create Figure III.

One potential concern with estimating the stacked regression specification in Equation (3) is that it restricts the impact of policies to be the same across subsidy levels, which may lead to bias if there are interaction effects of the wage subsidies and other treatments (Muralidharan et al., 2019). We address this concern in two complementary ways. First, we present results restricting to the subset of hiring managers who were randomized to have no wage subsidy, a comparison that is free of possible concerns of treatment effect interactions but may be statistically underpowered. Our results are generally very similar in the no subsidy and full samples, unless otherwise noted. Second, we present estimates that fully interact the wage subsidies and the other treatments in column 3 of Appendix Tables B.5 and B.6. The point estimates from the no subsidy sample and this fully interacted specification are mechanically identical, but the fully interacted specification allows us to test the statistical significance of the interaction effects. We find no evidence for interaction effects for the crime and safety insurance and job history treatments, although the interactions are statistically significant for the limited screening treatment. The point estimate on the main limited screening effect is not meaningfully different when we include the interactions, however, suggesting that these interaction effects are not driving our findings.<sup>13</sup>

We find that providing even a modest level of crime and safety insurance significantly increases the level of demand for WCs, consistent with concerns about downside risk when hiring WCs. In both panels, we find that insurance that covers damages up to \$5,000 increases demand for WCs by 12 percentage points, to 51%. This 12 percentage point increase is equivalent to the effect of an 80% wage subsidy, based on a linear extrapolation of our baseline estimates from Figure I. We find largely similar effects of insurance coverage at higher amounts in Appendix Figure B.3, with insurance policies that cover damages up to \$100,000 and \$5m both increasing the share of businesses willing to work with WCs by about 16 percentage points from the baseline level of demand. The effects of these insurance coverage at higher amounts are statistically indistinguishable from the effect of insurance coverage at just \$5,000, however, even in the full-sample estimates.

The results from Figure I and Appendix Figure B.3 suggest that businesses are particularly concerned about the types of moderate risks covered by low levels of insurance (e.g., due to petty theft) and, to a less extent, with the more severe tail risk events that are only covered by the higher levels of insurance (e.g., due to violence towards customers or coworkers). Our results for the \$5,000 cap are particularly striking, as the \$5,000 cap is equal to that of the rarely-used U.S. Federal Bonding program. This estimate thus raises the possibility that the Bonding program’s low usage reflects non-demand-based reasons (e.g., that businesses do not know about it or that its use is stigmatized).

<sup>13</sup>We pre-registered that we would estimate both the stacked regression specification in Equation (3) and the fully interacted model described here. We also note that the regression estimates from Equation (3) represent weighted means of treatment effects in regard to various counterfactuals (Muralidharan et al., 2019). For completeness, we also show the entire labor demand curve for each of our main treatments in Appendix Figure B.2.

We can also interpret the magnitude of the insurance estimates using our simple conceptual framework with risk-neutral firms. We observe that median job on the Platform pays \$15 per hour, and our estimates show that firm’s act as though insurance that covers damages up to \$5,000 is roughly equivalent to an 80% wage subsidy. These results imply that firms act as though there is a 1.9% probability that there is a WC-related event that causes \$5,000 in damage in a job-day. Of course, firms may not be fully risk neutral or hiring managers may not have accurate beliefs regarding the risks of working with WCs. We provide some evidence on the impact of inaccurate beliefs below, but caution that it is difficult to adjudicate these different explanations with our data. The contribution of our paper is to carefully measure the impact of different treatments.

#### *IV.B. Screening Based on Performance History*

Figure III next plots the fraction of businesses that are willing to work with WCs if the Platform provides job history screening so that WCs can only accept jobs if they have satisfactorily completed at least one job on the Platform. Panel A reports these results for the subset of hiring managers randomized to have no wage subsidy and Panel B for the full sample of hiring managers using the stacked regression specification described above.

Screening by performance history substantially increases the demand for WCs.<sup>14</sup> In our sample with no wage subsidy, businesses are 11 percentage points more willing to work with WCs if they know that they successfully completed at least one prior job on the Platform, increasing total WC demand to 50%. This 11 percentage point increase is roughly equivalent to the effect of \$5,000 crime and safety insurance or an 80% wage subsidy. The increase in the fraction of businesses willing to work with WCs is similar in the full sample, with an average increase of 14 percentage points, to 53%. We find similar effects of requiring a longer job history in Appendix Figure B.4, with policies requiring that WCs satisfactorily complete five and 25 prior jobs on the Platform increasing the share of businesses willing to work with WCs by 20-22 percentage points and 13-19 percentage points, respectively. The effects of these more restrictive five and 25 job history requirements are statistically indistinguishable from requiring just one successfully completed job, however.

These results suggest that businesses see WCs as heterogeneous in their productivity or downside risk, with just a single positive or negative review providing valuable information on who to hire. These results are consistent with Pallais (2014), who similarly finds that randomly providing a single job’s worth of experience along with a positive review leads to economically large increases in future employment and wages for inexperienced workers on the online platform oDesk.

<sup>14</sup>This finding is consistent with concerns about both productivity and downside risk since such screening could reveal instances of both poor productivity and risky behavior on previous jobs. However, according to accounts from Platform designers and clients, past performance ratings are considered primarily interpreted as a measure of ability to satisfactorily complete the job rather than a measure of bad events that may have occurred on the job.

#### IV.C. Screening Based on Criminal Record History

Figure III finally plots the fraction of businesses that are willing to work with WCs if the Platform provides limited criminal record screening so that WCs can only accept jobs if it has been at least one year since the most recent arrest or conviction. Panel A again reports these results for the subset of hiring managers randomized to have no wage subsidy and Panel B for the full sample of hiring managers using the stacked regression specification described above.

We find that offering businesses the opportunity to screen just the most recent arrests can substantially increase the demand for WCs, again consistent with concerns about both worker productivity and downside risk. In our no wage subsidy sample, screening WCs so that they are only permitted to accept jobs if it has been at least one year since their most recent arrest or conviction increases demand by 21 percentage points, increasing total WC demand to 60%. The increase in the fraction of businesses willing to work with WCs is again similar in the full sample, with an average increase of 20 percentage points, to 59%. Both estimates are considerably larger than the crime and safety and job history estimates and the effects of a 100% wage subsidy. We also find even larger effects of requiring longer periods without an arrest or conviction in Appendix Figure B.5, particularly at the 7-year level where the share of businesses willing to work with WCs increases by a statistically significant 31-41 percentage points from the baseline level of demand.

We further explore the importance of screening on criminal record history in Figure IV, which plots the fraction of businesses that are willing to work with WCs convicted of different types of crime. Businesses are most willing to work with WCs convicted of less serious or drug-related crimes and least willing to work with WCs convicted of more serious and violence-related crimes. In our no wage subsidy sample, for example, only 6% of businesses are willing to work with WCs convicted of a violent felony and only 10% are willing to work with WCs convicted of a violent misdemeanor. By comparison, 27% of businesses are willing to work with WCs convicted of a drug-related felony and 51% are willing to work with WCs convicted of a drug-related misdemeanor.

We also see consistent evidence that businesses prefer WCs with misdemeanor to felony convictions. In the full sample, willingness to hire is much higher for property misdemeanors compared to property felonies (29% vs. 11%,  $p < 0.01$  on the difference), much higher for drug misdemeanors than drug felonies (50% to 28%,  $p < 0.01$ ), and much higher for violent misdemeanors than violent felonies (9% to 5%,  $p = 0.01$ ).

Taken together, the results in Figures III and IV show that targeted screening based on criminal record history can significantly increase the demand for WCs. This finding could reflect the concern that WCs with recent or more serious convictions may have a higher risk of recidivism (e.g., given that the hazard of recidivism is downward-sloping) or that these individuals may be less productive or have higher no-show rates. These results are also broadly aligned with other research findings in this area. For example, Holzer et al. (2007) find that employers report being more willing to hire workers with drug and property convictions compared to other types of convictions, while audit studies show that there are relatively small effects of having a misdemeanor arrest (Uggen et al., 2014) but large effects of any type of both felony drug and property convictions on call-back

rates (Agan and Starr, 2017). We again note, however, that limiting the pool of WCs mechanically reduces the level of demand to zero for the screened out individuals.

## V. OBJECTIVE PERFORMANCE INFORMATION

The final policy we consider is providing hiring managers with objective information about the average performance of WCs on the Platform in an effort to correct mistaken perceptions about the productivity and risk of WCs. We first describe the misperceptions of hiring managers, we then measure the effect of providing objective information on hiring managers’ beliefs, and finally, we examine the effects on hiring decisions.

### *V.A. Correcting Misperceptions in Beliefs*

We collect beliefs about the relative performance of WCs and non-WCs using the incentive-compatible guessing game described above. We elicit beliefs before and after sharing objective information about historical performance on the platform. The reason for collecting prior and posterior beliefs is that not all respondents pay attention to all information presented, and some do not fully update because they think the data shown is not the whole story or credible. More importantly, those who are not shown information can still update their beliefs either by searching for info or making inferences based on how the Platform asked the question. For these reasons, we do not assume the posterior is identical to the truth for those shown info nor completely unchanged for those not shown info.

Panels A and B of Figure V plot the distribution of prior beliefs and posterior beliefs following the provision of the objective information, with both sets of beliefs elicited using the incentive-compatible guessing game. The solid lines show posterior beliefs for the respondents who were shown objective information about WC performance. The dashed lines show the prior beliefs of these same participants. The vertical dash-dotted lines show the true average performance of WCs on the Platform.

We find that prior beliefs about the performance of WCs vary substantially, but on average, hiring managers overestimate the likelihood of a low-performance rating or no-show by 14 percentage points, off a 3 percentage point likelihood, and underestimate the likelihood of a WC earning a five-star rating by 12 percentage points, off an 87% likelihood. These mistaken and pessimistic beliefs are consistent with businesses incorrectly using a criminal record as a signal of lower average productivity and higher average downside risk, creating the potential for more explicit performance information to replace WC status signals.

We find that providing objective information about the average performance of WCs on the Platform led participants to update their beliefs toward the truth, as indicated graphically by the compression, of posterior beliefs around the truth in Figure V. On average, treated participants shifted their beliefs downwards about the likelihood of receiving a no-show or low rating by 5 percentage points and upwards about the likelihood of receiving a five-star rating by 7 percentage

points.

### *V.B. Revisions in Hiring*

We follow Cullen and Perez-Truglia (2018) and estimate the impact of correcting misperceptions in beliefs on the willingness to work with WCs using the following first- and second-stage specifications that allow for the effect of shifting beliefs on behavior even in the presence of updating by the group of participants who are not shown information:

$$\begin{aligned} p_{\text{posterior},i} &= \pi_0 + \pi_1(p_{\text{signal},i} - p_{\text{prior},i}) + \pi_2(p_{\text{signal},i} - p_{\text{prior},i}) * \text{Info}_i + \eta H_{\text{prior},i} + \xi_i \\ H_{\text{posterior},i} &= \beta_0 + \beta_1 \hat{p}_{\text{posterior},i} + \beta_2(p_{\text{signal},i} - p_{\text{prior},i}) + \gamma H_{\text{prior},i} + v_i \end{aligned} \quad (4)$$

where the information shock,  $(p_{\text{signal},i} - p_{\text{prior},i})$ , interacted with the treatment indicator,  $\text{Info}_i$ , is an instrument for hiring managers' posterior beliefs.  $H_{\text{prior},i}, H_{\text{posterior},i} \in \{0, 1\}$  are the hiring manager's prior and posterior willingness to work with WCs, respectively. We express prior and posterior beliefs in log terms throughout this subsection. Following Armantier et al. (2016), Cullen and Perez-Truglia (2018), and Fuster et al. (2018), this log model assumes that the relationship between outcomes and beliefs are linear in log beliefs and symmetric.<sup>15</sup> We estimate Equation (4) separately for respondents assigned to the high- and low-rating treatment arms in the full sample of hiring managers.<sup>16</sup>

Table IV reports regression estimates of the impact of high- and low-performance information on firm beliefs and willingness to work with WCs. Panel A provides results for the high-rating treatment arm while Panel B provides results for the low-rating treatment arm. Column 1 presents first-stage estimates of the effect of information on (log) posterior beliefs, column 2 presents OLS estimates of the relationship between (log) posterior beliefs and hiring decisions, column 3 presents the main IV estimates of the effect of the information treatment on hiring decisions, and column 4 presents reduced form estimates.

The first stage results in column 1, Panel A, show that, on average, treated hiring managers close the gap between their prior beliefs and the truth by 33% more than the control group for the high-performance treatment group. The main IV results in column 3, Panel A, show that the

<sup>15</sup>Appendix Figure B.6 reports binned scatter-plot estimates of the impact of high- and low-performance information on business beliefs and hiring decisions to assess the log-linearity assumption and illustrate the variation underlying specification (4). Panels A and B plot the difference between the reported performance beliefs at the end of the experiment and prior beliefs against the perception gap for the low- and high-performance treatments, respectively. These results show that treated hiring managers, by and large, eliminated nearly all of the error in their initial guesses about WC performance. Control hiring managers also partially eliminated the error in their initial guesses. We do not expect this partial updating to bias our IV estimates given our direct measures of posterior beliefs for all participants. Panels C and D plot the willingness to work with WCs at the end of the experiment against the fitted posterior belief predictions from the first stage regression, again for the low- and high-performance treatments, respectively. The results graphically corroborate the assumed linear relationship between (randomly) shocked beliefs about WC performance and hiring demand.

<sup>16</sup>Estimates of Equation (4) are qualitatively similar but much less precisely estimated in the no subsidy sample. For example, we estimate an IV coefficient of 0.818 (s.e. = 0.414) in the full sample for the high-performance treatment, compared to 0.494 (s.e. = 0.655) in the no subsidy sample. For the low-performance treatment, we estimate IV coefficients of 0.067 (s.e. = 0.074) in the full sample and 0.161 (s.e. = 0.152) in the no subsidy sample.



elasticity of hiring with respect to high-performance beliefs about WCs is 0.8, meaning that a 10% increase in managers’ beliefs about WCs’ performance leads to an 8 percentage point increase in willingness to work with WCs. Given that treated participants shifted their beliefs upwards about the likelihood of receiving a five-star rating by 7 percentage points, this IV estimate implies that providing objective information on the true share of high-performance ratings increased the share of firms willing to hire WCs by about 6 percentage points on average, roughly equivalent to the effect of an 40% wage subsidy.

In contrast, the IV results in column 3 of Panel B of Table IV shows that changing perceptions in the low-performance group has no impact on hiring decisions. This null result is despite column 1, Panel B, showing that treated hiring managers close the gap between their prior beliefs and the truth by 45% more than the control group. Our interpretation of these results is that the share of low ratings or no-shows is less salient and less relevant for WC hiring decisions. Consistent with this interpretation, hiring managers also have more dispersed priors about low ratings and no-shows at baseline.

## VI. THREATS TO VALIDITY

In this section, we describe how the details of the experimental design and setting may affect the interpretation of our results.

### VI.A. *Social Desirability Bias*

One important consideration is that hiring managers may express interest in hiring WCs out of a desire to appear socially conscious. The *ex-ante* incentive-compatible structure of our experiment and the fact that hiring managers were not aware that they were part of a research study directly addresses this concern. Our study is based on businesses making real, high-stakes choices. From a participating hiring manager’s perspective, the Platform—to whom they had ceded discretion over circulating their posted jobs—was asking direct questions about whether their business would allow WCs to accept their jobs. The hiring managers were also not aware that they were part of a research study at any point during the outreach or experiment. There is therefore no reason for the hiring managers to express interest in hiring WCs other than if they are actually interested in allowing WCs to match to their jobs.

### VI.B. *Screening Expectations*

A second consideration is that hiring managers could have mistakenly assumed that WCs on the Platform would be pre-screened in some way. This concern is alleviated by the fact that we asked the direct question about whether businesses would allow WCs to perform their jobs twice, once at the start of the experiment and again at the end of the experiment. In between these two questions, hiring managers were asked to consider WCs who were convicted of specific crimes, ranging from drug-related misdemeanors to violent felonies. As a result, it was likely very clear that the pool of

WCs included individuals convicted of more serious crimes. Yet, 85% of hiring managers that did not receive performance information gave the same answer to the direct questions about whether businesses would allow WCs to perform their jobs. The consistency of the answers at the start and end of the experiment suggests that mistaken beliefs about screening cannot explain our results, as well as any concerns about measurement error from participants who may not have been paying full attention.

### *VI.C. Multiple Hypothesis Testing*

A third consideration is that we are detecting false positives due to multiple hypothesis-testing, i.e., that many of our results are statistically significant due to chance alone. To control for the family-wise error rate, defined as the probability of making one or more false discoveries when performing multiple hypothesis tests, we employ a step-down algorithm similar to those described by Westfall and Young (1993). For a given family of  $k$ -hypothesis tests, the algorithm uses a step-down procedure that employ permutation calculations to estimate dependence relationships and provide corrected p-values. Westfall-Young adjusted p-values calculated using 10,000 bootstrap samples are reported in brackets in Appendix Tables B.5 and B.6 for the effect of each level of each policy (e.g., \$5k crime and safety insurance or one previous job) as compared to the baseline or to an omitted group. In both tables, column 1 reports the adjusted p-values for the sample with no wage subsidy and column 2 reports the adjusted p-values for the mean effect across all wage subsidy levels in the full sample. The adjusted p-values are calculated based on families of hypotheses grouped by each column in each figure. For the adjusted p-values in Appendix Table B.5, the family of nine hypotheses include three for the mean effect of crime and safety insurance, three for the mean effect of performance history, and three for the mean effect of screening by years since arrest or conviction. For the adjusted p-values in Appendix Table B.6, the family of five hypotheses are for the effect of restrictions based on crime type relative to the effect of restricting to WCs with a violent felony conviction. In the smaller no subsidy sample, the adjusted p-values for crime and safety insurance relative to the baseline, for the job history screening relative to the baseline, for the one and three year limited record screening relative to the baseline, and for violent misdemeanor relative to violent felony are all greater than 0.05 due to the small number of observations. By comparison, our main findings are qualitatively unchanged in the full sample. The adjusted p-values for the mean effect of \$5,000 in crime and safety insurance relative to the baseline and the mean effect of requiring the completion of at least one previous job relative to the baseline are 0.055, while the adjusted p-values for the other mean effects relative to the baseline remain below 0.05.

### *VI.D. COVID-19 Pandemic*

A fourth consideration is that we conducted the experiment in March and April, 2020, at the start of the COVID-19 pandemic. The onset of the pandemic may have impacted the willingness of firms to work with WCs by, for example, leading to more stringent hiring standards for temporary workers due to the increased uncertainty. To explore how the onset of the pandemic may have impacted our

findings, Appendix Figure B.7 shows baseline results for hiring managers (1) responding before and after the national declaration of state of emergency on March 13, 2020; and (2) located in counties that were less and more exposed to the first wave of the pandemic. Our results are qualitatively unchanged across these groups, with nearly identical levels of demand for WCs in all specifications. None of the estimates suggest that our findings are driven by the timing of the experiment and the COVID-19 pandemic. In unreported results, we also find no statistically significant differences in our main treatment effects across these groups. In addition, our sample is not especially overrepresented or underrepresented in regard to the industries that were hit hardest at the start of the pandemic (Bartik et al., 2020).

### *VI.E. External Validity*

A final and more general consideration is the external validity of our estimates to other settings. The setting for our study is a leading online labor platform that thousands of traditional businesses use to source workers for temporary staffing across a wide range of entry-level roles. We expected this setting to offer a large and concentrated pool of appropriate jobs for WCs re-entering the workplace. When extrapolating to other settings, it is important to keep in mind that the demand for WCs may be very different for permanent positions or more senior roles. While we believe the jobs offered through the Platform often convert to longer term work based on the high share of repeat hires (e.g., over 54% of workers paired with a business once through the Platform will return to a job at the same business), the Platform does not track long-term outcomes. We also cannot speak to the evolution of demand after businesses gain experience working with WCs or over the business cycle. Finally, our finding that the level of customer interactions and presence of high-value inventory affect demand for WCs suggests that role-specific traits are meaningfully correlated with demand and must be taken into consideration when extrapolating to other settings.

To partially explore these issues, Appendix Tables B.7 and B.8 show treatment effects relative to the baseline where we (1) restrict our sample to firms with information on firm size, industry type, and WC hiring policies; (2) weight observations to more closely match the distribution of answers to nationwide responses to SHRM questions about attitudes and policies in place with respect to WCs; and (3) weight observations to more closely match the distribution of firms in the U.S. economy based on industry and firm size. We weight observations using the iterative proportional fitting (IPF) algorithm following Deming and Stephan (1940) to adjust for consistency with the marginal distributions of WC hiring policies and then industry shares and firm sizes. The weights were calculated through stepwise adjustment and was repeated for 50 iterations using the implementation of the IPF algorithm developed by Bergmann (2011). Our results are qualitatively unchanged across these conditions, with nearly identical levels of demand in all specifications.

## VII. CONCLUSION

This paper uses information from a discrete choice field experiment on a nationwide staffing platform to test several approaches to increasing the demand for WCs, each of which is intended to directly address a potential underlying reason that employers choose to conduct criminal background checks. We find that 39% of businesses on the Platform are willing to work with WCs at baseline, with higher levels of demand for jobs that do not involve customer interactions or high-value inventory and when the Platform is having a hard time filling a job. The level of demand also increases to 50% or higher when businesses are offered a modest level of crime and safety insurance, a single performance review, or screening of the most recent criminal records. All of our results suggest that policymakers may affect WC demand by directly addressing the underlying reasons that employers choose to conduct background checks, rather than simply prohibiting or delaying questions about job applicants' arrest and conviction record during the hiring process.

An important open question is whether these alternative approaches are more cost-effective than wage subsidies, which can achieve similar gains at high enough subsidy levels. While a comprehensive cost comparison is beyond the scope of this paper, we can calculate the direct costs of increasing the demand for WCs for each of our main treatments under reasonable assumptions. These calculations reveal that all of these policies can significantly increase demand for WCs at a fraction of the cost of wage subsidies. Performance screening, for example, can achieve notable gains in the share of businesses willing to work with WCs at near-zero cost because a large number of businesses are willing to work with and provide WCs with their first performance review, opening the door to businesses that highly value that first positive review. Providing objective information on the average productivity of WCs can similarly increase the share of businesses hiring WCs at essentially zero additional cost to the Platform. Revising background check matrices to only exclude candidates with the most recent criminal records requires no new costs for the Platform. Finally, we calculate that crime and safety insurance can increase the demand for WCs at  $\frac{1}{2}$  to  $\frac{1}{10}$  the cost of wage subsidies under plausible assumptions of the probability of damages due to WC misbehavior.<sup>17</sup> These calculations suggest that all of the options we consider are substantially more cost-effective than wage subsidies, at least in this context.

A second important but more speculative question is whether the treatments we consider can substantially increase the number of WCs employed nationwide. There are two challenges to answering this question with our study. First, we would need to extrapolate the results from our experimental sample to the country as a whole. Second, the unit of observation in our experiment is a business, not a worker. The impact of the business-level changes we study on WC employment depends on several additional factors, including the total demand for WCs at each firm, how long each job lasts, and the supply of WCs for each type of job. We view the careful modeling and

<sup>17</sup>For example, increasing the number of businesses willing to work with WCs by approximately 10% would require a 50% wage subsidy. Using a typical Platform wage of \$15 per hour and an 8-hour work day, the subsidy approach would thus cost \$60 per worker per day. Providing a \$5,000 crime and safety insurance policy could also increase WC demand by approximately 10%. Assuming that WCs have either a one in 1,000 or one in 200 daily chance of incurring \$5,000 in damages, this insurance policy would thus have an expected cost of \$5 to \$25 per worker per day.

estimating such models of labor supply and demand is an important avenue for future work.

Based on the findings from our study, the Platform is changing its user interface nationwide. Businesses that join after the close of our experiment will also have the option to allow WCs accept their jobs, with crime and safety insurance coverage provided by the Platform. To date, demand from our study participants combined with the permanent policy changes made following the result of our experiment led to over 12,000 jobs being made available to WCs through August 2021. This rapid expansion in the number of jobs available to WCs opens new questions for future research, including the evolution of demand as businesses gain experience working with WCs and the long-term employment opportunities created for WCs that accept jobs on the Platform. While we cannot examine the impact of policies such as Ban the Box in our context, it would also be interesting for future work to compare such policies against the ones that we consider here.

HARVARD BUSINESS SCHOOL AND NBER

HARVARD KENNEDY SCHOOL, J-PAL, AND NBER

U. TORONTO ROTMAN SCHOOL OF MANAGEMENT, NBER, AND CEPR

## REFERENCES

- Agan, A. and S. Starr (2017). The Effect of Criminal Records on Access to Employment. *American Economic Review: Papers & Proceedings* 107(5), 560–564.
- Agan, A. and S. Starr (2018). Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment. *Quarterly Journal of Economics* 133(1), 191–235.
- Allcott, H. (2011). Social Norms and Energy Conservation. *Journal of Public Economics* 95(9-10), 1082–1095.
- Armantier, O., S. Nelson, G. Topa, W. van der Klaauw, and B. Zafar (2016). The Price Is Right: Updating Inflation Expectations in a Randomized Price Information Experiment. *Review of Economics and Statistics* 98(3), 503–523.
- Bartik, A. W., M. Bertrand, Z. Cullen, E. L. Glaeser, M. Luca, and C. Stanton (2020). The impact of covid-19 on small business outcomes and expectations. *Proceedings of the National Academy of Sciences* 117(30), 17656–17666.
- Bergmann, M. (2011). IPFWEIGHT: Stata Module to Create Adjustment Weights for Surveys. Statistical Software Components, Boston College Department of Economics.
- Bollinger, C. R. and J. L. Hotchkiss (2003). The Upside Potential of Hiring Risky Workers: Evidence from the Baseball Industry. *Journal of Labor Economics* 21(4), 923–944.
- Brandts, J. and G. Charness (2011). The Strategy Versus the Direct-Response Method: A First Survey of Experimental Comparisons. *Experimental Economics* 14(3), 375–398.
- Bursztyn, L., A. L. Gonzalez, and D. Yanagizawa-Drott (2020). Misperceived Social Norms: Female Labor Force Participation in Saudi Arabia. *American Economic Review* 110(10), 2997–3029.
- Bushway, S. D. (2004). Labor Market Effects of Permitting Employer Access to Criminal History Records. *Journal of Contemporary Criminal Justice* 20(3), 276–291.


- Carson, R. T. and T. Groves (2007). Incentive and Informational Properties of Preference Questions. *Environmental and Resource Economics* 37(1), 181–210.
- Charness, G., U. Gneezy, and B. Halladay (2016). Experimental Methods: Pay One or Pay All. *Journal of Economic Behavior & Organization* 131, 141–150.
- Couloute, L. and D. Kopf (2018). Out of Prison & Out of Work. Prison Policy Initiative Report.
- Craigie, T. A. (2020). Ban the Box, Convictions, and Public Employment. *Economic Inquiry* 58(1), 425–445.
- Cullen, Z. and R. Perez-Truglia (2018). How Much Does Your Boss Make? The Effects of Salary Comparisons. *NBER Working Paper No. 24841*.
- Deming, W. E. and F. F. Stephan (1940). On a Least Squares Adjustment of a Sampled Frequency Table When the Expected Marginal Totals are Known. *The Annals of Mathematical Statistics* 11(4), 427–444.
- Doleac, J. L. and B. Hansen (2020). The Unintended Consequences of “Ban the Box”: Statistical Discrimination and Employment Outcomes When Criminal Histories Are Hidden. *Journal of Labor Economics* 38(2), 321–374.
- Fuster, A., R. Perez-Truglia, M. Wiederholt, and B. Zafar (2018). Expectations with Endogenous Information Acquisition: An Experimental Investigation. *NBER Working Paper No. 24767*.
- Holzer, H. J. (2007). Collateral Costs: The Effects of Incarceration on the Employment and Earnings of Young Workers. *IZA Discussion Paper No. 3118*.
- Holzer, H. J., S. Raphael, and M. A. Stoll (2004). How Willing are Employers to Hire Ex-Offenders? *Focus* 23(2), 40–43.
- Holzer, H. J., S. Raphael, and M. A. Stoll (2006). Perceived Criminality, Criminal Background Checks, and the Racial Hiring Practices of Employers. *The Journal of Law and Economics* 49(2), 451–480.
- Holzer, H. J., S. Raphael, and M. A. Stoll (2007). The Effect of an Applicant’s Criminal History on Employer Hiring Decisions and Screening Practices: Evidence from Los Angeles. In S. D. Bushway, M. A. Stoll, and D. F. Weiman (Eds.), *Barriers to Reentry? The Labor Market for Released Prisoners in Post-Industrial America*. New York: Russell Sage Foundation.
- Hunt, P., R. Smart, L. Jonsson, and F. Tsang (2018). Breaking Down Barriers Experiments into Policies That Might Incentivize Employers to Hire Ex-Offenders. *RAND Research Report No. 2142*.
- Infogroup (2016). Infogroup U.S. Historical Business Data.
- Jackson, O. and B. Zhao (2017). The Effect of Changing Employers’ Access to Criminal Histories on Ex-Offenders’ Labor Market Outcomes: Evidence from the 2010–2012 Massachusetts CORI Reform. *Federal Reserve Bank of Boston Research Department Working Papers No. 16–30*.
- Kessler, J. B., C. Low, and C. Sullivan (2019). Incentivized Resume Rating: Eliciting Employer Preferences without Deception. *American Economic Review* 109(11), 3713–3744.
- Lazear, E. P. (1998). Hiring Risky Workers. In *Internal Labour Markets, Incentives and Employment*,



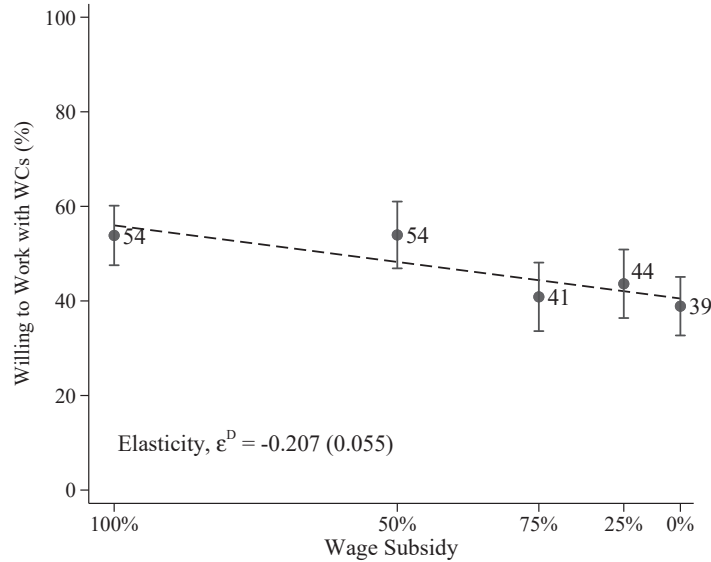
- pp. 143–158. London: Palgrave Macmillan.
- Low, C. (2017). A “Reproductive Capital” Model of Marriage Market Matching. *Manuscript, Wharton School of Business*.
- Mas, A. and A. Pallais (2017). Valuing Alternative Work Arrangements. *American Economic Review* 107(12), 3722–3759.
- Minor, D., N. Persico, and D. M. Weiss (2018). Criminal Background and Job Performance. *IZA Journal of Labor Policy* 7(1), 1–49.
- Muralidharan, K., M. Romero, and K. Wüthrich (2019). Factorial Designs, Model Selection, and (Incorrect) Inference in Randomized Experiments. *NBER Working Paper No. 26562*.
- Oyer, P. and S. Schaefer (2011). Personnel Economics: Hiring and Incentives. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics* (1 ed.), Volume 4B, Chapter 20, pp. 1769–1823. Elsevier.
- Pager, D. (2003). The Mark of a Criminal Record. *American Journal of Sociology* 108(5), 937–975.
- Pallais, A. (2014). Inefficient Hiring in Entry-Level Labor Markets. *American Economic Review* 104(11), 3565–3599.
- Raphael, S. (2011). Improving Employment Prospects for Former Prison Inmates: Challenges and Policy. In *Controlling Crime: Strategies and Tradeoffs*, pp. 521–565. University of Chicago Press.
- Rose, E. (2021). Does Banning the Box Help Ex-Offenders Get Jobs? Evaluating the Effects of a Prominent Example. *Journal of Labor Economics* 39(1), 79–113.
- Sabia, J. J., T. Mackay, T. T. Nguyen, and D. M. Dave (2020). Do Ban the Box Laws Increase Crime? *NBER Working Paper No. 24381*.
- Schnepel, K. T. (2018). Good Jobs and Recidivism. *Economic Journal* 128(608), 447–469.
- Sherrard, R. (2021). “Ban the Box” Policies and Criminal Recidivism. *Manuscript*.
- Shoag, D. and S. Veuger (2016). No Woman No Crime: Ban the Box, Employment, and Upskilling. *HKS Working Paper No. 16-015*.
- Society for Human Resource Management (2018). Workers with Criminal Records: A Survey by the Society for Human Resource Management and the Charles Koch Institute. Report.
- Stoll, M. A. (2009). Ex-Offenders, Criminal Background Checks, and Racial Consequences in the Labor Market. *University of Chicago Legal Forum* 2009(1), 381–419.
- Uggen, C., M. Vuolo, S. Lageson, E. Ruhland, and H. K. Whitham (2014). The Edge of Stigma: An Experimental Audit of the Effects of Low-Level Criminal Records on Employment. *Criminology* 52(4), 627–654.
- Westfall, P. H. and S. S. Young (1993). *Resampling-Based Multiple Testing: Examples and Methods for p-Value Adjustment*, Volume 279. John Wiley & Sons.
- Yang, C. S. (2017). Local Labor Markets and Criminal Recidivism. *Journal of Public Economics* 147, 16–29.

FIGURE I  
Labor Demand for Workers with a Criminal Record

A. Effect of Wage Subsidies




B. Labor Demand Curve



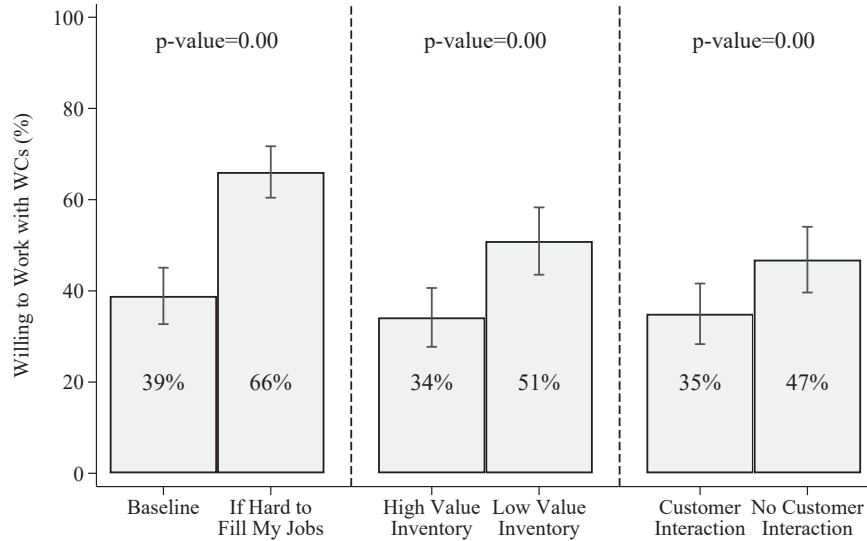
*Notes.* This figure presents estimates of the mean willingness to work with WCs (and 95% confidence interval) by randomized wage subsidy. The estimates are based on the 1,095 hiring managers from 913 businesses that completed the experiment. Panel A reports the fraction of respondents answering “Yes” when asked if they are willing to work with WCs at a given randomized wage subsidy. We estimate a regression model similar to Equation (1), but where we include dummies for the different subsidy levels instead of a pseudo-continuous effective wage. Each p-value corresponds to the test of the null hypothesis that there is no difference in demand between a given subsidy level and the baseline condition (e.g., no difference between Baseline and 10% Subsidy). Panel B reports the same estimates as a labor demand curve along with the OLS line of best fit and average labor demand elasticity. The p-values and 95% confidence intervals are calculated using robust standard errors clustered by business.

FIGURE II  
Heterogeneity by Labor Market Conditions and Job Characteristics

A. Effect in the No Wage Subsidy Sample




B. Mean Effect in the Full Sample




*Notes.* This figure presents estimates of the mean willingness to work with WCs (and 95% confidence interval) by labor market conditions and job characteristics. The “If Hard to Fill My Jobs” column reports estimates from a stacked regression specification similar to Equation (3) for the fraction of respondents who answer “Yes” or “Only if it’s hard to fill my jobs” when asked if they are willing to work with a WC. All other estimates report the estimates of the fraction of respondents who answer “Yes” to the same question from regression specifications described by Equation (2). Each p-value corresponds to the test of the null hypothesis that there is no difference in demand by labor market conditions or job characteristic. Panel A reports results for the sample of 234 hiring managers from 203 businesses in the no wage subsidy sample. Panel B reports results for the full sample of 234 hiring managers from 203 businesses that completed the experiment. The p-values and 95% confidence intervals are calculated using robust standard errors clustered by business. Appendix Table B.1 provides details on the coding of job characteristics.

FIGURE III  
 Crime and Safety Insurance, Job History Screening, and Limited Criminal Record Screening

A. Effect in the No Wage Subsidy Sample



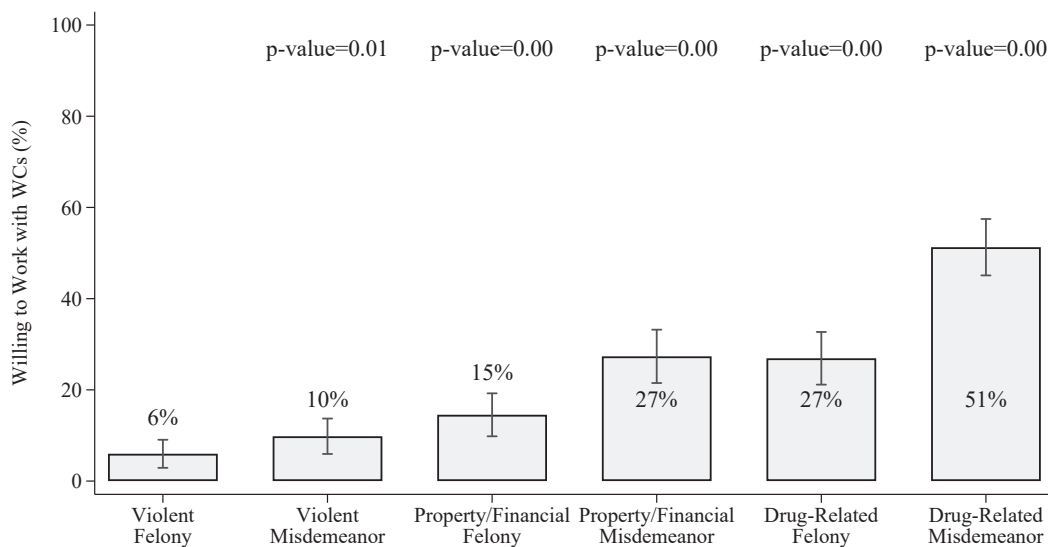
B. Mean Effect in the Full Sample



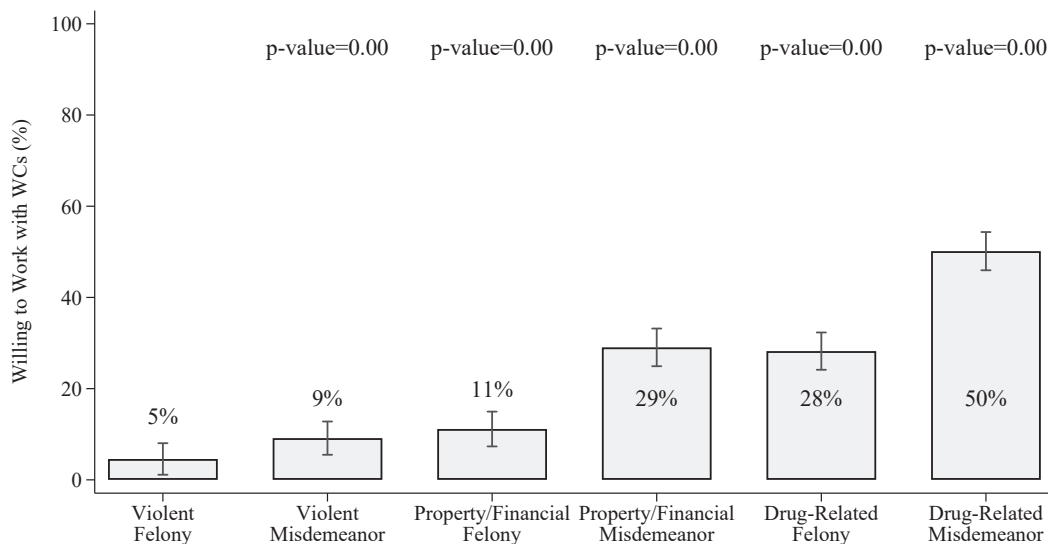
*Notes.* This figure presents estimates of the mean willingness to work with WCs (and 95% confidence interval) if the Platform provides crime and safety insurance, job history screening, or limited criminal record screening. The estimates are based on results from regression specifications described by Equation (3). Each p-value corresponds to the test of the null hypothesis that each treatment has no effect relative to the baseline. Panel A reports results for the sample of 234 hiring managers from 203 businesses in the no wage subsidy sample. Panel B reports results for the full sample of 1,095 hiring managers from 913 businesses that completed the experiment. The first bar in both panels reports the fraction of respondents willing to work with a WC at baseline, the second if the Platform provides crime and safety insurance policy that covers damages up to \$5,000, the third if the Platform provides job history screening so that WCs can only accept jobs if they have satisfactorily completed at least one job on the Platform, and the fourth if the Platform provides limited criminal record screening so that WCs can only accept jobs if it has been at least one year since the most recent arrest or conviction. The estimates are presented as effects relative to the baseline in columns 1 and 2 of Appendix Table B.5. The p-values and 95% confidence intervals calculated using robust standard errors clustered by business.

FIGURE IV  
Criminal Record Screening by Conviction Type

A. Effect in the No Wage Subsidy Sample




B. Mean Effect in the Full Sample



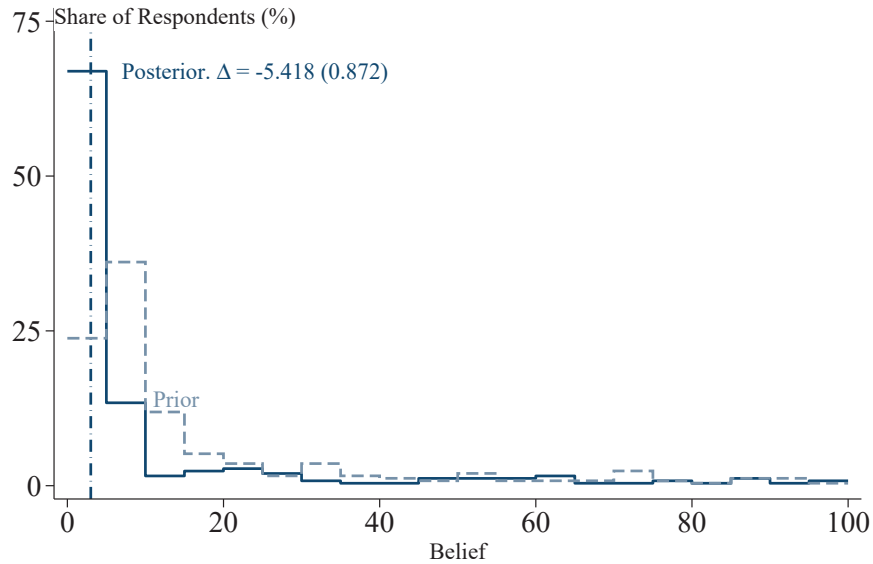
*Notes.* This figure presents estimates of the mean willingness to work with WCs (and 95% confidence interval) by conviction type. The estimates are based on results from regression specifications described by Equation (3). Each p-value corresponds to the test of the null hypothesis that there is no difference in demand for a given type of conviction type relative to violent felony. Panel A reports results for the sample of 234 hiring managers from 203 businesses in the no wage subsidy sample. Panel B reports results for the full sample of 1,095 hiring managers from 913 businesses that completed the experiment. The estimates are presented as effects relative to the willingness to work with WCs who have a violent felony conviction in columns 1 and 2 of Appendix Table B.6. The p-values and 95% confidence intervals calculated using robust standard errors clustered by business.

FIGURE V  
Effects of Receiving Information on Posterior Beliefs

A. High-Performance



B. Low-Performance



*Notes.* This figure reports the prior and posterior distributions of business beliefs about WC productivity for the full sample of respondents. Panel A reports the distribution of prior and posterior beliefs about the share of WCs who receive high-performance ratings (five-star ratings). Panel B reports the distribution of prior and posterior beliefs about the share of WCs who receive low performance ratings (no-shows and either one- or two-star ratings). The dotted lines present the distribution of prior beliefs. The solid lines present the distribution of posterior beliefs. The  $\Delta$  is the impact of the treatment on beliefs, with the standard error of this estimate in parentheses.



TABLE I  
Description of Main Treatments

Treatment Name	Survey Question	Values
Wage Subsidy	If the [Platform] gave you a {Wage Subsidy} discount for [Platform Workers] with a criminal record, would you permit such [Workers] to perform jobs you post? This means you would only pay (100 - {Wage Subsidy}) of the wage for those with a criminal record.	0%; 5%; 10%; 25%; 50%; 100%
Crime and Safety Insurance	If the [Platform] could cover damages up to {Crime and Safety Insurance Cap} related to theft or safety incurred by workers with a criminal record, would you permit such [Workers] to perform jobs you post?	\$1k; \$5k; \$100k; \$5m
Performance History	If the [Platform] required [Platform Workers] with a criminal record to have satisfactorily completed {Performance History} job(s), receiving >85% positive reviews (5 stars), would you permit such [Workers] to perform jobs you post?	1 job; 5 jobs; 25 jobs
Clean Record Length	If the [Platform] required users with a criminal record to have maintained a clean record for at least {Clean Record Length} would you permit such users to perform jobs you post?	1 year; 3 years; 7 years
Conviction Type	Please indicate whether you would permit [Platform Workers] with these types of convictions to perform jobs you post. The [Platform] would still give you a {Wage Subsidy} discount, but no other supplementary policies would apply.	Violent Felony/Misd; Prop/Financial Felony/Misd; Drug-Related Felony/Misd
Objective Information	The truth is that {Share}% of jobs completed by people with a criminal record resulted in a {Rating} on the same or a similar platform—actually better than everyone else.	13% share for low rating; 87% share for high rating

*Notes.* This table summarizes the main experimental treatments. The text in square brackets is redacted information identifying the Platform. The text in curly brackets is a placeholder for the randomized values of each treatment.

TABLE II  
Descriptive Statistics

A. Firm Characteristics	Experimental Sample	Infogroup Database
Median Firm Age	19.0	16.0
Median Number of Employees	40	2.5
Service	0.31	0.37
Manufacturing	0.19	0.06
Retail	0.15	0.21
Public Administration	0.10	0.02
Transportation & Public Utilities	0.10	0.03
Wholesale Trade	0.09	0.08
Finance, Insurance, & Real Estate	0.03	0.07
Construction	0.01	0.08
Nonclassifiable	0.01	0.08
Firms with Nonmissing Age or Number Employees	666	3,260,733
with Nonmissing Industry Classification	518	1,245,145
B. Hiring Policies and Views	Experimental Sample	SHRM Survey
Firm-Wide WC Hiring Policy	0.45	0.66
Consider WCs Because Best Candidate	0.46	0.53
Consider WCs Because Second Chances Are Important	0.50	0.38
Consider WCs Because of Financial Incentives	0.08	0.02
Concerned About Customer Reactions	0.49	0.30
Concerned About Local, State, or Federal Regulations	0.26	0.22
Concerned About Performance	0.15	0.04
Firms with Nonmissing Hiring Policy Information	900	1,228

*Notes.* This table reports descriptive statistics for the experimental sample comprised of the 1,095 hiring managers from 913 businesses that completed the experiment. Panel A reports statistics for the 913 firms in our sample matched to the Infogroup Historical Business Database (column 1) and all firms in the Infogroup Historical Business Database (column 2), which contains basic profile data for more than a million U.S. businesses. The industry characteristics are further limited to the 518 firms in our sample with that data available in the Infogroup Database. Panel B reports information on WC hiring policies, where information for the broader set of U.S. businesses comes from a nationwide survey of over 1,000 HR professionals commissioned by the Society for Human Resource Management.

TABLE III  
Randomization Assessment  
p-values from Regressions of Covariates on Treatment Indicators

	Wage Subsidy	Crime Insurance	Performance History	Clean Record	Unemp. Rate	Shown Info.
A. Firm Characteristics and Policies						
Firm Age	0.291	0.646	0.407	0.268	0.271	0.347
Employees	0.257	0.858	0.099	0.009	0.613	0.424
Service	0.240	0.956	0.576	0.287	0.286	0.711
Manufacturing	0.045	0.277	0.949	0.877	0.269	0.386
Retail	0.393	0.138	0.036	0.684	0.873	0.231
Transportation & Public Utilities	0.691	0.908	0.625	0.434	0.765	0.988
Nonclassifiable & Misc. Industries	0.494	0.454	0.697	0.552	0.937	0.261
Firm-Wide WC Hiring Policy	0.779	0.673	0.728	0.513	0.491	0.229
Platform Tenure (Years)	0.099	0.075	0.602	0.518	0.224	0.947
Job Vacancy Rate	0.857	0.569	0.151	0.928	0.733	0.913
B. Firm Characteristics and Policies						
Job Involves Customer Interactions	0.675	0.595	0.429	0.356	0.628	0.710
Job Involves High-Value Inventory	0.277	0.450	0.242	0.531	0.467	0.437
Modal Job is Fulfillment / Warehousing	0.767	0.325	0.559	0.474	0.284	0.052
Modal Job is General Labor	0.523	0.085	0.499	0.258	0.231	0.117
Modal Job is Event Staff	0.341	0.950	0.762	0.457	0.718	0.559
Modal Job is Delivery	0.965	0.097	0.324	0.964	0.726	0.808
Modal Job is Washing & Cleaning	0.940	0.732	0.738	0.051	0.726	0.174
Firms	913	913	913	913	913	913
Managers	1,095	1,095	1,095	1,095	1,095	1,095

*Notes.* This table reports balance tests for the estimation sample described in Table II. Each cell reports the p-value of an F-statistic from a separate regression of the baseline covariates listed in the rows on indicator variables for each value of the treatments listed in the columns. Standard errors are clustered at the firm level. Nonclassifiable & Misc. Industries is an aggregation of Nonclassifiable, Construction, Finance, Public Administration, and Wholesale Trade industries. See the Table II notes for additional details on the outcomes and sample.

TABLE IV  
High- and Low-Performance Information

	First Stage	OLS	IV	Reduced Form
<b>A. Impact of High-Performance Information</b>				
Shown Info $\times$ (Signal - Prior Belief)	0.334 (0.0805)			0.273 (0.131)
ln(Posterior Belief)		0.338 (0.108)	0.818 (0.414)	
Mean: Dependent Variable	4.36	0.51	0.51	0.51
Kleibergen-Paap: Weak Identification F-Stat			17.25	
Firms	490	490	490	490
Managers	558	558	558	558
<b>B. Impact of Low-Performance Information</b>				
Shown Info $\times$ (Signal - Prior Belief)	0.447 (0.0474)			0.0302 (0.0327)
ln(Posterior Belief)		-0.0348 (0.0294)	0.0676 (0.0740)	
Mean: Dependent Variable	1.68	0.51	0.51	0.51
Kleibergen-Paap: Weak Identification F-Stat			88.85	
Firms	430	430	430	430
Managers	502	502	502	502

*Notes.* This table reports estimates of the impact of high- and low-performance information on firm beliefs and willingness to work with WCs. Panel A reports results for managers who were shown information on the fraction of five-star ratings. Panel B reports results for managers who were shown information on the fraction of no-shows and either one- or two-star ratings. Column 1 reports first stage estimates of the effect of information on posterior beliefs. Column 2 reports OLS estimates of the cross-sectional relationship between posterior beliefs and willingness to work with WCs. Column 3 reports IV estimates of the causal impact of a change in posterior beliefs on willingness to work with WCs. Column 4 reports reduced form estimates of the effect of information on willingness to work with WCs. See Section V of the text for additional details.