

INCREASING THE DEMAND FOR WORKERS WITH A CRIMINAL RECORD*

ZOË CULLEN
WILL DOBBIE
MITCHELL HOFFMAN

JULY 2022

We experimentally test several approaches to increasing the demand for workers with a criminal record on a nationwide staffing platform by addressing potential downside risk and productivity concerns. The staffing platform asked hiring managers to make a series of hypothetical hiring decisions that impacted whether workers with a criminal record could accept their jobs in the future. 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 a 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, Raphael, and Stoll 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 policies that delay questions about a job applicant’s arrest

*We are grateful to four anonymous referees and Lawrence Katz and Andrei Shleifer, whose guidance greatly improved this paper. We also thank Amanda Agan, Paul Goldsmith-Pinkham, Sara Heller, Judd Kessler, Eddie Lazear, Melvin Stephens, Crystal Yang, and numerous seminar participants for helpful comments and suggestions. Dylan Balla-Elliott, Kenny Gu, 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.

and conviction record. These policies are meant to increase hiring and employment 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, Ban the Box policies do not address the underlying reasons that employers may conduct criminal background checks, such as the potential for downside risk or lower productivity. 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 (Bushway 2004; Holzer, Raphael, and Stoll 2006; Stoll 2009; Agan and Starr 2018; Doleac and Hansen 2020). Perhaps as a result, existing work shows mixed effects of Ban the Box policies on the employment of WCs (Jackson and Zhao 2017; Craigie 2020; Rose 2021) and larger populations that include both WCs and non-WCs (Shoag and Veuger 2016; Doleac and Hansen 2020).

In this paper, we use a field experiment to test several approaches to increasing the demand for WCs that more directly address the underlying reasons that employers 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 risk and productivity concerns. We also provide objective information on the average performance of WCs and non-WCs to address 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 approaches to 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, the pay for the job, and qualifying criteria. The Platform sends out the job offer to workers who meet the qualifying criteria, then workers can accept the job on a first-come, first-serve basis. Businesses rarely cancel jobs after they have been accepted by workers due to the 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.¹ Hiring managers were already familiar with submitting qualifying criteria for workers who could accept their jobs. The Platform truthfully informed these managers that, in the same way, the hypothetical hiring decisions made during the experiment could impact whether WCs could accept their jobs in the future. For example, WCs could accept a business’s jobs if the hiring manager indicated that they would be willing to work with WCs under a particular insurance policy and the Platform made such a policy available in the future. The high-stakes nature of these choices was not just theoretical—the decisions that hiring managers made during the experiment actually impacted whether WCs could accept their posted jobs after the experiment.

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, we find that a sizable share of businesses in our sample, 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 customer safety and theft concerns. 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 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 are required to successfully complete one prior job, or when the pool of WCs is limited to those who have not been convicted or arrested in the past year. Wage subsidies can achieve similar increases but only at relatively high subsidy levels that may be cost prohibitive.² 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

¹Workers 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 do not legally apply to the Platform or the workers who use it.

²We 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 VII that these estimates imply that all of the non-subsidy policies we consider can increase the demand for WCs at $\frac{1}{10}$ to $\frac{1}{2}$ the cost of wage subsidies under reasonable assumptions.

the effects of an 80% wage subsidy according to a linear extrapolation of our experimental subsidy estimates. Requiring that WCs successfully complete 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 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 and non-WCs. We exploit the fact that some WCs inadvertently completed their first job on the Platform while their background check is pending, allowing us to objectively compare the performance ratings of WCs and non-WCs in their first job. We use an incentive-compatible elicitation of performance beliefs to show that hiring managers underestimate the performance of WCs in terms of both high- and low-performance ratings. We then show that providing objective information on the true share of high-performance ratings leads to more accurate beliefs and increases the level of demand for WCs by about 6 percentage points on average, approximately equivalent to the effect of a 40% wage subsidy. Providing objective information on the true share of low-performance ratings, typically from no-shows, also leads to more accurate beliefs but does not have a statistically significant impact on the level of demand. The muted effect of the low-performance information compared to the high-performance information suggests that businesses may be less concerned with no-shows and more concerned with general performance when deciding whether or not to work with WCs.

Based on our findings, the Platform carried out a staged roll-out to relax the criminal background check procedures it coordinates for businesses. First, the Platform extended job offers to WCs on behalf of the businesses that indicated a willingness to work with WCs under then-current Platform conditions. Second, the Platform allowed thousands of businesses posting new jobs to include WCs in their pool of potential workers with crime and safety insurance covering damages up to \$1 million, one of the most promising randomized conditions tested in this study. Eventually, the Platform plans to change the default option so that WCs are included in the pool of potential workers for all businesses unless they pay an additional fee to explicitly exclude WCs. Through August 2021, demand from our study participants and the staged roll-out led to over 12,000 jobs being made

available to WCs. 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 effect of these new job opportunities on the outcomes of WCs.

Beyond demonstrating that a range of policies can increase the demand for WCs, our paper provides new evidence on why employers 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 positive effect of crime and safety insurance, suggest that some employers are concerned by the potential for downside risk. Other results, including the large positive effect of objective performance information, suggest that other employers are concerned by the potential for lower productivity. The positive effects of the wage subsidies, performance screening, and screening of the most recent records are consistent with both risk and productivity concerns.³

Our paper also builds on important work by Holzer (2007), Holzer, Raphael, and Stoll (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 businesses report being more willing to hire WCs if there are wage subsidies, certificates of validated work performance history, or guaranteed replacement workers. These businesses 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 low-stakes nature of these surveys makes it difficult to know whether businesses are expressing their true preferences or just their aspirations. We add to this literature by measuring the demand for WCs using the 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 documents the baseline level of demand for WCs in our sample. Section IV presents our main experimental results. Section V presents results from providing objective performance information. Section VI discusses alternative explanations for our results, and Section VII concludes. The Online Appendix provides additional details of the experimental design and results.

³These results relate to work in personnel economics on hiring (Oyer and Schaefer 2011). 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 employers from the perceived downside risk of hiring disadvantaged workers for short-term positions.

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, Raphael, and Stoll 2004; Raphael 2011; Yang 2017; Schnepel 2018). The Platform is hosted on both the web and a mobile app, but the work they support generally does not involve computers or the internet, nor 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 additionally 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. The job is then reserved for the worker who accepted it and no other workers can 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 an accepted job request within 12 hours of the job’s start time, the business generally must pay a cancellation fee based on 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 who could accept their jobs, 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.

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, hiring managers are asked to rate each worker’s overall performance on a scale of one to five. Hiring managers are also asked to rate workers on specific attributes such as timeliness, cooperation, and quality of work. In practice, however, hiring managers only complete these more specific reviews 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 potential Platform applicants are currently screened out by a criminal background check. The Platform incurs a cost to run each background check but, beyond this check, does not extensively screen 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 businesses’ 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 complied 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 and the commitment of the Platform’s executives and board members, allowed the Platform to truthfully inform hiring managers that their responses during the experiment may determine whether WCs could accept their jobs in the future.

The *ex-ante* incentive-compatible nature of the hiring managers’ choices was reinforced by two *ex-post* actions taken by the Platform after the experiment. First, the Platform immediately allowed WCs to accept jobs posted by the businesses that were willing to work with WCs under then-current Platform conditions when there was a pool of available WCs. Second, the Platform used all of the businesses’ choices to set a new Platform-wide policy, where businesses posting new jobs could include WCs in their pool of potential workers with crime and safety insurance covering damages up to \$1 million. The Platform is also planning to change the default option so that WCs are included in the pool of potential workers for all businesses, 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.

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.

Methodologically, we build on work by Mas and Pallais (2017), Low (2017), and Kessler, Low, and Sullivan (2019) that generates incentive-compatible responses in field experiments by asking respondents to make hypothetical decisions that can potentially impact later outcomes. Our approach also builds on the strategy method in lab experiments (Brandts and Charness 2011), where participants make multiple conditional decisions and one decision is randomly implemented. Like this previous work, we rely on the fact that the exact probability that a choice is implemented is generally unimportant in generating incentive-compatible responses (Carson and Groves 2007; Charness, Gneezy, and Halladay 2016).

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 who had at least three months of experience on the site 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 hypothetical hiring 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 signature (“Sincerely, [Platform] 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.

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 questions. Eighty-six percent of hiring managers completed the hiring questions conditional on opening the email, with 91% of managers completing the questions conditional on reaching the first question related to WCs. 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) and note that 80% 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. Our results are qualitatively unchanged if we limit the sample to businesses with only one respondent or to hiring managers with the unilateral authority to allow WCs to perform the jobs they post.

The Platform emailed 7,450 hiring managers during the outreach, meaning that about 14% of hiring managers opened the email and completed all of the hiring questions. According to the Platform, this is a typical or slightly higher response rate for when it contacts the full pool of hiring managers to make decisions or request input, reflecting that many hiring managers are not actively

looking for workers at a particular time. Appendix Table B.2 shows that businesses responding to the survey and entering our experimental sample are similar in terms of industrial composition to businesses that did not respond to the survey, but skew younger (19 years vs. 22 years) and larger (40 employees vs. 21 employees). Businesses responding to the survey are also more active on the Platform compared to businesses that did not respond to the survey (2,828 posted jobs vs. 522 posted jobs over the past two years), consistent with the higher-stakes nature of the survey for more active users. We are unable to compare businesses based on WC hiring policies, as we gather this information during the survey experiment.

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:

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 the no subsidy condition described above. 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 the displayed values.

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 wage subsidy with the other randomized treatments. 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, and the \$5 million insurance level provides liability against much more serious crimes like sexual assault and murder. 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.

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

questions. Here, we exploit the fact that the background check is not initiated until a worker is matched to their first job on the Platform. Since jobs often begin shortly after the match occurs and workers are allowed to remain active while the background check is pending, we observe that about 5% of WCs complete the first job before their background check is returned and they are removed from the Platform. The fact that some WCs complete their first job on the Platform allows us to compare the first-job performance ratings of otherwise qualified WCs and non-WCs. These performance data reveal that WCs modestly out-perform non-WCs in their first job 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, Persico, and Weiss 2018). We caution, however, that the performance of WCs and non-WCs may differ over longer time horizons and in other settings.

We first measure baseline performance beliefs using an incentive-compatible guessing game about the performance of WCs on the Platform that rewards accuracy. We provide information on the performance of non-WCs and ask respondents to report the relative performance difference between WCs and non-WCs. We focus on relative performance because participants may be unfamiliar with the mapping between the performance metric in question (5-star ratings) and underlying productivity. Rewards in the guessing game ranged between \$2 and \$10 for an answer within 5% of the truth, where 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. We collected posterior beliefs from everyone because we expected that some treated respondents would not pay attention to the information while other treated respondents may not see the information as credible or complete. We also expected that the control group, who saw no objective performance information, might update their beliefs with the passing of time and new knowledge that the Platform elicited beliefs and shared the truth with some. 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 explore heterogeneity across labor market conditions by giving hiring managers the option of selecting “Only if it’s hard to fill my jobs” for all

of the questions in the experiment, providing a targeted measure of labor market tightness that is specific to each business’s context. We also 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 explore heterogeneity across job characteristics by asking hiring managers about the typical jobs they post, including whether there are any customer interactions or access to high-value inventory. These questions allow us to explore the extent to which businesses are concerned about violence or theft when considering downside risk. We also asked hiring managers whether their business has a hiring policy related to WCs to understand the potential constraints to working with 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 downside risk or lower productivity. The framework also 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 are 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, θ , 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. 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, θ , 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 less likely to have a business-wide WC policy (45% vs. 66%) and to report wanting 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 more closely match the distribution of firms in the U.S. economy based on these characteristics.

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.⁴

III.A. Baseline Results

Figure I reports the estimated share of businesses that are willing to work with WCs at each randomized wage subsidy level. 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

⁴Our empirical analysis closely follows our pre-analysis plan, available at AEARCTR-0005200.

business, the level of random assignment.

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.⁵ Estimates 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, as seen in Appendix Figure B.1.

Appendix Table B.3 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 businesses 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.

⁵The point estimates shown in Figure I are non-monotonic over some ranges, with a slightly lower share 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 $p = .60$. The simplest explanation for these results is that our non-parametric estimates include sampling error due to having only $\frac{1}{5}$ of the sample at each subsidy level.

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 reports the estimated share of businesses that are always willing to work with WCs, where we code businesses as willing to work with WCs only if they responded “Yes.” The second bar instead reports the estimated share 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 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 25 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.2. 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:

$$(1) \quad \text{Hire}_i = \alpha_0 \text{No HVI}_i + \alpha_1 \text{HVI}_i + \sum_{k \in K} \lambda_k \cdot \text{Subsidy}_{ik} + e_i$$

where No HVI_i is a dummy for not having high-value inventory and HVI_i is a dummy for having high-value inventory (and likewise for having or not having customer interactions). 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 downside risk and productivity and concerns.

IV.A. *Crime and Safety Insurance*

Figure III reports the estimated share 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. 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 :

$$(2) \quad \text{Hire}_{ij} = \sum_{l \in L} \lambda_l \cdot \text{PolicyLevel}_{il} + \sum_{k \in K} \lambda_k \cdot \text{Subsidy}_{ik} + e_{ij}$$

where the first observation codes willingness to work with WCs in the baseline case (Hire_{i0}) and the second observation codes willingness to work with WCs under the counterfactual policy (Hire_{i1}). 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. 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, with the coefficient on the no policy indicator measuring

the share of businesses that are willing to work with WCs at baseline.⁶

One potential concern with estimating the stacked regression specification in Equation (2) 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, Romero, and Wüthrich 2020). 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 with 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. The interaction effects are statistically insignificant for the crime and safety insurance and job history treatments ($p = .78$ and $p = .52$, respectively), but statistically significant for the limited screening treatment ($p = .02$). 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. For completeness, we show the entire labor demand curve for each of our main treatments in Appendix Figure B.3.

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% ($p = .02$ and $p < .01$ for Panel A and B, respectively). 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.4, 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

⁶Including flexible controls for the wage subsidy amount can lead to bias if the non-subsidy treatments are only conditionally random and any treatment effect heterogeneity is correlated with the wage subsidy controls (Goldsmith-Pinkham, Hull, and Kolesár 2022). But the subsidy and non-subsidy treatments are unconditionally random in our experiment and, as a result, we find little scope for such contamination bias in our full-sample regression estimates (see Appendix Table B.4). There is also no scope for such contamination bias when we focus on the subset of hiring managers who were randomized to have no wage subsidy, as these estimates do not include the wage subsidy controls.

percentage points from the baseline level of demand. However, the effects of the insurance coverage at \$100,000 and \$5m are statistically indistinguishable from the effect of insurance coverage at just \$5,000 in both the no subsidy sample ($p = .45$ and $p = .45$, respectively) and in the full sample ($p = .19$ and $p = .14$, respectively).

The results from Figure III and Appendix Figure B.4 suggest that businesses are particularly concerned about the types of moderate risks covered by lower levels of insurance (e.g., due to petty theft) and, to a lesser 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 insurance cap are particularly striking, as the \$5,000 cap is equal to that of the rarely used U.S. Federal Bonding Program. For example, the Federal Bonding Bond Program’s website reports that there were only 1,691 insurance policies issued to 1,068 individuals through the program last year (Federal Bonding Program 2022), while individuals managing the program report that the number of insurance policies issued was less than 1,000 per year in the early 2000s (Holzer, Raphael, and Stoll 2003). Our estimates raise the possibility that the Bonding Program’s low usage reflects non-demand-based reasons, such as employers being unaware of the program or having difficulty navigating the program requirements.

We can also interpret the magnitude of our results for the \$5,000 insurance cap using our simple conceptual framework. We observe that the median job on the Platform pays \$15 per hour, and our estimates show that firms act as though insurance that covers damages up to \$5,000 is roughly equivalent to an 80% wage subsidy. In our framework, a risk-neutral firm will equally value a \$5,000 insurance policy and an 80% wage subsidy when there is a 1.9% probability of a negative event covered by the insurance in a workday. In practice, however, hiring managers may be risk-averse or have inaccurate 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 test these different explanations with our data.

IV.B. Screening Based on Performance History

Figure III next reports the estimated share 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. 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 share 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.5, 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. However, the effects of these more restrictive five and 25 job history requirements are statistically indistinguishable from requiring just one successfully completed job in both the no subsidy sample ($p = .17$ and $p = .33$, respectively) and in the full sample ($p = .10$ and $p = .86$, respectively).

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.

IV.C. Screening Based on Criminal Record History

Figure III finally reports the estimated share 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 share 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.6, particularly at the 7-year level where the share of businesses willing to work with WCs increases by 41 percentage points from the baseline level of demand in the no subsidy sample ($p < .01$) and 28 percentage points from the baseline level of demand in the full sample ($p < .01$).

We further explore the importance of screening on criminal record history in Figure IV, which reports the estimated share 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 a misdemeanor conviction compared to a felony conviction. In the full sample, willingness to hire is much higher for property misdemeanors compared to property felonies (29% to 11%, $p < .01$), much higher for drug misdemeanors than drug felonies (50% to 28%, $p < .01$), and much higher for violent misdemeanors than violent felonies (9% to 5%, $p = .01$).

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, Raphael, and Stoll (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 to correct potentially 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 for the full sample of hiring managers following the provision of the objective information, with both sets of beliefs elicited using the incentive-compatible guessing game. The solid lines show the distribution of posterior beliefs for the respondents who were shown objective information about WC performance. The dashed lines show the distribution of prior beliefs for 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% likelihood, and underestimate the likelihood of a five-star rating by 12 percentage points, off an 87% likelihood. These mistaken and pessimistic beliefs are consistent with businesses inaccurately 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 (2022) 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:

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

where the information shock, $(p_{\text{signal},i} - p_{\text{prior},i})$, interacted with the treatment indicator, Info_i , is an instrument for hiring managers' posterior beliefs. $p_{\text{signal},i}$ equals the information that the hiring manager is actually assigned in the treatment group and what would have been assigned in the control group. $H_{\text{prior},i}, H_{\text{posterior},i} \in \{0, 1\}$ 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 (2022), and Fuster et al. (2020).⁷

⁷Appendix Figure B.7 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 underlying our empirical design and build intuition for our approach. 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 initial errors.

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. We focus on the larger sample of all hiring managers throughout to maximize statistical precision.⁸

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 elasticity of hiring for 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 a 40% wage subsidy.

By comparison, the IV results in column 3 of Panel B of Table IV show that changing perceptions in the low-performance group has a statistically insignificant impact on hiring decisions ($p = .36$). 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.

Control hiring managers also partially eliminated their initial errors. 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.

⁸Estimates are qualitatively similar but much less precisely estimated in the smaller no subsidy sample. 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 an IV coefficient of 0.067 (s.e. = 0.074) in the full sample and 0.161 (s.e. = 0.152) in the no subsidy sample.

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 address this concern. Our study is based on businesses making what are perceived as 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.

VI.B. *Screening Expectations*

A second consideration is that some hiring managers may have been confused, not paying attention, or made incorrect assumptions about whether the Platform would pre-screen the WCs in some way. We can explore the importance of these issues by examining the consistency of the answers at the start and end of the experiment for the subset of hiring managers not receiving the performance information. We asked the direct question about whether businesses would allow WCs to accept their jobs twice, once at the start of the experiment and once at the end of the experiment. 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, making it likely that by the end of the experiment most hiring managers better understood the questions and that the Platform would not pre-screen the WCs. We find that 85% of the hiring managers that did not receive the objective performance information answered in the same way at the start and end of the experiment, suggesting that confusion, inattention, or mistaken beliefs about screening cannot explain our results.

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. Appendix Tables B.5 and B.6 explore this concern by reporting adjusted p -values that control for the probability of making one or more false discoveries when performing multiple hypothesis tests. We employ a step-down algorithm that uses permutation calculations to estimate dependence relationships and provide corrected p -values for a given family of k -hypothesis tests, similar to the step-down algorithm described by Westfall and Young (1993). For the adjusted p -values in Appendix Table B.5, the family of nine hypotheses includes three for the mean effect of crime and safety insurance relative to the baseline, three for the mean effect of job history relative to the baseline, and three for the mean effect of screening by years since arrest or conviction relative to the baseline. For the adjusted p -values in Appendix Table B.6, the family of five hypotheses includes the mean effect of each crime type restriction relative to the violent felony conviction restriction. Our main findings are qualitatively unchanged in the full sample when adjusting for multiple hypothesis testing. The adjusted p -values for the \$5,000 crime and safety insurance policy and one job history are both .055, while the adjusted p -values for the other treatments remain below .05. In the smaller no subsidy sample, the adjusted p -values for the crime and safety insurance effects, the job history effects, two of the three limited record screening effects, and the violent misdemeanor restriction are all greater than .05 due to the small number of observations.

VI.D. COVID-19 Pandemic

A fourth consideration is that the experiment occurred in March and April 2020 at the start of the COVID-19 pandemic. 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. Appendix Figure B.8 explores how the timing of the experiment may have impacted our findings by showing the baseline level of demand for hiring managers responding before and after the national declaration of a state of emergency and located in counties that were less and more exposed to the first wave of COVID-19 infections. We observe nearly identical levels of demand for WCs across these groups, with none of the estimates suggesting that our findings are driven by the timing of the experiment

and the start of the pandemic.

VI.E. External Validity

A final and more general consideration is the external validity of our estimates to other settings (List 2020). 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 the jobs offered through the Platform often lead to additional 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 and thus there is no way to determine if any of the jobs convert to full-time work. 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 customer interactions and high-value inventory impact the demand for WCs suggests that role-specific traits are meaningfully correlated with demand and must be taken into consideration when extrapolating to other settings.

Appendix Tables B.7 and B.8 partially explore these issues by showing treatment effects relative to the baseline when we restrict our sample to firms with information on firm size, industry type, and WC hiring policies; when we weight observations to more closely match the distribution of answers to nationwide responses to SHRM questions about attitudes and policies in place concerning WCs; and when we 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 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 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 a job applicant’s arrest and conviction records 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 show 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, crime and safety insurance can increase the demand for WCs at $\frac{1}{10}$ to $\frac{1}{2}$ the cost of wage subsidies under plausible assumptions of the probability of damages due to WC misbehavior.⁹ These calculations

⁹Providing a \$5,000 crime and safety insurance policy increases the share of businesses willing to work with WCs by approximately 10%, about the same effect as a 50% wage subsidy. The cost of providing a 50% wage subsidy is about \$60 per worker per day, assuming a typical Platform wage of \$15 per hour and an 8-hour workday. By comparison, the expected cost of a \$5,000 crime and safety insurance policy is \$5 to \$25 per worker per day, assuming WCs have a one in 1,000 to a one in 200 daily chance of incurring \$5,000 in damages.

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, the supply of WCs for each type of job, and any heterogeneity in the demand and supply for different types of WCs. We view the careful modeling and estimation of such labor supply and demand models as 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 to 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, the effect of the short-term employment opportunities created for WCs that accept jobs on the Platform, and the consequences of the Platform eliminating the criminal background check entirely. Finally, 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, Amanda, and Sonja Starr, “The Effect of Criminal Records on Access to Employment,” *American Economic Review: Papers & Proceedings*, 107 (2017), 560–564.

- , “Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment,” *Quarterly Journal of Economics*, 133 (2018), 191–235.
- Armantier, Olivier, Scott Nelson, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar, “The Price Is Right: Updating Inflation Expectations in a Randomized Price Information Experiment,” *Review of Economics and Statistics*, 98 (2016), 503–523.
- Bergmann, Michael, “IPFWEIGHT: Stata Module to Create Adjustment Weights for Surveys,” Statistical Software Components, Boston College Department of Economics, 2011.
- Bollinger, Christopher R., and Julie L. Hotchkiss, “The Upside Potential of Hiring Risky Workers: Evidence from the Baseball Industry,” *Journal of Labor Economics*, 21 (2003), 923–944.
- Brandts, Jordi, and Gary Charness, “The Strategy Versus the Direct-Response Method: A First Survey of Experimental Comparisons,” *Experimental Economics*, 14 (2011), 375–398.
- Bushway, Shawn D., “Labor Market Effects of Permitting Employer Access to Criminal History Records,” *Journal of Contemporary Criminal Justice*, 20 (2004), 276–291.
- Carson, Richard T., and Theodore Groves, “Incentive and Informational Properties of Preference Questions,” *Environmental and Resource Economics*, 37 (2007), 181–210.
- Charness, Gary, Uri Gneezy, and Brianna Halladay, “Experimental Methods: Pay One or Pay All,” *Journal of Economic Behavior & Organization*, 131 (2016), 141–150.
- Couloute, Lucius, and Daniel Kopf, “Out of Prison & Out of Work,” Prison Policy Initiative Report, 2018.
- Craigie, Terry-Ann L., “Ban the Box, Convictions, and Public Employment,” *Economic Inquiry*, 58 (2020), 425–445.
- Cullen, Zoë, and Ricardo Perez-Truglia, “How Much Does Your Boss Make? The Effects of Salary Comparisons,” *Journal of Political Economy*, 130 (2022), 766–822.
- Doleac, Jennifer L., and Benjamin Hansen, “The Unintended Consequences of “Ban the Box”: Statistical Discrimination and Employment Outcomes When Criminal Histories Are Hidden,” *Journal of Labor Economics*, 38 (2020), 321–374.
- Federal Bonding Program, “About the FBP,” available at <https://bonds4jobs.com/about-us>, accessed on July 13, 2022.
- Fuster, Andreas, Ricardo Perez-Truglia, Mirko Wiederholt, and Basit Zafar, “Expectations with

- Endogenous Information Acquisition: An Experimental Investigation,” *Review of Economics and Statistics*, forthcoming, 2020.
- Goldsmith-Pinkham, Paul, Peter Hull, and Michal Kolesár, “Contamination Bias in Linear Regressions,” NBER Working Paper No. 30108, 2022.
- Holzer, Harry J., “Collateral Costs: The Effects of Incarceration on the Employment and Earnings of Young Workers,” IZA Discussion Paper No. 3118, 2007.
- Holzer, Harry J., Steven Raphael, and Michael A. Stoll, “Employment Barriers Facing Ex-Offenders,” Urban Institute Reentry Roundtable, 2003.
- , “How Willing are Employers to Hire Ex-Offenders?,” *Focus*, 23 (2004), 40–43.
- , “Perceived Criminality, Criminal Background Checks, and the Racial Hiring Practices of Employers,” *The Journal of Law and Economics*, 49 (2006), 451–480.
- , “The Effect of an Applicant’s Criminal History on Employer Hiring Decisions and Screening Practices: Evidence from Los Angeles,” In Bushway, Shawn D., Michael A. Stoll, and David F. Weiman, eds., *Barriers to Reentry? The Labor Market for Released Prisoners in Post-Industrial America*, (New York: Russell Sage Foundation, 2007).
- Hunt, Priscillia, Rosanna Smart, Lisa Jonsson, and Flavia Tsang, “Breaking Down Barriers Experiments into Policies That Might Incentivize Employers to Hire Ex-Offenders,” RAND Research Report No. 2142, 2018.
- Infogroup, “Infogroup U.S. Historical Business Data,” 2016.
- Jackson, Osborne, and Bo Zhao, “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, 2017.
- Kessler, Judd B., Corinne Low, and Colin Sullivan, “Incentivized Resume Rating: Eliciting Employer Preferences without Deception,” *American Economic Review*, 109 (2019), 3713–3744.
- Lazear, Edward P., “Hiring Risky Workers,” In *Internal Labour Markets, Incentives and Employment*, (London: Palgrave Macmillan, 1998), 143–158.
- List, John, “Non est Disputandum de Generalizability? A Glimpse into The External Validity Trial,” *Artefactual Field Experiments* 00711 (2020).
- Low, Corinne, “A ‘Reproductive Capital’ Model of Marriage Market Matching,” Manuscript, 2017.

- Mas, Alexandre, and Amanda Pallais, “Valuing Alternative Work Arrangements,” *American Economic Review*, 107 (2017), 3722–3759.
- Minor, Dylan, Nicola Persico, and Deborah M. Weiss, “Criminal Background and Job Performance,” *IZA Journal of Labor Policy*, 7 (2018), 1–49.
- Muralidharan, Karthik, Mauricio Romero, and Kaspar Wüthrich, “Factorial Designs, Model Selection, and (Incorrect) Inference in Randomized Experiments,” NBER Working Paper No. 26562, 2020.
- Oyer, Paul, and Scott Schaefer, “Personnel Economics: Hiring and Incentives,” In Ashenfelter, Orley, and David Card, eds., *Handbook of Labor Economics* (1 ed.) vol. 4B, (Amsterdam: Elsevier, 2011), 1769–1823.
- Pager, Devah, “The Mark of a Criminal Record,” *American Journal of Sociology*, 108 (2003), 937–975.
- Pallais, Amanda, “Inefficient Hiring in Entry-Level Labor Markets,” *American Economic Review*, 104 (2014), 3565–3599.
- Raphael, Steven, “Improving Employment Prospects for Former Prison Inmates: Challenges and Policy,” In *Controlling Crime: Strategies and Tradeoffs*, (Chicago: University of Chicago Press, 2011), 521–565.
- Rose, Evan, “Does Banning the Box Help Ex-Offenders Get Jobs? Evaluating the Effects of a Prominent Example,” *Journal of Labor Economics*, 39 (2021), 79–113.
- Schnepel, Kevin T., “Good Jobs and Recidivism,” *Economic Journal*, 128 (2018), 447–469.
- Shoag, Daniel, and Stan Veuger, “No Woman No Crime: Ban the Box, Employment, and Upskilling,” HKS Working Paper No. 16-015, 2016.
- Society for Human Resource Management, “Workers with Criminal Records: A Survey by the Society for Human Resource Management and the Charles Koch Institute,” Report, 2018.
- Stoll, Michael A., “Ex-Offenders, Criminal Background Checks, and Racial Consequences in the Labor Market,” *University of Chicago Legal Forum*, 2009 (2009), 381–419.
- Uggen, Christopher, Mike Vuolo, Sarah Lageson, Ebony Ruhland, and Hilary K. Whitham, “The Edge of Stigma: An Experimental Audit of the Effects of Low-Level Criminal Records on Employment,” *Criminology*, 52 (2014), 627–654.

- Westfall, Peter H., and S. Stanley Young, *Resampling-Based Multiple Testing: Examples and Methods for p-Value Adjustment*, vol. 279, (New York: John Wiley & Sons, 1993).
- Yang, Crystal S., “Local Labor Markets and Criminal Recidivism,” *Journal of Public Economics*, 147 (2017), 16–29.

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	.291	.646	.407	.268	.271	.347
Employees	.257	.858	.099	.009	.613	.424
Service	.240	.956	.576	.287	.286	.711
Manufacturing	.045	.277	.949	.877	.269	.386
Retail	.393	.138	.036	.684	.873	.231
Transportation & Public Utilities	.691	.908	.625	.434	.765	.988
Nonclassifiable & Misc. Industries	.494	.454	.697	.552	.937	.261
Firm-Wide WC Hiring Policy	.779	.673	.728	.513	.491	.229
Platform Tenure (Years)	.099	.075	.602	.518	.224	.947
Job Vacancy Rate	.857	.569	.151	.928	.733	.913
B. Firm Characteristics and Policies						
Job Involves Customer Interactions	.675	.595	.429	.356	.628	.710
Job Involves High-Value Inventory	.277	.450	.242	.531	.467	.437
Modal Job is Fulfillment / Warehousing	.767	.325	.559	.474	.284	.052
Modal Job is General Labor	.523	.085	.499	.258	.231	.117
Modal Job is Event Staff	.341	.950	.762	.457	.718	.559
Modal Job is Delivery	.965	.097	.324	.964	.726	.808
Modal Job is Washing & Cleaning	.940	.732	.738	.051	.726	.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
Effects of High- and Low-Performance Information

	First Stage	OLS	IV	Reduced Form
A. Impact of High-Performance Information				
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. Impact of Low-Performance Information				
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. Standard errors clustered by business are in parentheses. Panel A reports results for managers who were shown information on the fraction of 5-star ratings. Panel B reports results for managers who were shown information on the fraction of no-shows and either 1- or 2-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.


FIGURE I
Labor Demand for Workers with a Criminal Record




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. The figure reports the share 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 where we regress an indicator for whether a business is willing to work with WCs on dummies for the different subsidy levels. 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). 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 first bar reports the estimate from a stacked regression specification similar to Equation (2) for the share of businesses willing to work with WCs if we code businesses as willing to work with WCs only if they responded “Yes” when asked if they are willing to work with a WC and the second bar reports the estimate from the stacked regression for the share of respondents who answer “Yes” or “Only if it’s hard to fill my jobs” to the question. All other estimates report the estimates of the share of respondents who answer “Yes” to the same question from regression specifications described by Equation (1). 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 1,095 hiring managers from 913 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



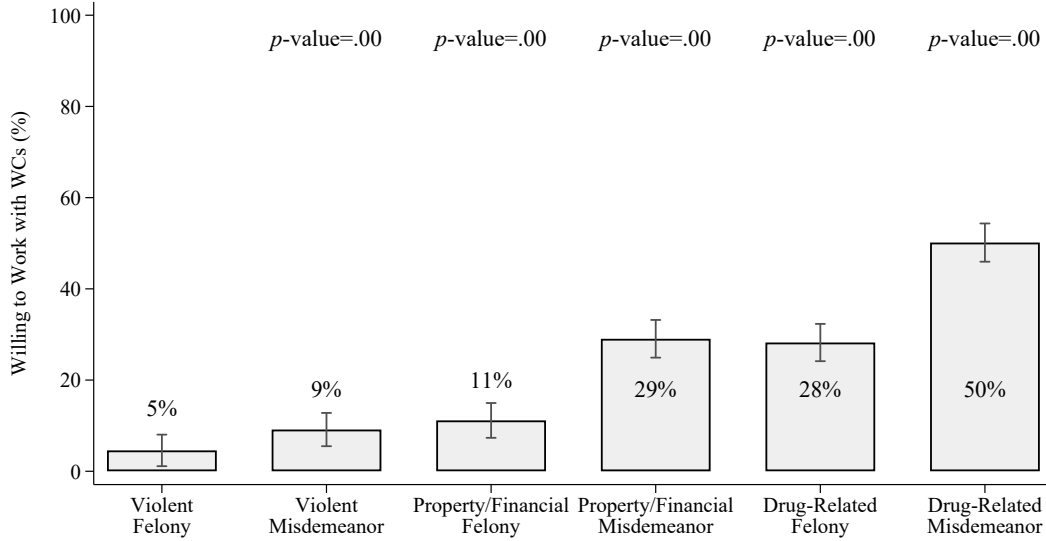
Notes. This figure presents estimates of the mean willingness to work with WCs (and 95% confidence interval) if the Platform provides crime and safety insurance, job history screening, or limited criminal record screening. The estimates are based on results from regression specifications described by Equation (2). 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 share 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



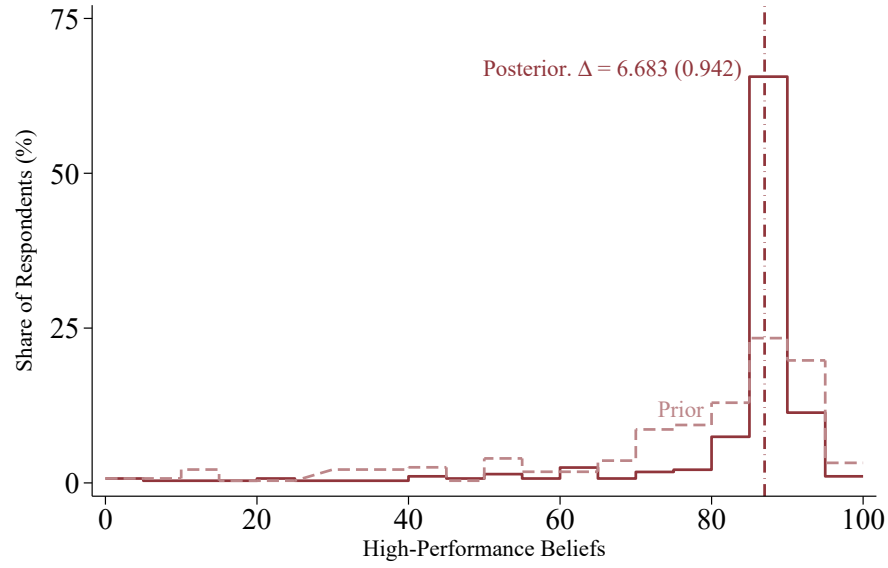
B. Mean Effect in the Full Sample



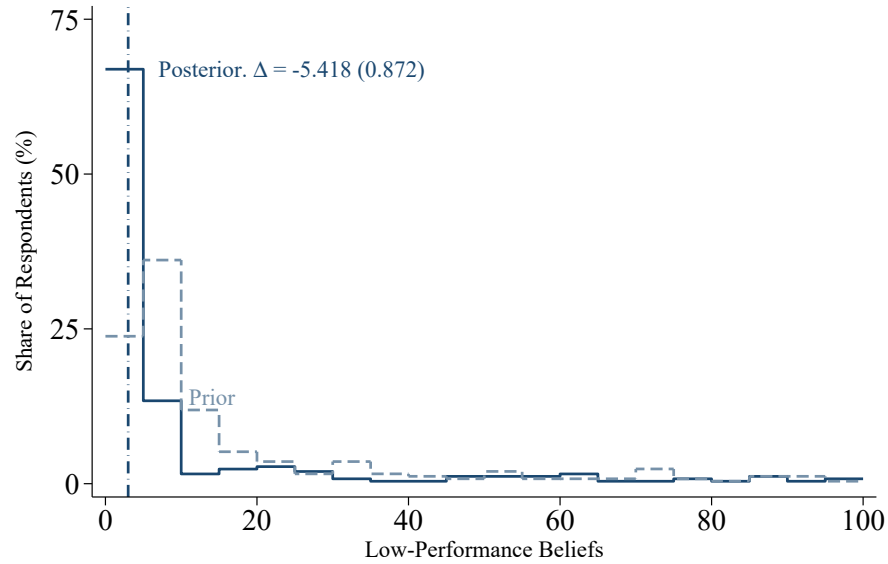
Notes. This figure presents estimates of the mean willingness to work with WCs (and 95% confidence interval) by conviction type. The estimates are based on results from regression specifications described by Equation (2). Each p-value corresponds to the test of the null hypothesis that there is no difference 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 High- and Low-Performance Information on Beliefs

A. High-Performance Beliefs



B. Low-Performance Beliefs



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 Δ is the impact of the treatment on beliefs, with the standard error of this estimate in parentheses.