Disclosure statement

Mar 26, 2026

We have no sources of research support, financial relationships, or potential conflicts of interest to report.

Piero Gottardi, Benjamin Lester, Ryan Michaels, and Ronald Wolthoff

Signals, Interviews, and Hiring in Labor Market Equilibrium*†

Piero Gottardi
University of Essex

Benjamin Lester
Federal Reserve Bank of Philadelphia

Ryan Michaels
Federal Reserve Bank of Philadelphia

Ronald Wolthoff
University of Toronto

March 27, 2026

Abstract

We present an equilibrium search model of the labor market with a multi-stage hiring process. Workers with different attributes apply for jobs, and firms choose the number and composition of workers to interview and which worker to hire. The model yields analytical solutions for callback, job-finding, and unemployment rates across workers with different observable attributes. We calibrate the model to study how race and criminal record affect workers' labor market experience. Then, motivated by the widespread adoption of Ban the Box legislation, we study the equilibrium implications of a policy that prohibits firms from observing a candidate's criminal record at the application stage. We find that the policy has economically significant effects on firms' job creation decisions *and* optimal interviewing strategies. We explore the dynamic implications of these effects across workers, and for the labor market as a whole.

Keywords: Labor Markets, Job Search, Interviews, Crime, Ban the Box
JEL Classification: J15, J63, J64, J71, J78

*We are particularly grateful to Teddy Kim for helpful comments.

†**Disclaimer**: The views expressed in this paper are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

1 Introduction

Hiring a worker is a complicated undertaking that involves multiple, distinct stages. A firm typically receives interest from many candidates for a single position. However, gathering information about applicants, including conducting interviews, is time consuming and costly. Therefore, it can be necessary to screen out many candidates using an initial (written) application; only a subset of applicants is interviewed.

This screening process has attracted increasing attention from policymakers and scholars. Concerned that some workers do not receive a "fair shake", lawmakers have enacted policies designed to restrict screening based on certain characteristics. Many jurisdictions now limit firms' access to applicants' criminal records, credit scores, and even some aspects of their employment history, such as past wages. A growing empirical literature studies the effects of such policies on several labor market outcomes, including callback (i.e., interview) rates.[1]

However, our understanding of the effects of these policies remains incomplete. A key challenge is to reconcile the estimated effects within a single micro-founded framework that captures the many margins along which firms can adjust. For example, employers may change the total number of interviews; the composition of the workers interviewed; and the number of vacancies for which they recruit. Can we account for the *joint* dynamics of interviewing and vacancy posting and their implications for the distributions of job finding rates and wages across different workers? What is the ultimate effect on the size and composition of unemployment? Finally, should we expect the effects to unfold quickly or gradually?

To answer these questions, we develop a tractable theory of aggregate labor market dynamics in which firms receive (potentially many) applications from ex ante heterogeneous workers and optimally choose which of these applicants to interview and, ultimately, hire. Our framework thus introduces a micro-founded intensive margin into the canonical search and matching market of equilibrium unemployment. Firms' screening decisions will affect—and be affected by—the size and composition of the pool of unemployed workers. The model can speak to *both* micro data on workers' and firms' search experiences and macro data on job creation, worker flows, employment, and wages. Finally, the framework can be easily adapted to the case where firms can only observe a subset of worker characteristics, which enables us to study the effects of policies that restrict information during certain stages of the hiring process—on impact *and* along the transition path to a new steady state.

To elaborate, we consider an environment with workers who differ along an arbitrary number of different dimensions or "attributes." When unemployed, these workers search for a job by sending an application to a randomly selected (homogeneous) firm. Each firm receives

[1] Among others, see Agan and Starr (2018) on criminal records, Cortés et al. (2022) on credit reports, and Davis et al. (2022) on salary histories.

potentially many applications and must choose which workers to interview, or "call back", based on their observable attributes, along with payoff-relevant information contained in their application. However, applications alone do not reveal all of the relevant information; the firm needs to interview a worker to learn the total idiosyncratic productivity of the match.[2]

We characterize the optimal interviewing and hiring strategy for a firm conditional on *any* realization of both the number and types of applicants. Despite the complexity introduced by rich worker heterogeneity and a multi-stage screening process, we show that the model remains highly tractable: it admits analytical expressions for core objects of interest, including the callback rates and job-finding rates of each type of worker and the expected profits of firms. Importantly, we establish that all of these properties carry over to an environment in which firms can only observe some of a worker's attributes. The result is a parsimonious, tractable, and flexible framework that generates endogenous relationships between the information available to firms at various stages of the hiring process, their optimal recruiting decisions, and both the size and composition of the pool of unemployed.

We exploit this framework to study a policy commonly referred to as Ban the Box. This legislation was designed to improve the labor market prospects of exoffenders by prohibiting firms from asking if a candidate has a criminal record at the initial application stage of the hiring process, while still allowing for a background check later if a firm conducts an interview. This policy has garnered significant interest from policymakers and spurred a growing empirical literature. However, to the best of our knowledge, it has not been analyzed within the context of a dynamic equilibrium model of the labor market.

Ban the Box (BTB) is an instructive application for our framework. The policy places the role of screening center stage. Accordingly, the extent to which the model is consistent with the estimated labor market impact of the policy represents an important test of the theory's treatment of the recruiting process and the interplay between screening and job creation. In addition, understanding the Ban's many, and at times seemingly conflicting, effects is a worthwhile exercise in its own right. Over 35 U.S. states and some 150 cities and counties have adopted BTB, but its impact remains unclear.[3] While studies have found clear evidence that BTB has significant effects on callback rates across race and criminal record, the ultimate effects on key labor market outcomes, such as job creation and employment rates across workers, are mixed.

Our model is well-suited to contribute to the ongoing discussion of BTB and potentially reconcile some of the reduced-form results. To start, a key feature of this policy is that it

[2]Our environment shares features with important contributions by Jarosch and Pilossoph (2019) and Meier and Obermeier (2025). One key distinction is that, in our setting, the optimal ordering of interviews cannot be done based on a single observable, such as unemployment duration, or even on expected productivity.

[3]See Avery and Lu (2021) for a survey of Ban the Box policies.

conceals criminal records *before* firms choose which applicants to interview, but allows firms to observe it *after* the interviewing process (before offers are extended). Hence, studying this policy requires an explicit model of multi-stage search. Second, the key objects we derive in our model have clear counterparts in the data. In contrast to more abstract models, our framework can be used to make direct contact with moments from worker surveys and audit studies, including the contact, callback, and job-finding rates of workers with different characteristics. Third, by accounting for the equilibrium interaction between firms' vacancy creation and screening behavior and the evolution of unemployment, our model can tie together moments from the micro data with aggregate labor market flows. Lastly, by analyzing this policy within the context of a dynamic equilibrium setting, we are also able to conduct counterfactual exercises. For example, we can compare the effects of BTB with *other* policy interventions, such as subsidies to firms who hire workers with a criminal record.

The rest of the paper is organized as follows. Section 2 introduces a static, one-period model in which an exogenously specified pool of heterogeneous unemployed workers randomly send applications to a fixed set of vacancies. We derive the firms' optimal interviewing and hiring behavior, along with three key objects that follow from this behavior: the callback and job-finding rates of workers with different types, and the firms' expected profits. Then, in Section 3, we embed this model into a dynamic equilibrium framework. In doing so, we endogenize both the vacancy creation decision and the evolution of the composition and stock of unemployment. Next, in Section 4, we generalize the model to allow for incomplete information. Specifically, we show how our results extend to the case when firms can only observe a subset of applicants' payoff-relevant attributes. Finally, in Section 5, we use this framework to quantitatively evaluate the effects of Ban the Box. We show that our model can simultaneously generate a large change in callback rates across workers with different races and criminal records, while movements in aggregate unemployment are more modest. Nevertheless, a key insight is that, over time, the composition of unemployed workers changes and feeds back onto vacancy posting and recruiting decisions. This interplay creates nontrivial and economically interesting dynamics in job creation and interviewing intensity. An important empirical implication is that reduced-form estimates of the BTB effects can be sensitive to timing: an evaluation of BTB six months after the policy is implemented can reveal vastly different results than an identical study conducted two years later.

2 The Multi-Stage Search Process

In this section, we present a one-period model that describes the multiple stages of the job search process, with particular attention to the decisions that firms make in screening their

pool of applicants. At this stage, we take as given both the measure and composition of the pool of unemployed workers, along with the measure of vacancies that firms have created. Then, in Section 3, we embed this multi-stage job search process into an otherwise standard dynamic model of the labor market, in which the pool of unemployed workers and the number of vacancies are endogenously determined.

2.1 Agents, Preferences, and Technology

Consider an environment with a measure v of risk-neutral homogeneous firms who demand one unit of labor and a measure u of risk-neutral unemployed workers who supply one unit of indivisible labor. Workers are heterogeneous along N dimensions or "attributes," and we denote a worker's type by her vector of attributes, $\mathbf{x} \equiv (x_1, ..., x_N) \in \mathcal{X}$. Let $L = |\mathcal{X}| > 1$ denote the number of distinct worker types, and let $u(\mathbf{x})$ denote the measure of unemployed workers of type $\mathbf{x} \in \mathcal{X}$, with $\sum_{\mathbf{x} \in \mathcal{X}} u(\mathbf{x}) = u$.

A worker's potential output at a firm is stochastic and depends on her type. In particular, the match-specific output or "productivity" of a worker of type $\mathbf{x}$ at a randomly selected firm is the sum of three terms:

$$z_{\mathbf{x}} + \eta + \varepsilon.$$

The first term, $z_{\mathbf{x}}$, is the portion of a worker's productivity that a firm can deduce from observing the worker's attributes $\mathbf{x}$ alone. This term captures the common component of productivity among all workers with a particular set of (potentially observable) attributes, such as age, gender, race, educational attainment, and criminal record. The second term, η, is an idiosyncratic component of productivity that the firm learns at the *application stage*. This term captures information that is not common across workers with the same attributes $\mathbf{x}$, but is observable from the worker's application, resumé, or reference letters. For example, η can be thought to capture previous work experience, fluency in a foreign language, or proficiency with a certain type of computer software. We assume that η is drawn from a discrete, M-dimensional distribution with support $\{\eta_1, ..., \eta_M\}$, and let $g(\eta|\mathbf{x})$ denote the probability of a type $\mathbf{x}$ worker drawing η, with associated cumulative density function (CDF) given by $G(\eta|\mathbf{x})$. The final term, ε, is an idiosyncratic component of productivity that the firm learns at the *interview stage*, such as a worker's demeanor or ability to answer questions. We assume that ε is a continuous random variable drawn from the CDF $F(\cdot|\mathbf{x})$, which is independent of η, has mean zero, and support $[\underline{\varepsilon}, \overline{\varepsilon}]$.

All workers receive utility b from unemployment, which can be interpreted as unemploy-

ment insurance, home production, utility from leisure, or some combination of the three.[4] Hence, we denote the match surplus between a firm and a worker of type $\mathbf{x}$ who draws productivity shocks (η, ε) by

$$y = z_{\mathbf{x}} + \eta + \varepsilon - b. \tag{1}$$

2.2 The Job Search Process

In this section, we describe the job search process—how workers apply to firms, and how firms screen applicants in an attempt to find the most productive worker.

2.2.1 Worker Behavior

The first step in the job search process is that each unemployed worker sends a single application to a randomly selected firm. As is common in the literature, we assume an *urn-ball* meeting technology, so that the total number of applicants at each firm, denoted by n, is drawn from a Poisson distribution with mean $\lambda \equiv u/v$, which represents aggregate market tightness. Given the properties of the Poisson distribution, the number of applicants with attributes $\mathbf{x}$ at each firm will also be drawn from a Poisson distribution with mean $\lambda(\mathbf{x}) \equiv u(\mathbf{x})/v$, where $\sum_{\mathbf{x}\in\mathcal{X}} \lambda(\mathbf{x}) = \lambda$. Following common practice, we will often refer to the market tightness of a particular type of worker, $\lambda(\mathbf{x})$, as the "queue length."

2.2.2 Firm Behavior

After workers send applications, each firm observes (i) the number n of applications it has received, (ii) the type $\mathbf{x}$ of each applicant, and (iii) the realization of the shock η of each applicant. In order to hire a worker, we assume that the firm first has to *interview* him or her. For each interview, the firm incurs a cost $c > 0$ and learns the final idiosyncratic component of an applicant's output, ε.

After completing the interview process, we assume that firms make a take-it-or-leave-it offer to a single applicant. Given this assumption, it is well known (see, e.g., Diamond, 1971) that, in equilibrium, it is optimal for all firms to offer a wage equal to the outside option of the workers, b. Hence, the hiring decision of the firm at this final stage is trivial: they will hire the worker that they interviewed with the largest realized match surplus y, so long as $y \geq 0$.[5]

[4] It is straightforward to allow workers' utility from unemployment to depend on their type, $\mathbf{x}$.

[5] Naturally, our results extend to an environment where firms pay a fixed wage $w \neq b$; for example, one could interpret our framework as a labor market in which the minimum wage is binding. In ongoing work,

We now turn to the question of which workers the firm should interview. To start, it is well known (see, e.g., Lippman and McCall, 1976) that the firm should interview workers *sequentially*, so as to preserve the option value of stopping its search if it discovers an applicant with sufficiently high productivity. The following lemma describes the precise order and stopping rules that a firm should follow. Since this result is a direct application of Weitzman (1979), we refer the reader to Weitzman for a proof of the optimality of what he termed "Pandora's rule".

Lemma 1. *Consider a firm that receives $n \geq 1$ applicants, where the type $\mathbf{x} \in \mathcal{X}$ and the realization of $\eta \in \{\eta_1, ..., \eta_M\}$ is observed for each applicant. For each $\mathbf{x}$ and η, define $y^\star(\mathbf{x}, \eta)$ as the solution to:*

$$c = \int_{y^\star}^{\infty} (y - y^\star) \, dH(y|\mathbf{x}, \eta), \tag{2}$$

where

$$H(y|\mathbf{x}, \eta) = F\left(y - z_{\mathbf{x}} - \eta + b|\mathbf{x}\right). \tag{3}$$

The optimal interviewing and hiring strategy is as follows.

1. ***Ranking applicants:*** *The firm should place the n applicants in descending order according to the values of $y^\star(\mathbf{x}, \eta)$. Let $y_1^\star \geq y_2^\star \geq ... \geq y_n^\star$ denote the ordered values of the set $\{y^\star(\mathbf{x}^i, \eta^i)\}_{i=1}^n$.*[6]

2. ***Stopping rule:*** *For each $i \in \{1, 2, ..., n\}$, the firm should interview applicant i and observe $y^i = z_{\mathbf{x}}^i + \eta^i - b + \varepsilon^i$ if $y_i^\star \geq \max\{y^0, ..., y^{i-1}\}$, where $y^0 \equiv 0$ denotes the option of not hiring. If $y_i^\star < \max\{y^0, ..., y^{i-1}\}$, the firm should stop the interviewing process.*

3. ***Hiring rule:*** *Among the $\tilde{n} \leq n$ applicants that were interviewed, the firm should hire applicant $j = \arg\max\{y^0, ..., y^{\tilde{n}}\}$, with the convention that $j = 0$ denotes a firm that does not hire any workers.*

Intuitively, the "reservation value" $y^\star(\mathbf{x}, \eta)$ is defined as the hypothetical productivity of a worker who has already been interviewed that would make the firm indifferent between conducting an additional interview of a candidate with observable characteristics $(\mathbf{x}, \eta)$ or not (and just hiring the hypothetical worker). The crucial property of Weitzman's solution is that the optimal search process is completely summarized by these reservation values,

we incorporate Nash bargaining into the model.

[6] When multiple applicants have the same $y^\star$, we assume that the firm randomizes to determine their order.

which are *independent* of both the number and the observable characteristics of all other applications that the firm receives. Hence, the optimal search strategy is relatively simple. Our contribution is to solve the technically challenging problem of deriving the probability of various labor market outcomes for each type of worker—and for the firm—when both the number and the characteristics of applications are stochastic.

2.3 Characterizing Key Outcomes

Given firms' optimal interviewing strategy, described above, we now derive analytical expressions for three key objects: the probability that a worker with a specific set of attributes receives an interview (the "callback rate"); the probability that the firm hires a worker with a specific set of attributes (the "hiring rate"); and the distribution of payoffs that the firm receives ("profits"). These objects are important inputs into our equilibrium characterization in the dynamic model, analyzed in Section 3.

2.3.1 Callback rate

Lemma 1 describes how workers are interviewed in order according to a sufficient statistic, $y^\star$, which depends on the worker's ex ante type $\mathbf{x} \in \mathcal{X}$ with $L = |\mathcal{X}|$ and the realization of $\eta \in \{\eta_1, \ldots, \eta_M\}$. Therefore, to calculate the ex ante probability that a worker of type $\mathbf{x}$ is interviewed, we first need to derive the distribution over realizations of $y^\star$ for each type of applicant. To do so, let $K = L \times M$ and let $\{Y_1^\star, \ldots, Y_K^\star\}$ denote the ordered set of all possible values of $y^\star(\mathbf{x}, \eta)$, with $Y_1^\star \geq \ldots \geq Y_K^\star$. For now, let us assume that each realization is associated with a single type of worker, i.e., that there exists a unique $\mathbf{x}$ such that $Y_k^\star = z_\mathbf{x} + \eta - b + \varepsilon_\mathbf{x}^\star$ and hence $Y_1^\star > \ldots > Y_K^\star$. We relax this assumption later in 3.

Letting $\mathbf{x}_k \in \mathcal{X}$ and $\eta_k \in \{\eta_1, \ldots, \eta_M\}$ denote the pair $(\mathbf{x}, \eta)$ associated with $Y_k^\star$, we can define by

$$\lambda_k = g(\eta_k|\mathbf{x}_k)\lambda(\mathbf{x}_k) \tag{4}$$

the expected number of applicants with $\mathbf{x} = \mathbf{x}_k$ and $\eta = \eta_k$. The following lemma reports the probability that such an applicant is interviewed.

Lemma 2. *The probability that an applicant of type $\mathbf{x}_k$ who draws η_k is interviewed is:*

$$\Gamma_k = e^{-\sum_{j=1}^{k-1} \lambda_j \left(1 - H_j\left(Y_k^\star\right)\right)} \times \frac{1 - e^{-\lambda_k \left(1 - H_k\left(Y_k^\star\right)\right)}}{\lambda_k \left(1 - H_k\left(Y_k^\star\right)\right)}, \tag{5}$$

where, in a slight abuse of notation,

$$H_k(y) \equiv H(y|\mathbf{x}_k, \eta_k). \tag{6}$$

Proof. See Appendix A.1. □

Intuitively, two events need to occur for a firm to interview a worker with $y^\star = Y_k^\star$. First, it must be that the firm did not receive any applications with $y^\star \in \{Y_1^\star, \ldots, Y_{k-1}^\star\}$ and a realization $y^i \geq Y_k^\star$, as otherwise the firm would have stopped interviewing before reaching the applicant with $y^\star = Y_k^\star$. The probability of this event occurring is represented by the first term in equation (5). To see why, note that, for any $j \in \{1, \ldots, K\}$, the probability that there are no applicants with $y^\star = Y_j^\star$ with realized match surplus $y^i \geq y$ is $e^{-\lambda_j(1-H_j(y))}$. Accordingly, the *joint* probability of there being no applicants with $y^\star \in \{Y_1^\star, \ldots, Y_{k-1}^\star\}$ and surplus $y^i \geq y$ is $e^{-\sum_{j=1}^{k-1}\lambda_j(1-H_j(y))}$.

However, even if the firm does not stop interviewing before it reaches applicants with $y^\star = Y_k^\star$, there may be *multiple* applicants of type $\mathbf{x}_k$ who draw shock η_k. The second term in equation (5) represents the probability that each applicant with $y^\star = Y_k^\star$ is interviewed, conditional on at least one such applicant being interviewed.

The following corollary uses the expression in (5) to formalize the callback rate for individuals with a particular vector of attributes.

Corollary 1. *The callback rate of workers of type* $\mathbf{x}$ *is given by*

$$\Gamma(\mathbf{x}) = \sum_{k=1}^{K} g(\eta_k|\mathbf{x})\mathbb{I}\{\mathbf{x}_k = \mathbf{x}\}\Gamma_k,$$

where $\mathbb{I}\{\cdot\}$ *denotes the indicator function.*

2.3.2 Hiring rate

We now derive the probability that a firm hires a worker of type $\mathbf{x}$. As in the analysis of callback rates, we first derive hiring probabilities conditional on the realization of $y^\star(\mathbf{x}, \eta) \in \{Y_1^\star, \ldots, Y_K^\star\}$. Then, below, we map these probabilities onto ex ante worker types, $\mathbf{x} \in \mathcal{X}$.

Lemma 3. *For any* $k \in \{1, ..., K\}$*, the probability that a firm hires a worker with* $y^\star = Y_k^\star$ *is given by*

$$\Phi_k = e^{-\sum_{j=1}^{k-1}\lambda_j\left(1-H_j\left(Y_k^\star\right)\right)}\left(1 - e^{-\lambda_k\left(1-H_k\left(Y_k^\star\right)\right)}\right) + \sum_{\ell=k}^{K}\int_{Y_{\ell+1}^\star}^{Y_\ell^\star} e^{-\sum_{j=1}^{\ell}\lambda_j(1-H_j(y'))}\lambda_k dH_k(y'). \tag{7}$$

Proof. See Appendix A.2. □

To understand the expression in (7), note that there are two possible sequences of events that are consistent with a firm hiring a worker of type $\mathbf{x}_k$ who draws realizastion η_k, such that $y^\star = Y_k^\star$. First, it's possible that the firm did not receive any applications from workers with $y^\star \in \left\{Y_1^\star, \ldots, Y_{k-1}^\star\right\}$ and a realization $y^i \geq Y_k^\star$, but that it *did* receive an application with $y^\star = Y_k^\star$ and a realization $y^i \geq Y_k^\star$. Using the logic from the callback rates above, one can see that the first term in equation (7) captures the probability of this first event.

The second possibility is that the firm interviews a worker with $y^\star = Y_k^\star$ and a realization $y^i < Y_k^\star$, continues to interview additional workers, but does not find any applicants with a payoff $y' > y^i$. To derive the probability of this event occurring, consider first the case of $k = 1$. With probability

$$e^{-\lambda_1\left(1-H_1\left(Y_1^\star\right)\right)} \int_{Y_3^\star}^{Y_2^\star} e^{-\lambda_2(1-H_2(y'))} d\left(\frac{e^{-\lambda_1(1-H_1(y'))}}{e^{-\lambda_1\left(1-H_1\left(Y_1^\star\right)\right)}}\right) = \int_{Y_3^\star}^{Y_2^\star} e^{-\sum_{j=1}^{2}\lambda_j(1-H_j(y'))}\lambda_1 dH_1(y'),$$

the firm interviews an applicant with $y^\star = Y_1^\star$ and $y^i \in (Y_3^\star, Y_2^\star)$ such that no other applicant with $Y_2^\star$ is interviewed and discovered to have a surplus greater than y^i. This same logic can be applied to account for the possibility that $y^\star = Y_1^\star$ and $y^i \in (Y_4^\star, Y_3^\star)$ or $y^i \in (Y_5^\star, Y_4^\star)$ and so on, leading to the second term in equation (7).

Keeping in mind that Φ_k characterizes the probability that a firm hires a worker with $y^\star = Y_k^\star$, the following corollary uses this expression to characterize the probability that a worker of type $\mathbf{x}$ is hired by the firm.

Corollary 2. *The probability that a worker of type* $\mathbf{x}$ *is hired is given by*

$$\Psi(\mathbf{x}) = \sum_{k=1}^{K} g(\eta_k|\mathbf{x})\mathbb{I}\{\mathbf{x}_k = \mathbf{x}\}\Psi_k,$$

where $\Psi_k = \Phi_k/\lambda_k$.

2.3.3 Firm Profits

Finally, we derive the distribution of match surplus across firms. As in the analysis above, we first derive this distribution conditional on the realization of $y^\star(\mathbf{x}, \eta) \in \{Y_1^\star, \ldots, Y_K^\star\}$.

Lemma 4. *Define* $Y_{K+1}^\star \equiv 0$ *and, for any* $y \geq 0$, *let* $k(y) = \min\left\{k \in \{0, ..., K-1\} | Y_{k+1}^\star \leq y\right\}$. *Then, for any* $k \in \{1, ..., K\}$ *and* $y \geq 0$, *the probability that a firm hires a worker with* $y^\star = Y_k^\star$

with realized surplus $y^i \geq y$ is given by

$$\overline{\Phi}_k(y) = \sum_{\ell=k}^{k(y)} \int_{\max\{y, Y^\star_{\ell+1}\}}^{Y^\star_\ell} e^{-\sum_{j=1}^{\ell} \lambda_j (1-H_j(y'))} \lambda_k dH_k(y') + e^{-\sum_{j=1}^{k-1} \lambda_j \left(1-H_j\left(Y_k^\star\right)\right)} \left(1 - e^{-\lambda_k \left(1-H_k\left(Y_k^\star\right)\right)}\right) \frac{1 - H_k\left(\max\left\{y, Y_k^\star\right\}\right)}{1 - H_k\left(Y_k^\star\right)}. \tag{8}$$

Proof. See Appendix A.3. □

It is now straightforward to compute the unconditional distribution of firm profits across all active matches. This enables us to calculate the expected payoff from a filled job.

Corollary 3. *The expected profits of a firm are*

$$\int_0^\infty y d\Phi(y) - c \sum_{\mathbf{x}} \lambda(\mathbf{x}) \Gamma(\mathbf{x}), \tag{9}$$

where

$$\Phi(y) = 1 - \sum_{k=1}^{K} \overline{\Phi}_k(y) \tag{10}$$

denotes the probability that a firm hires any worker with $y^i \leq y$.

Intuitively, the first term in (9) is the expected revenue of the firm, while the second term is the expected interviewing costs. To understand the latter, note that the total number of type $\mathbf{x}$ workers who are interviewed is $u(\mathbf{x})\Gamma(\mathbf{x})$, i.e., the product of the mass of type $\mathbf{x}$ workers and the interview probability, $\Gamma(\mathbf{x})$. Summing across worker types and dividing by v yields the expected number of interviews each firm conducts, and multiplying by c then yields the expected interview costs per firm.

3 Multi-Stage Search in Dynamic Equilibrium

We now extend the one-period model from Section 2 to a discrete-time, infinite horizon setting. In this setting, an interesting interaction arises between firms' hiring behavior and the pool of unemployed workers. As in the analysis in Section 2, the size and composition of the pool of unemployed affects firms' optimal screening process, i.e., the number of interviews they conduct and the probability that they hire each type of workers. However, in an equilibrium of the dynamic model the opposite is true as well: the firms' optimal screening process determines the size and composition of the pool of unemployed workers.

3.1 The Dynamic Setting

We assume that there is a measure of workers, normalized to one, and a large measure of firms who can post a vacancy by paying a flow cost of χ per period. All agents are risk neutral and discount between periods at rate $\beta = (1+r)^{-1}$. As in the previous section, we assume that workers are heterogeneous with respect to their underlying characteristics, and we denote the (exogenous) measure of workers with characteristics $\mathbf{x}$ by $m_{\mathbf{x}}$. Given initial conditions, the measure of unemployed workers of type $\mathbf{x}$ in period $t \geq 1$ is denoted by $u_t(\mathbf{x})$.

Each period consists of four phases. The first phase features job destruction. Among those workers who enter the period employed, we assume that a fraction $\delta_{\mathbf{x}}$ of type $\mathbf{x}$ workers receive a separation shock and become unemployed. As is often the case, we assume that a worker who loses her job in period t cannot search for a new job until period $t+1$.[7] The second phase features free entry: in each period t, firms post vacancies until the value of a vacancy, denoted by Π_t, is equal to zero. We let v_t denote the measure of vacancies posted in period t. Given $\lambda_t(\mathbf{x}) = u_t(\mathbf{x})/v_t$, the third phase features the matching process described in Section 2.

The final phase consists of production. In order to adapt our analysis from Section 2, we must redefine the match surplus to account for the value of current *and future* output. To that end, consider a match in period t between a firm and a worker of type $\mathbf{x}$ who has drawn (η, ε). In a slight abuse of notation, let y_t denote the firm's expected discounted value of this match, so that

$$y_t = z_{\mathbf{x}} + \eta + \varepsilon - b + \beta \left[\delta_{\mathbf{x}} \Pi_{t+1} + (1 - \delta_{\mathbf{x}}) y_{t+1}\right].$$

Since $\beta = (1+r)^{-1}$ and, in equilibrium $\Pi_t = 0$ for all $t \geq 0$, it follows that the expected discounted value of the match surplus can be written

$$y_t = y_{t+1} \equiv y = \Delta_{\mathbf{x}} \left(z_{\mathbf{x}} + \eta + \varepsilon - b\right).$$

where

$$\Delta_{\mathbf{x}} = \frac{1+r}{r + \delta_{\mathbf{x}}}$$

denotes the effective discount factor of the match surplus. Note that we allow the heterogeneity in y across workers of type $\mathbf{x}$ to stem from several different sources, including the flow productivity $z_{\mathbf{x}}$, the separation rate $\delta_{\mathbf{x}}$, and the distribution from which the shocks η and ε are drawn.

Following our earlier analysis, we define $y^{\star}(\mathbf{x}, \eta)$ as the solution to (2), with a minor

[7] See, e.g., Menzio and Shi (2011) for a recent example that makes this assumption.

modification to

$$H(y|\mathbf{x},\eta) = F\left(\frac{y}{\Delta_{\mathbf{x}}} - z_{\mathbf{x}} - \eta + b\middle|\mathbf{x}\right). \tag{11}$$

With $K = L \times M$, we again let $Y_1^\star > \ldots > Y_K^\star$ denote the ordered set of all possible values of $Y^\star$.

The distribution of applicants over k in any period $t \geq 0$ can be constructed as in Section 2, with

$$\lambda_{k,t} = g(\eta_k|\mathbf{x}_k)\lambda_t(\mathbf{x}_k). \tag{12}$$

Substituting (11) and (12) into (5), (7), and (8) yields the callback rate $\Gamma_t(\mathbf{x})$, the job-finding rate $\Psi_t(\mathbf{x})$, and the distribution of firm profits $\Phi_t(y)$, respectively. Given these minor modifications, we now show how our analysis from the one-period model can be applied *mutatis mutandis* to characterize equilibrium in the dynamic setting.

3.2 Equilibrium in the Dynamic Setting

In each period $t \geq 0$, firms decide to post a vacancy, taking as given the size and composition of the pool of unemployed, $\{u_t(\mathbf{x})\}_{\mathbf{x}\in\mathcal{X}}$, and the measure of vacancies posted by other firms, v_t. As we established in the one-period model, the (now endogenous) ratios $\lambda_t(\mathbf{x}) = u_t(\mathbf{x})/v_t$ determine firms' hiring rate and expected profits, along with workers' job-finding rates. We now establish that these are key inputs into the definition of an equilibrium in the dynamic model.

To start, we can write the expected discounted value of an unfilled vacancy as

$$\Pi_t = -\chi - c\sum_{\mathbf{x}\in\mathcal{X}} \lambda_t(\mathbf{x})\Gamma_t(\mathbf{x}) + [1-\Phi_t(0)]\int_0^\infty y\frac{d\Phi_t(y)}{1-\Phi_t(0)} + \beta\Phi_t(0)\Pi_{t+1}, \tag{13}$$

where $\Phi_t(y)$ is given by (10). In words, a firm pays an entry cost χ and expected interviewing costs $c\sum_{\mathbf{x}\in\mathcal{X}} \lambda_t(\mathbf{x})\Gamma_t(\mathbf{x})$. With probability $\Phi_t(0)$ the firm does find a match with $y \geq 0$ in period t—and hence does not hire—but can post a vacancy again in the following period $t+1$. With probability $1-\Phi_t(0)$, the firm matches with a productive worker and receives payoff y with conditional probability $\frac{d\Phi_t(y)}{1-\Phi_t(0)}$.

In equilibrium, free entry implies that $\Pi_t = \Pi_{t+1} = 0$ so that, for all $t \geq 0$,

$$\chi + c \sum_{\mathbf{x} \in \mathcal{X}} \lambda_t(\mathbf{x}) \Gamma_t(\mathbf{x}) = \int_0^\infty y d\Phi_t(y). \tag{14}$$

Finally, the evolution of the pool of unemployed workers must satisfy the law of motion:

$$u_{t+1}(\mathbf{x}) - u_t(\mathbf{x}) = [m_{\mathbf{x}} - u_t(\mathbf{x})] \delta_{\mathbf{x}} - u_t(\mathbf{x}) \Psi_t(\mathbf{x}). \tag{15}$$

In words, the first term on the right-hand side of (15) represents the inflow of workers into the pool of unemployed: a measure $m_{\mathbf{x}} - u_t(\mathbf{x})$ of workers of type $\mathbf{x}$ are employed at the beginning of period t, and each of these workers exogenously separates into the pool of unemployed with probability $\delta_{\mathbf{x}}$. The second term represents the outflow of workers from the pool of unemployed, since unemployed workers of type $\mathbf{x}$ match with probability $\Psi_t(\mathbf{x})$.

Definition 1. *An equilibrium is characterized by a time path for the measure of unemployed* $\{u_t(\mathbf{x})\}_{\mathbf{x} \in \mathcal{X}}$ *and the measure of vacancies* v_t*, with associated market tightness* $\lambda_t(\mathbf{x}) = u_t(\mathbf{x})/v_t$*, such that, for all* $t \geq 0$*:*

1. *The measure of vacancies* v_t *satisfies the free entry condition* (14)*;*
2. *The distribution of unemployed workers* $u_t(\mathbf{x})$ *satisfies* (15) *for all* $\mathbf{x} \in \mathcal{X}$*.*

On the surface, the two equations characterizing an equilibrium bear a strong resemblance to standard expressions from the labor search literature. The free entry condition in (14) equates the costs associated with posting a vacancy and searching for a worker with the expected benefits that accrue when the firm matches with one of its applicants. The law of motion in (15) describes how the composition of the pool of unemployed workers evolves as some matches are destroyed and new matches are formed.

What is new to our framework is that the key functions underlying these conditions depend on the firm's interviewing process. Specifically, the firm's optimal screening policy—i.e., which workers to interview and which of the interviewed applicants to hire—determines the functions $\Gamma_t(\mathbf{x})$, $\Psi_t(\mathbf{x})$, and $\Phi_t(y)$. These functions, in turn, determine the expected costs and benefits associated with posting a vacancy, along with the probabilities that each type of worker finds a job.

There are at least three advantages from deriving these key functions endogenously. First, since firms' interviewing decisions depend on the quantity and types of applications they receive, we are able to study the *interaction* between firms' optimal screening process and the composition of the pool of unemployed workers. For example, as firms search more intensely for good matches, they are more likely to hire the most productive workers, shifting

the pool of unemployed to less productive workers. This, in turn, changes firms' incentives to screen workers. As we show below, exploring this interaction between the intensive margin of hiring and the pool of unemployed—in and outside of the steady state—reveals a number of new insights. Second, while the functions $\Gamma_t(\cdot)$, $\Psi_t(\cdot)$, and $\Phi_t(\cdot)$ are key inputs for matching standard labor market moments, like the unemployment rate, our derivations also generate testable predictions for moments that are not targeted in most existing calibrations, such as callback rates across worker types. This is an appealing property of the model since moments like these are becoming increasingly available in data sets derived from worker-level surveys. Lastly, armed with a theory of how firms decide which applicants to interview and hire based on the workers' attributes observed at the application stage, our model provides a new laboratory to study the *equilibrium* implications of policies that change the type of information that firms are allowed to observe or collect at different stages of the hiring process.

4 Incomplete Information

Thus far, we have assumed that the firm perfectly observes all of the workers' payoff-relevant attributes contained in the vector $\mathbf{x}$. We now show that the model can be readily adapted to the case in which only a subset of these attributes are available to the firm.

This extension is important for a few reasons. First, it addresses the more general case in which workers have payoff-relevant attributes that are difficult to communicate through a job application. Second, it allows us to analyze an important class of policies that *restrict* the set of applicant characteristics that can be observed by the firm. A variety of such laws have become commonplace across U.S. states. For instance, there are now bans on firms' use of credit reports (Cortés et al. (2022)), salary history (Davis et al. (2022)), and drug tests (Wozniak (2015)) during various stages of the hiring process. Our focus in Section 5 concerns another area of widespread legislative interest: criminal justice. In particular, we will quantitatively evaluate the effects of Ban the Box, which prohibits firms from observing workers' criminal record at the application stage.

4.1 Setup

Suppose that firms are not able to observe at least one payoff-relevant worker attribute at the application stage of the hiring process. Formally, let us order worker attributes (w.l.o.g.) so that only the first $\widetilde{N} < N$ attributes, $\mathbf{x}^o = (x_1, x_2, ..., x_{\widetilde{N}}) \in \mathcal{X}^o$, are observed at the application stage. The remainding $N - \widetilde{N}$ attributes, $\mathbf{x}^u = (x_{\widetilde{N}+1}, ..., x_N) \in \mathcal{X}^u$, are unobserved.

Importantly, the firm's inability to observe $\mathbf{x}^u$ at the application stage represents the

only departure from the model in Section 2. In particular, we maintain the assumption that the match-specific shock η is observed with the application. However, since this shock is drawn from a type-dependent distribution, $g(\eta|\mathbf{x})$, the realization of η is informative about the applicant's type and, hence, about $z_{\mathbf{x}}$ and the distribution $F(\varepsilon|\mathbf{x})$.

4.2 Optimal search with partial information

Characterizing the firm's optimal interviewing strategy when some information is unobservable requires a minor modification from the full information benchmark studied in Section 2. To start, using Bayes' rule, the probability that an applicant with observable attributes $\mathbf{x}^o$ and realization η is of true type $\mathbf{x} \equiv (\mathbf{x}^o, \mathbf{x}^u)$ is given by

$$\mathbb{P}(\mathbf{x}^u|\mathbf{x}^o, \eta) = \frac{\mathbb{P}\left(\mathbf{x}^u|\mathbf{x}^o\right) g\left(\eta|\mathbf{x}^o, \mathbf{x}^u\right)}{\sum_{\widetilde{\mathbf{x}}^u} \mathbb{P}\left(\widetilde{\mathbf{x}}^u|\mathbf{x}^o\right) g\left(\eta|\mathbf{x}^o, \widetilde{\mathbf{x}}^u\right)}, \tag{16}$$

where

$$\mathbb{P}\left(\mathbf{x}^u|\mathbf{x}^o\right) = \frac{u\left(\mathbf{x}^o, \mathbf{x}^u\right)}{\sum_{\tilde{\mathbf{x}}^u} u\left(\mathbf{x}^o, \tilde{\mathbf{x}}^u\right)} \tag{17}$$

denotes the probability that a randomly selected job seeker with observable attributes $\mathbf{x}^o$ has unobservable attributes $\mathbf{x}^u$.

In a slight abuse of notation, let $y^\star(\mathbf{x}^o, \eta)$ denote the solution to

$$c = \int_{y^\star}^{\infty} (y - y^\star)\, d\hat{H}\left(y|\mathbf{x}^o, \eta\right), \tag{18}$$

where we use $H(\cdot)$ in (11) to now define the distribution of y among the set of applicants with observable type $\mathbf{x}^o$ and realization η by

$$\hat{H}\left(y|\mathbf{x}^o, \eta\right) = \sum_{\tilde{\mathbf{x}}^u} \mathbb{P}\left(\tilde{\mathbf{x}}^u|\mathbf{x}^o, \eta\right) H\left(y|\mathbf{x}, \eta\right). \tag{19}$$

Re-defining the number of *observable* types by $L = |\mathcal{X}^o|$ and again setting $K = L \times M$, we denote the ordered set of possible values of $y^\star(\mathbf{x}^o, \eta)$ by $Y_1^\star > Y_2^\star > \ldots > Y_K^\star$.

4.3 Callback rates, hiring rates, and profits

As in the perfect information benchmark, each $k \in 1, ..., K$ is associated with a pair $(\mathbf{x}_k^o, \eta_k)$. Letting $H_k(y) = \hat{H}\left(y|\mathbf{x}_k^o, \eta_k\right)$, the definitions of Γ_k, Φ_k, and $\Phi_k(y)$ are essentially unchanged from (5), (7), and (8), respectively.

Since all workers with the same $(\mathbf{x}^o, \eta)$ are equally likely to be interviewed, the callback

rate of a type $\mathbf{x} = (\mathbf{x}^o, \mathbf{x}^u)$ worker can be calculated with a only a minor modification:

$$\Gamma(\mathbf{x}^o, \mathbf{x}^u) = \sum_{k=1}^{K} g(\eta_k|\mathbf{x}) \mathbb{I}\{\mathbf{x}_k^o = \mathbf{x}^o\} \Gamma_k.$$

However, since all applicants with the same $(\mathbf{x}^o, \eta)$ are *not* equally likely to be hired, we need to make an important adjustment to calculate the job-finding rate of a worker of type $\mathbf{x}$. To do so, it will be helpful to define the probability that an applicant with observable attribute $\mathbf{x}^o$ and shock η who draws output y is of unoberved type $\mathbf{x}^u$. A simple application of Bayes' rule reveals:

$$\mathbb{P}(\mathbf{x}^u|\mathbf{x}^o, \eta, y) = \frac{u(\mathbf{x}^o, \mathbf{x}^u) g(\eta|\mathbf{x}^o, \mathbf{x}^u) dH(y|\mathbf{x}^o, \mathbf{x}^u)}{\sum_{\tilde{\mathbf{x}}^u} u(\mathbf{x}^o, \tilde{\mathbf{x}}^u) g(\eta|\mathbf{x}^o, \tilde{\mathbf{x}}^u) dH(y|\mathbf{x}^o, \tilde{\mathbf{x}}^u)}. \tag{20}$$

Using equation (20), the probability that a firm hires a worker with type $(\mathbf{x}_k^o, \mathbf{x}^u)$ and η_k can be calculated as

$$\Phi_k(\mathbf{x}^u) = \int \mathbb{P}(\mathbf{x}^u|\mathbf{x}_k^o, \eta_k, y) d\Phi_k(y).$$

The probability that a worker of type $(\mathbf{x}_k^o, \mathbf{x}^u)$ and η_k is hired then equals

$$\Psi_k(\mathbf{x}^u) = \frac{\Phi_k(\mathbf{x}^u)}{\lambda_k(\mathbf{x}^u)},$$

where $\lambda_k(\mathbf{x}^u) = \lambda(\mathbf{x}_k^o, \mathbf{x}^u) g(\eta_k|\mathbf{x}_k^o, \mathbf{x}^u)$. The matching probability for a worker of type $\mathbf{x} = (\mathbf{x}^o, \mathbf{x}^u)$ is now given by

$$\Psi(\mathbf{x}^o, \mathbf{x}^u) = \sum_{k=1}^{K} g(\eta_k|\mathbf{x}) \mathbb{I}\{\mathbf{x}_k^o = \mathbf{x}^o\} \Psi_k(\mathbf{x}^u). \tag{21}$$

5 Quantitative Application: Ban the Box

Over 35 states have enacted Ban the Box (BTB) laws that regulate firms' access to information regarding job applicants' criminal records. Moreover, in many of these states, BTB laws have a rather subtle feature for which our model is particularly well-suited to study: a candidate's criminal history is shielded only through the *application* stage, but may be accessed by the firm at (or after) the *interview* stage. Hence, understanding the effects of BTB laws requires an equilibrium model with distinct application and interviewing stages, where the choice of which applicants to interview has a well-defined allocative role.

In what follows, we apply our framework to study the effects of the introduction of Ban the Box. As a first step, we calibrate the model to the *pre*-Ban the Box environment in which

all worker attributes are perfectly observed. Next, leveraging the extended model outlined in Section 4, we show quantitatively how firms' interviewing policies adjust when criminal record is hidden at the application stage. We also trace through the implications of this change for firms' vacancy posting decisions, worker job finding rates, and unemployment.

5.1 Worker Types

The first task of the calibration is to define the discrete worker attributes that make up the vector $\mathbf{x}$. For our application, there are two particularly salient characteristics: a worker's criminal record and race.

We distill a "criminal record" down to a binary variable: a worker is said to have a record if he has ever been convicted of a crime. Admittedly, this represents a substantial simplification since the vector of attributes in the model could encompass many dimensions of a criminal record, including the number, types, and dates of all offenses. We make this simplification for two reasons. The first is that there are readily available, and credible, estimates of the incidence of a past conviction based on large-scale administrative datasets. More detailed estimates—i.e., by the type and number of offenses—are subject to considerably more uncertainty. The second reason is parsimony: to highlight the mechanisms in the model, and how they shape the dynamic response to Ban the Box, it is helpful to restrict the set of attributes to the fewest required for this application.

The second attribute is race. While the model could accommodate a rich specification of race, we focus again on a binary indicator: an applicant is either black or white. We make this simplification because the most credible empirical evidence on interview rates is available only for these two races. Thus, our choice better aligns the model with the data that we will use for calibration.

In summary, the vector $\mathbf{x} = (x_1, x_2)$ consists of two attributes: (1) race, which is represented by attribute x_1; and (2) criminal record, which is represented by attribute x_2. In what follows, let $x_1 = \mathcal{B}$ if a worker is black and $x_1 = \mathcal{W}$ if white; and let $x_2 = 1$ if a worker has a past conviction and $x_2 = 0$ if a nonoffender. Thus, in total, there are four types of applicants: black exoffenders and nonoffenders, and white exoffenders and nonoffenders.

5.2 Externally Calibrated Parameters

As is typical, we partition the set of parameters into two groups. One consists of parameters with clear empirical analogues whose values can be set before the model is solved. The second consists of parameters whose values are calibrated "internally" so that the model reproduces certain observed moments. We will begin our discussion with the former.

Among the parameters that are set externally, the first to consider is the discount factor of the firm, β. We set β in line with a 5 percent annual real interest rate. At biweekly frequency, this implies an interest rate of 0.0019, so that $\beta = 1/(1 + 0.0019) \approx 0.998$.[8]

The remainder of the externally calibrated parameters fall into two classes. The first is the population shares $m(\mathbf{x})$. We restrict attention to men, since conviction rates of women of all races have generally been very low.[9] As noted above, we also restrict the population to two races, black and white. The size of each racial group is taken from Census Bureau population estimates of self-identified black and white male respondents. Thus, the shares sum to 1.

Within each race, we also must pin down the share with at least one past conviction, which we refer to as the *exoffender share*. We utilize two sources to estimate these shares. First, Shannon et al. (2017) estimate the incidence of a past felony conviction based on prisoner release data—prison sentences are only given for felonies—and outflows from other forms of felony supervision (e.g., probation).[10] The authors' tabulations imply that approximately 40 percent of adult black men are felony exoffenders, whereas roughly 10 percent of adult white men have at least one felony conviction.

These estimates neglect the incidence of nonfelony, or misdemeanor, convictions.[11] Therefore, we supplement estimates from Shannon et al with data from the National Longitudinal Survey of Youth 1997 (NLSY97), which provides self-reported felony and misdemeanor convictions. In fact, we could have relied solely on the NLSY97, but we found that, relative to Shannon et al, the NLSY97 understates by half the felony exoffender share of black men. It is possible (and perhaps probable) that this understatement extends to nonfelony (misdemeanor) convictions. For this reason, we apply a simple "correction" to the NLSY97 based on Shannon et al: for black men, we scale up the overall exoffender share in the NLSY97 by a factor of two. This measures the share *ever* convicted for *any* (felony or nonfelony) offense.

The difference in exoffender shares by race is stark: roughly 48 percent of black men are exoffenders, whereas the exoffender share of white men is nearly 20 percent. This racial gap in offense rates will play a crucial role in the quantitative model. Notably, the former estimate is not much higher than the felony exoffender share based on Shannon et al. This result reflects

[8] The time period in search models is typically short to better capture the high rates of turnover in the U.S. labor market. We choose a biweekly frequency, though we will sometimes translate results into monthly rates for ease of interpretation.

[9] In this sense, Ban the Box laws do not conceal information about female applicants. Nevertheless, female labor market activity is not necessarily invariant to the ban. Rather, firms could favor women since the latter can be identified almost surely as nonoffenders even under the ban. We return to this point later.

[10] To prevent double-counting of exoffenders with multiple convictions, these data are adjusted using estimates of recidivism.

[11] Agan et al. (2024) find that even a misdemeanor conviction significantly reduces the probability of future employment. This result is consistent with firms conditioning hiring on nonfelony convictions.

that, among black respondents in the NLSY97, most misdemeanor offenders have a felony conviction. Therefore, the addition of adults ever convicted of only a misdemeanor does not dramatically increase the overall black exoffender share. Proportionately, the addition of misdemeanors has a more significant effect on the white exoffender share.

We are now in a position to report the population shares. White nonoffenders constitute the largest type: they make up around 67 percent of this population, $m(\mathcal{W}, 0) = 0.674$. White exoffenders make up 14 percent. Of the remainder, 9.5 percent is comprised of black nonoffenders. Just under 9 percent is made up of black exoffenders (specifically, $m(\mathcal{B}, 1) = 0.088$).

A second class of parameters consists of the type-specific separation rates, $\delta_{\mathbf{x}}$. In general, we could assign a different separation rate to each of the four types. However, there is very limited evidence on separation rates conditional on a criminal conviction, and the few available estimates are derived from specialized populations. These observations alone make us hesitant to calibrate separation rates separately by criminal record.

Morover, among the few published results, there does not appear to be a significant gap between exoffenders and nonoffenders. Minor et al. (2018) examined human resources records from low-wage, white collar jobs, the most typical of which were in call centers. The authors find that exoffenders' separation rates are *lower* but not dramatically so: in their sample of high-turnover jobs, the implied monthly separation rates of exoffenders and nonoffenders were around 15 and 16 percent, respectively. Lundquist et al. (2018) is perhaps the only other paper that examines employment records of exoffenders, but they report results only for separations initiated due to poor performance. They follow U.S. military enlistees and find that the rate of separation for cause is essentially independent of criminal record.

Given these results, we elect not to differentiate separation rates by criminal record. Rather, we set the same separation rate for all workers of a given race. Empirically, we measure the separation rate as the average employment-to-nonemployment transition rate in the longitudinally matched Current Population Survey micro data.[12] We estimate that black employees separate at a rate of roughly 2.5 percent per month, which is twice as high as the rate observed among white employees.

5.3 Baseline Calibration

The model features a rich payoff structure. The match value consists of three components, two of which are drawn from distributions that potentially depend on the full vector of worker attributes. In general, we see this scope for heterogeneity as an important feature

[12] As discussed below, we pool the unemployed and (most) nonparticipants to make up the nonemployed.

of the model. However, for the application at hand, the available data is somewhat limited. Therefore, we will need to make several restrictions on payoffs to be just identified.[13]

Our point of departure is a normalization, namely, we fix the expected flow profit from a match with a white nonoffender such that the expected discounted *present value* of the match is one, $E[y|\mathbf{x} = (\mathcal{W}, 0)] = 1$. Recall from (1) that expected flow profit is $z(\mathbf{x}) - b$ (since η and ε are mean-zero random variables). Since b is treated as a constant, we will, with some abuse of notation, typically express the properties of expected flow profit in terms of *only* $z(\mathbf{x})$.

It is convenient to normalize with respect to white nonoffenders for a simple reason. Since this type has the highest interview rate in the data, it will follow that the expected present value surplus of all other types will be bounded above by one. In addition, if we anticipate, plausibly, that the *average* surplus of any type is positive, we may also bound $z(\mathbf{x})$ below by zero for all $\mathbf{x}$.

Next, to reduce the dimension of the parameter space, we restrict differences in average surplus *conditional on criminal record*. Given the normalization for white nonoffenders, a simple option is to assume the same average surplus for *all* nonoffenders. Thus, we fix the flow surplus of black nonoffenders such that their mean present value payoff to the firm is also one.[14]

Insofar as a restriction is needed, average surplus is a relatively attractive target. This approach aligns with a dominant view in the theoretical literature that models of racial groups should not, in general, be grounded in assumptions about average differences (in ability, for instance). Rather, Fang and Moro (2011) advocate for models where employers' *uncertainty* over payoffs varies by applicant type, a strategy that we will also pursue and describe shortly.

Of course, this argument could be extended to assume no average racial differences among exoffenders as well. We are sympathetic to this view and plan to pursue this alternative. For the moment, since a further restriction along these lines is not needed, we do not make one. Thus, we will choose parameters $z(\mathcal{B}, 1)$ and $z(\mathcal{W}, 1)$.

Next, we return to the uncertainty over payoffs *given* an application. This is regulated by the variance of ε, the component of the payoff revealed at the interview. Fang and Moro, among others, have argued that the precision of employers' beliefs is likely to vary by applicant type. For instance, employers may be less familiar with the quality of educational and employment experiences of minority applicants. In this context, interviews may be especially informative. More formally, the component ε would account for a relatively large

[13] We will use "firm payoff" and "firm profit" interchangably with "match surplus" since the wage is fixed at b.

[14] Since separation types vary by race, the flow payoff does vary by race such that the expected present value is one for both groups.

share of the variance of payoffs of minority applicants. To capture this idea, we allow that the variance of ε depends on race.

In our current calibration, in fact, we assume the variance of ε is uniform *within* race, that is, it does not vary by criminal record. As we will discuss, we can accommodate a richer structure of ε but at a cost: some other restriction(s) must be placed elsewhere on the distribution of payoffs.

For now, though, our restriction on the distribution of ε implies just two additional parameters to identify.[15] The mapping between these parameters and the moments will be better understood by thinking of the variance of ε as having a *baseline*, or common, component, which will be denoted by $\bar{\sigma}_\varepsilon$. We then introduce a second parameter, $\hat{\sigma}_\varepsilon(\mathcal{B})$, to measure the difference in dispersion between black and white applicants. The standard deviation faced by black workers is then $\sigma_\varepsilon(\mathcal{B}) = \bar{\sigma}_\varepsilon + \hat{\sigma}_\varepsilon(\mathcal{B})$. Note that, in this rendering, $\bar{\sigma}_\varepsilon$ coincides exactly with the standard deviation among white applicants, $\sigma_\varepsilon(\mathcal{W}) = \bar{\sigma}_\varepsilon$.

The final component of payoffs is η, which is drawn at the application stage. Given restrictions on z and ε, we can be somewhat more general in our specification of the distribution of η. Specifically, we allow that the variance of η differs by race and criminal record, but in a certain, controlled manner.

Our calibration of the distribution of η begins with the choice of a parameter $\bar{\sigma}_\eta$, which is the baseline standard deviation of η. This parameter coincides with the standard deviation of η for white nonoffenders but also represents the common component of the dispersion across all types. The standard deviation of η for black nonoffenders is then constructed according to

$$\sigma_\eta(\mathcal{B}, 0) = \bar{\sigma}_\eta + \hat{\sigma}_\eta(\mathcal{B}),$$

where $\hat{\sigma}_\eta(\mathcal{B})$ is a free parameter that will measure the difference in standard deviations between black and white applicants for a given criminal history. Similarly, the standard deviation of η for white exoffenders is constructed according to

$$\sigma_\eta(\mathcal{W}, 1) = \bar{\sigma}_\eta + \hat{\sigma}_\eta(1),$$

where $\hat{\sigma}_\eta(1)$ is a free parameter that will measure the difference in standard deviations between exoffenders and nonoffenders for a given race. Finally, the standard deviation of η for black exoffenders is

$$\sigma_\eta(\mathcal{B}, 1) = \bar{\sigma}_\eta + \hat{\sigma}_\eta(\mathcal{B}) + \hat{\sigma}_\eta(1).$$

[15] Observe, furthermore, that this restriction will simplify the implementation of Ban the Box: if $\text{var}(\varepsilon)$ is independent of criminal record, the introduction of Ban the Box does not affect the firm's inference of an applicant's ε

One may now confirm that, for any race $x_1 = (\mathcal{B}, \mathcal{W})$, the difference between $\sigma_\eta(x_1, 1)$ and $\sigma_\eta(x_1, 0)$ depends only on $\hat{\sigma}_\eta(1)$, and, for any criminal record $x_2 = (0, 1)$, the difference between $\sigma_\eta(\mathcal{B}, x_2)$ and $\sigma_\eta(\mathcal{W}, x_2)$ depends only on $\hat{\sigma}_\eta(\mathcal{B})$. In total, we introduce three new parameters—namely, $\bar{\sigma}_\eta$, $\hat{\sigma}_\eta(\mathcal{B})$, and $\hat{\sigma}_\eta(1)$.

Importantly, this specification would have to be simplified if we sought a more general distribution of z and/or ε. For instance, if we allowed that ε varies by race and criminal record, we would have to "move" a free parameter from the distribution of η to the distribution of ε. In that context, we would eliminate $\hat{\sigma}_\eta(1)$.

In addition to payoff-related parameters, there are two parameters that represent the main frictions in the model: the marginal cost of an interview, c, and the parameter χ that captures all other (flow) costs of recruitment. This brings the total number of unknown parameters to nine. We next describe the moments we will use to identify these parameters.

5.4 Moments and Identification

As the interview decision is at the core of the firm's problem, the model naturally generates predictions for several interview-related moments that we aim to match. The first is the average number of interviews per applicant, or the probability that any given applicant is interviewed. Barron et al. (1997) estimate this to be a little over one third.

The model also yields a cross-section of interview probabilities over worker types, which have empirical analogues in Agan and Starr (2018). However, the levels of the probabilities in Agan and Starr will not necessarily be reconcilable with the levels of other labor market aggregates from very distinct sources (see below). Therefore, we target interview probabilities *relative* to the rate at which white nonoffenders are interviewed. This adds three moments to the list of targets. As shown in Table 1, Agan and Starr find that exoffender applicants of both races are interviewed at approximately 60 percent of the rate experienced by white nonoffenders, while black nonoffenders are interviewed only slightly less than their white counterparts.

Another core function of the model is to connect interview activity to employment outcomes. This observation suggests several outcomes of interest. To start, we seek to match nonemployment rates by race. The nonemployment rate—the complement of the employment rate—is defined in this context as the jobless share of men of a given race between the ages of 18-34. We choose this age group to zero in on the portion of the lifecycle where criminal activity is relatively likely (Rocque et al. (2016)).[16] We exclude students but do not otherwise drop non-labor force participants, a point to which we return below. These

[16] Criminal offending has generally been found to peak during ages 18-24. Using more recent data, Lofstrom et al. (2023) document a shift in the distribution of offending toward adults in their late 20s and early 30s.

moments are readily computed from Current Population Survey data. The black nonemployment rate of nearly 25 percent is roughly twice as high as the analogue for white men (see again Table 1).

Employment rates by conviction history are more difficult to observe. We rely on evidence from Agan et al. (2024), who have administrative panel data on offenders from four states. Their results imply that a conviction reduces the probability of employment by roughly 11.5 percent, or 8.5 percentage points. We target this (percent) difference in employment rates between ex- and nonoffenders, which we term the employment penalty of a record.

Table 1: Parameters and targeted moments

Parameter	**Description**	**Value**	**Closest Moment**	**Value**	**Source**
$z(\mathcal{W},1)$	Avg surplus $(\mathcal{W},1)$	0.815	Relative callback rate $(\mathcal{W},1)$	0.593	Agan & Starr (2018)
$z(\mathcal{B},1)$	Avg surplus $(\mathcal{B},1)$	0.847	Relative callback rate $(\mathcal{B},1)$	0.614	Agan & Starr (2018)
$\sigma_\varepsilon(\mathcal{B})$	Std dev ε black	0.510	Relative callback rate $(\mathcal{B},0)$	0.936	Agan & Starr (2018)
$\sigma_\varepsilon(\mathcal{W})$	Std dev ε white	0.528	P(interview) all applicants	0.373	Barron et al. (1997)
$\sigma_\eta(\mathcal{W},1)$	Std dev η $(\mathcal{W},1)$	0.263	Employment gap btw. x_2=0 & x_2=1	0.116	Agan et al. (2024)
$\sigma_\eta(\mathcal{B},0)$	Std dev η $(\mathcal{B},0)$	0.347	Black nonemployment rate	0.251	CPS
χ	Entry cost	1.809	White nonemployment rate	0.125	CPS
c	Cost of interview	0.018	$\frac{\text{hiring cost}}{\text{market value/worker}}$	0.044	Villena Roldan (2012)

Note: This table presents the calibrated parameters and reports a targeted moment to which each parameter is closely associated. A worker's race is denoted by $x_1 = \mathcal{B}$ if black and $x_1 = \mathcal{W}$ if white. A worker's criminal history is denoted by $x_2 = 1$ if he has a conviction and $x_2 = 0$ otherwise. The relative callback rate of a type measures the interview rate of that type divided by the interview rate of white nonoffenders ($\mathbf{x} = (\mathcal{W}, 0)$). The employment gap between nonoffenders ($x_2 = 0$) and exoffenders ($x_2 = 1$) is the percent difference between the employment rate of the former and the latter. Finally, the target corresponding to the interviewing cost, c, expresses the cost of hiring a worker, taken from Villena Roldan (2012), relative to an estimate of the market value per worker from Compustat.

Combining estimates from the CPS and Agan et al, we have three employment-related moments. It is worth noting here that the use of employment, rather than unemployment, enables this pooling of multiple data sources. The administrative records have no measure

of labor force participation, ruling out the measurement of unemployment.[17]

The final targeted moment concerns the marginal cost of interviewing, c. Given the role of interviewing in the model, this moment is a natural one to target. However, we face several challenges in deriving an estimate of the latter.

First, while there are estimates of the cost of recruiting in general, it is far more difficult to track down estimates of the cost of interviewing more specifically. We have found just one source that plausibly isolates interviewing costs, but this is from a now 45-year old survey. (We discuss this estimate at greater length below.)

Second, it is not necessarily clear if an estimate of the cost of interviewing captures the economic meaning of c in the model. For instance, from the perspective of the model, c is perhaps better understood more broadly as a cost of screening, which includes all the effort that a firm puts into assessing the suitability of an applicant. It is difficult to infer from the survey questions if this concept of a screening cost is being measured.

Third, estimates of recruiting or interviewing costs are typically expressed in hours of time or as a share of the wage bill. In our model, though, the wage is not separately identified from the rest of the flow value of a match and, moreover, it is the *present value* of the match that is the best measure of the firm's payoff. Thus, from the model's point of view, it is most natural to express recruiting costs as a share of the market value of the worker, which is not readily available from the literature.

In designing a response to these challenges, one important consideration is how the model dynamics depend on the choice of c. We will find (numerically) that, on many margins, the labor market response to Ban the Box is increasing in c. This finding is intuitive. With the Ban in place, more effort is required to find the best match, but if the marginal cost of screening is low, this additional effort will be supplied with little disruption to hiring. Exploiting this monotonicity in the model, we can bound the effects of the Ban by considering a reasonable range of values of c. We will start by implementing what is arguably an upper bound from Villena Roldan (2012), who estimates a cost of recruiting per hire of \$4200. Our understanding from Villena Roldan is that this estimate almost surely includes non-interviewing-related costs.

The next challenge is to express this estimate in terms of the market value of a worker. To this end, we use the following identity,

$$\frac{\text{cost of hiring}}{\text{market value per worker}} = \frac{\text{cost of hiring}}{\text{wage per worker}} \times \frac{\text{wage per worker}}{\text{revenue per worker}} \times \frac{\text{revenue per worker}}{\text{market value per worker}}.$$

[17] One other concern with the use of unemployment is that the line between it and nonparticipation is especially blurry for young men. For instance, the monthly nonparticipation-to-employment transition rate is only 2 percent in the general population but 10 percent for young men.

From Villena Roldan (2012), we have an estimate of the cost of hiring on the far left-hand side. Using data on compensation from the NIPAs, we calculate that this cost represents just over 10 percent of the average wage. Next, we assume a labor share of two thirds, which fills in the middle term. Finally, we draw on Compustat data for publicly traded U.S. firms to estimate the term on the far right-hand side. We find that, on average, revenue as a share of market value is also two thirds. It follows that the cost of hiring is on the order of 4.5 percent of market value per worker.

Taking stock, we have eight moments corresponding to the parameters in Table 1. However, there is one parameter that is unlisted, namely, the baseline dispersion of η, $\bar{\sigma}_\eta$. Lacking a moment that can compellingly identify this parameter, we will examine how the effects of the Ban vary over a range of values of $\bar{\sigma}_\eta$. This more exploratory approach also serves another purpose: the manner in which the equilibrium outcomes vary with the choice of $\bar{\sigma}_\eta$ will help elucidate the key economic mechanisms in the model.

While there is no "right" choice for the omitted parameter, there are several reasons why we designate this to be $\bar{\sigma}_\eta$. The first reason is that, since $\bar{\sigma}_\eta$ feeds into the dispersion of every type, the mapping between it and the targeted labor market outcomes is less clear-cut. Accordingly, the identification of $\bar{\sigma}_\eta$ is more fraught. By contrast, it is far easier to understand the mapping between a parameter and a moment when each is *specific* to one type — for example, mapping z for a white exoffender to the relative callback rate of white exoffenders. The second reason is that, given a different degree of dispersion in η, the calibration will recover a different degree of dispersion in the values of z. Specifically, we find that an increase in $\bar{\sigma}_\eta$ implies an increase in the dispersion of z, a point we discuss further below. Recalling that z is plausibly bounded between zero and one, we can then trace out virtually the entire feasible range of z by varying $\bar{\sigma}_\eta$.

The remaining eight parameters have a reasonably clear relationship with the eight moments that we target. To start, consider the cross section of callback rates over worker types. An applicant of type $\mathbf{x}$ is more likely to be interviewed if her expected payoff as of the application stage is relatively high.[18] All else equal, this tends to occur when $z(\mathbf{x})$ is relatively large. Accordingly, the values of $z(\mathcal{B}, 1)$ and $z(\mathcal{W}, 1)$ are highly informed by the relative callback rates of these types, as measured by Agan and Starr (2018).

Next, consider the uncertainty over payoffs as of the application stage. This uncertainty is what provides the impetus to interview applicants. Hence, the overall dispersion in ε, denoted by $\bar{\sigma}_\varepsilon$, shapes the probability that *any* given applicant is interviewed. Moreover, the *relative* callback rate of black applicants is strongly influenced by the difference in dispersion

[18] As we stressed earlier, the ranking of applicants accounts for more than the mean payoff, but a higher mean all else equal will elevate the probability of being interviewed.

between black and white applicants, which is governed by $\hat{\sigma}_\varepsilon(\mathcal{B})$. Given $\bar{\sigma}_\varepsilon$, the value of $\hat{\sigma}_\varepsilon(\mathcal{B})$ of course pins down the overall dispersion among black workers, $\sigma_\varepsilon(\mathcal{B})$.[19]

The final set of payoff-related parameters concern the distribution of η. A higher variance of η for a given type implies a higher expected maximum payoff from that type. Thus, it is more likely that an applicant from that type converts an interview into a job offer. For this reason, these dispersion parameters are highly informed by the moments on nonemployment. For instance, the parameter $\hat{\sigma}_\eta(\mathcal{B})$ is most closely related to the difference in nonemployment rates between black and white workers. Likewise, $\hat{\sigma}_\eta(1)$ is most closely related to the difference in nonemployment rates between nonoffenders and exoffenders. Note that, given $\bar{\sigma}_\eta = \sigma_\eta(\mathcal{W}, 0)$, the choice of $\hat{\sigma}_\eta(1)$ pins down the overall dispersion among white exoffenders, $\sigma_\eta(\mathcal{W}, 1)$.

Two parameters remain, namely, the parameters that govern labor market frictions. The cost χ regulates the rate of job creation, which in turn shapes aggregate nonemployment. Equivalently, we can think of choosing χ to reproduce white nonemployment, with $\hat{\sigma}_\eta(\mathcal{B})$ then pinning down black nonemployment. Identifying $\hat{\sigma}_\eta(\mathcal{B})$, in turn, yields the overall dispersion of η for black nonoffenders. Lastly, the marginal cost of an interview, c, is most closely related to the estimated cost of hiring (expressed in terms of the market value of a worker).

5.5 Calibration Results

We begin with a relatively low $\bar{\sigma}_\eta = \sigma_\eta(\mathcal{W}, 0)$ and, later, consider what will happen if we raise it. The parameters under this calibration are reported in Table 1. Since the match value is the present value of profits, we find it more natural to express the payoff parameters in terms of their present values. Therefore, for any type $\mathbf{x}$, values of $z(\mathbf{x})$ and the standard deviations of η and ε are scaled by $\Delta_{\mathbf{X}} = (1 + r)/(r + \delta(\mathbf{x}))$. To illustrate, our choice of $\bar{\sigma}_\eta$ implies that, for white nonoffenders, the standard deviation of the present value of η equals 0.325.

This choice of $\bar{\sigma}_\eta$ implies that the expected present value of a match with an exoffender is roughly 0.8. Since the analogue for nonoffenders is one, the dispersion in z in this calibration is relatively low. More generally, we find that dispersion in η and z go hand in hand: higher values of $\bar{\sigma}_\eta$ imply more dispersion in z, e.g., lower values of z for exoffenders. Intuitively, greater dispersion in η enables low-z exoffenders to leapfrog higher-z nonoffenders and thereby increases exoffender interviews (and hires). To match exoffenders' relative interview (and employment) rates, the calibration infers that z must be lower for exoffenders.

A few other aspects of the calibration are worth highlighting. First, within a given racial

[19] In Table 1, we present the levels of the standard deviations, such as $\sigma_\varepsilon(\mathcal{B})$, rather than the relative standard deviations, such as $\hat{\sigma}_\varepsilon(\mathcal{B})$.

group, exoffender applicants face a slightly less disperse distribution of η. Intuitively, a reduction in dispersion makes it less likely that an applicant will draw the highest payoff of the applicant pool. Hence, our result implies that the relatively high value of z needed to match exoffender interview rates implies too little nonemployment relative to nonoffenders. This tension is resolved by a reduction in the variance of η among exoffenders. Conversely, we find that a higher variance of η is consistent with the observed black nonemployment rate.

An important factor behind these results is the higher separation rate of black workers. Even if the model generates the "correct" interview rate of black applicants, the higher rate of separation diminishes the probability (all else equal) that a black applicant converts an interview into a job offer. In the model, this effect is so substantial that the model "needs" a higher variance of η among black applicants to moderate the racial employment gap. Importantly, the separation rate in this calibration varies only by race and not by criminal record. Essentially, this means that exoffenders convert too many interviews into jobs. A lower variance of η among exoffenders is "needed" to correct the nonemployment rates of exoffenders.

Second, we turn to ε, the other stochastic element of the payoff. In present value terms, the standard deviations of are very similar across race. To see why, recall that the relative uncertainty over ε is an important determinant of the relative callback rate. Intuitively, more uncertainty increases the incentive to interview. Since the callback rates of white and black applicants are nearly identical conditional on criminal record, the model infers that the degree of uncertainty over ε must be very similar.

5.6 Ban the Box

In the calibrated model, Ban the Box is implemented as an unanticipated and permanent change in policy governing the worker attributes that firms observe at the application stage. Specifically, the Ban conceals a worker's criminal record. This is the only restriction on a firm's information set at the application stage. Thus, we maintain that firms observe a worker's race as well as the match specific draw η.

Beginning from the pre-Ban steady state, we solve and simulate the dynamic paths of labor market aggregates after the imposition of the Ban. The out-of-steady-state dynamics are generally challenging to compute in random search models with heterogeneous agents. There are two features of our model environment that facilitate the out-of-steady-state solution.

First, since firms make take-it-or-leave-it offers, all workers earn the same constant wage. As a result, a firm does not have to forecast the path of future wages in order to calculate the expected payoff from filling a position this period. This prediction problem would have to be conditioned on a forecast for the entire future path of $\lambda(\mathbf{x})$ for each $\mathbf{x}$.

Second, we assume free entry. Each firm anticipates that other firms will create positions until the expected payoff from filling a vacancy is equated to χ in all periods. As a result, a firm considering creating a position now or later faces the same net payoff (of zero). In this sense, free entry essentially eliminates the intertemporal problem of timing entry. Otherwise, a firm would have to form beliefs about the entire path of expected payoffs.

Proceeding, we begin our discussion of results by inspecting the interviewing margin. Figure 1 presents the change in interview probabilities for each worker type. The black dashed lines illustrate the estimated effects of the Ban in Agan and Starr (2018). Importantly, these effects were not targeted in the calibration. While the model does understate the jump in exoffender interview probabilities, it nevertheless predicts a substantial increase that is comparable to the estimated responses. With respect to nonoffenders, the results are decidedly mixed. The model nearly matches, but slightly overstates, the decline in the black nonoffender interview probability. By contrast, the estimated response of the white nonoffender interview probability is very difficult to rationalize. Agan and Starr find that this probability increases after the Ban. In any model with Bayesian updating, this probability will almost surely fall since high-surplus nonoffenders are disadvantaged when they are pooled with low-surplus exoffenders. We return to this point in the conclusion of the paper.

A crucial feature of the model-predicted interview probabilities is that the gains for exoffenders, and losses for nonoffenders, are not uniform across race. In particular, the increase for black exoffenders is attenuated relative to gain for white exoffenders, and the decrease for black nonoffenders is amplified relative to the loss for white nonoffenders. This pattern reflects statistical discrimination: since the exoffender share of black applicants is relatively large, firms will tend to somewhat favor white applicants for interviews.

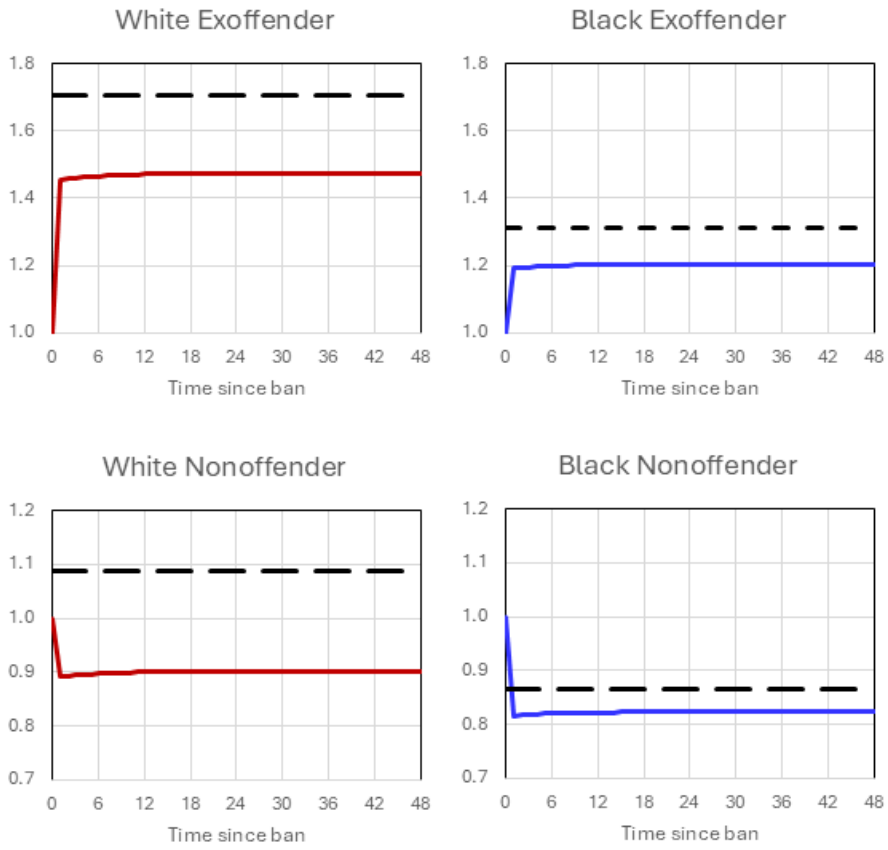
Formally, these findings reflect a change in the optimal interviewing policy, that is, in the reservation values y^*. A heuristic way of understanding what happens can be seen by first reconsidering optimal interviewing under perfect information. In this setting, since both z and η are observed, one can show that the reservation threshold y^* can be decoupled: it equals the sum of the known components of the payoff and a threshold value for the unknown component, namely, ε. Thus, under perfect information,[20]

$$y^*(\mathbf{x}, \eta) \;=\; z(\mathbf{x}) + \eta + \varepsilon^*(\mathbf{x}).$$

Under Ban the Box, this decoupling does not strictly hold, but it is still a useful device for interpreting y^*. Specifically, if the distribution of ε depends only on the *observed* attributes $\mathbf{x}^o$, then the most direct effect of the Ban is to alter the firm's estimate of z. It turns

[20] As noted, we interpret z here as net of b, so we omit the latter from this expression.

Figure 1: Effect of Ban the Box on interview rates



Note: This plots the change in the interview, or callback, rate of each type. Time period 0 refers to the pre-Ban steady state. The pre-Ban interview rate is normalized to 1. The dashed line presents the estimates from Agan and Starr (2018).

out that the revised reservation values can then be approximated reasonably well by simply applying Bayes' rule to z, taking as given the original values $\varepsilon^*(\mathbf{x}) = \varepsilon^*(\mathbf{x}^o)$. This yields an (approximate) reservation rule under the Ban,

$$y^*(\mathbf{x}^o, \eta) \;\approx\; \sum_{\tilde{\mathbf{x}}^u} \mathbb{P}\left(\mathbf{x}^u | \mathbf{x}^o, \eta\right) z(\tilde{\mathbf{x}}^u, \mathbf{x}^o) + \eta + \varepsilon^*(\mathbf{x}^o)\,,$$

where $\mathbb{P}\left(\mathbf{x}^u | \mathbf{x}^o, \eta\right)$ is given by (16).

This formulation helps to illustrate the mechanics behind the changes in the interviewing probabilities. First, exoffenders experience an increase in their reservation values, which are now constructed by averaging over nonoffenders' high z and exoffenders' low z. Therefore, exoffender interview probabilities rise. Conversely, nonoffender interview probabilities fall. Next, interview probabilities will rise more for white exoffenders even if z does not differ by race conditional on criminal record. The reason is that the probability of encountering an exoffender is higher among black applicants, which implies that the optimal estimate of z is lower for black applicants.

We turn next to the dynamics of nonemployment by type. Qualitatively, the path of nonemployment for each worker type mirrors the response of the interview probability. For instance, exoffender nonemployment rates fall but at the expense of higher nonoffender nonemployment. A more interesting observation concerns the quantitative magnitudes: a 20 percent increase over the long run in the interview probability of black exoffenders is consistent with just a 1.5 p.p. (or, 5 percent) decrease in their nonemployment rate. Likewise, an almost 50 percent increase in the interview probability of white exoffenders is consistent with a little over 2 p.p. (or almost 12 percent) decrease in their nonemployment rate. These changes in nonemployment are in turn supported by fairly modest changes in the job finding probability, which increases by 0.4 p.p. for black exoffenders and 0.8 p.p. for white exoffenders.

In other words, the model can generate (realistically) large changes in callback rates without predicting comparably large changes in employment. This strikes us as a potentially appealing property of the model. While there is no consensus in the reduced form literature, it is fair to say that large employment effects are unlikely. Indeed, the model responses, while seemingly modest in size, are still on the higher end of the range of estimates. Jackson and Zhao (2017) and Rose (2021) fail to recover any significant positive impact on exoffender employment. By contrast, Shoag and Veuger (2021) and Pyle (2024) report positive estimates, with Pyle finding that employment of exoffenders with a misdemeanor rose 5 percent following the Ban.

The contrast between interviewing and job finding shows that, while the Ban enables more exoffenders to receive interviews, exoffenders face a substantial barrier to employment.

There are at least two reasons for this. First, there is a lot of selection into jobs based on other components of the payoff, namely, η and ε. Indeed, in this calibration, the dispersion in z is fairly limited by comparison. One might have anticipated that exoffenders would be aided more by the Ban when their match value is not too low. However, when there is little dispersion in z, it follows that "hiding" z has little impact on the hiring decision, which is instead anchored by the variances of η and ε. Second, an important property of this calibration is that queues are relatively long: the firm receives on average 12 applications per vacancy. Hence, it can often correct the "mistake" of interviewing a low-surplus exoffender by choosing another applicant to interview.

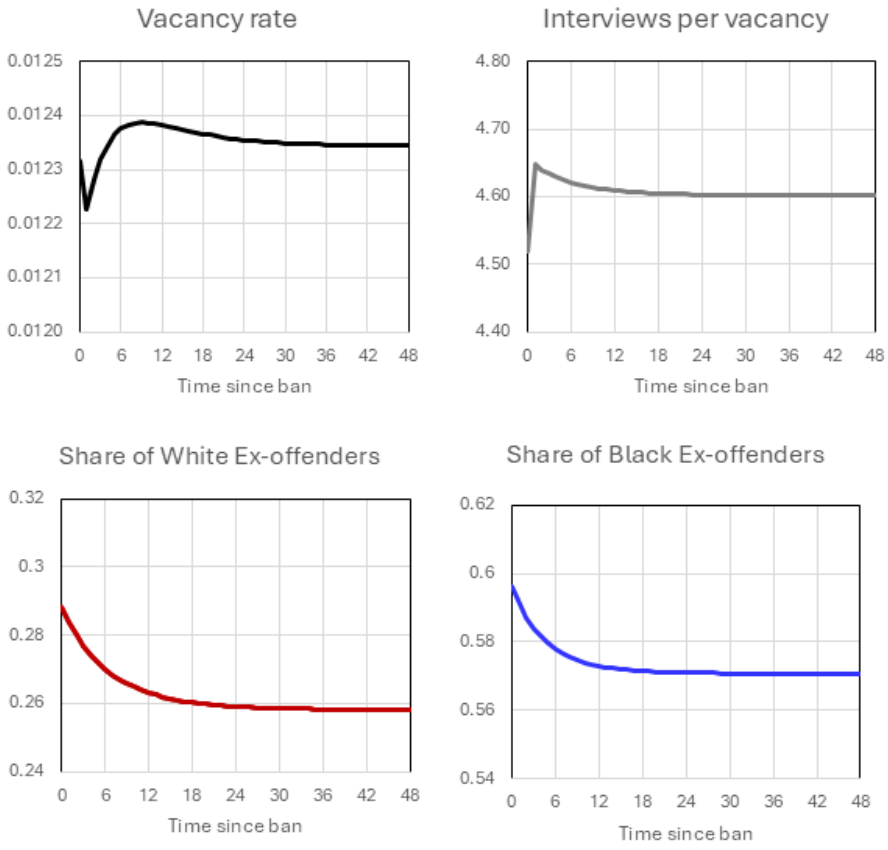
Considering the opposing paths of exoffender and nonoffender employment, it would seem that the response of *aggregate* employment is ambiguous. However, since separation rates are fixed, the driving force behind aggregate employment dynamics must be aggregate job creation. Thus, we can cut through some of the ambiguity surrounding aggregate employment by inspecting the behavior of the vacancy rate, as shown in Figure 2.

Interestingly, the Ban has very distinct short- and long-run effects on the vacancy rate. On impact, the vacancy rate declines. Intuitively, since firms must now interview to learn an applicant's criminal record, Ban the Box raises the cost of finding the best candidate. As a result, the Ban implies a lower expected payoff from filling a job and leads to an immediate drop in job creation. In the ensuing months, however, vacancies rebound. Indeed, the model makes a rather remarkable prediction: in the long run, the vacancy rate is *higher* than in the pre-Ban steady state.

This surprising long-run property of the model has a simple explanation. Despite the effort allocated to screening, firms will inevitably hire more exoffenders following the Ban. By shrinking the exoffender share of the nonemployed, these firms are elevating the average quality of future applicants. In other words, one firm's hiring of an exoffender creates a positive externality for all potential entrants. As a result, entry begins rebounding after its initial decline. A kind of virtuous feedback loop then develops: increased entry enables still more low-surplus exoffenders to find jobs, leading to further declines in the exoffender share. In the end, the exoffender share falls by 2.5-3 p.p. for each race, as shown in Figure 2.

The delayed rise in the vacancy rate shapes the dynamics of hiring and interviewing. Nonemployment within each racial group first rises, reflecting the initial dip in job creation, before falling below the pre-Ban steady state level. The long-run decline in nonemployment reflects the long-run increase in the vacancy rate in this calibration. This prediction is at odds with Doleac and Hansen (2020), who find that Ban the Box substantially increases aggregate black nonemployment. However, a re-evaluation of this finding using a longer sample period fails to find any robust statistically significant impact on nonemployment by

Figure 2: Effect of Ban the Box on vacancy rates, interviews, and unemployment composition



Note: The vacancy rate is measured as vacancies relative to population, e.g., in the pre-Ban steady state (period 0), vacancies were around 1.23 percent of the population. In the bottom panel, we plot the exoffender share of the nonemployed for each race. For example, in the pre-Ban steady state, just over 28 percent of the white nonemployed were exoffenders.

race (see Kaestner and Wang (2024)).

In addition, interviews per vacancy spike on impact of the policy. This occurs in part because the initial drop in vacancy posting lengthens queues, yielding more applicants to interview.[21] The interview rate later recedes somewhat because the rebound in vacancy posting unwinds the initial increase in queues.

These findings have a potentially important implication for empirical analysis. Consider implementing an event study design measuring changes in hiring and interviewing before and after the Ban takes effect. Setting too narrow a window around the implementation date would lead to inconsistent estimates. Specifically, one would underestimate the long-run positive effects on hiring, job creation, and worker interview rates.

We can gain further insight into the workings of the model by considering alternative calibrations. The results described thus far were conditioned on a certain value for the baseline dispersion, $\bar{\sigma}_\eta$. We noted that this value was on the lower end of the feasible range insofar as it implied a calibration where the expected payoffs of exoffenders and nonoffenders were relatively compressed. Specifically, under this calibration, the average payoff of exoffenders was 20 percent lower than for nonoffenders.

We are now in a position to examine how a higher value of $\bar{\sigma}_\eta$ shapes the model predictions. We raise $\bar{\sigma}_\eta$ such that the standard deviation of the present value of η among white nonoffenders is 0.8, e.g., $\bar{\sigma}_\eta \times (1+r)/(r+\delta(\mathcal{W},0)) = 0.8$. This compares to a value of 0.325 in the original calibration. As previewed, the model "needs" more dispersion in z in this case to match the moments. Specifically, this choice of $\bar{\sigma}_\eta$ implies a calibration where the average payoff of exoffenders is 50 percent lower than for nonoffenders.

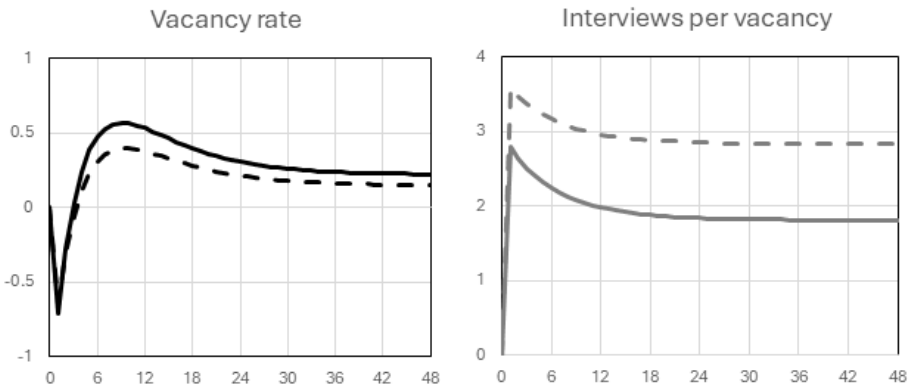
Comparing these results with what we found under the original calibration, a few key themes emerge. First and foremost, at higher values of $\bar{\sigma}_\eta$, the responses of job creation and, therefore, aggregate employment are attenuated (see Figure 3 and Figure 4). To see why, recall that even under the Ban, firms observe η. Thus, a higher value of $\bar{\sigma}_\eta$ implies that a larger share of the variance of the firm's payoff is accounted for by the observable component. Accordingly, "hiding" the true value of z has a smaller impact on the distribution of payoffs and, thus, on job creation decisions.

Perhaps surprisingly, though, interviewing activity is robust. Firms' interviewing intensity, as measured by interviews per vacancy, increases more given higher values of $\bar{\sigma}_\eta$. In addition, the interviewing probabilities by type look very similar to what we observed under the original calibration (see Figure 5). Why?

While the *direct* effect of a higher $\bar{\sigma}_\eta$ shows through clearly for most outcomes, the

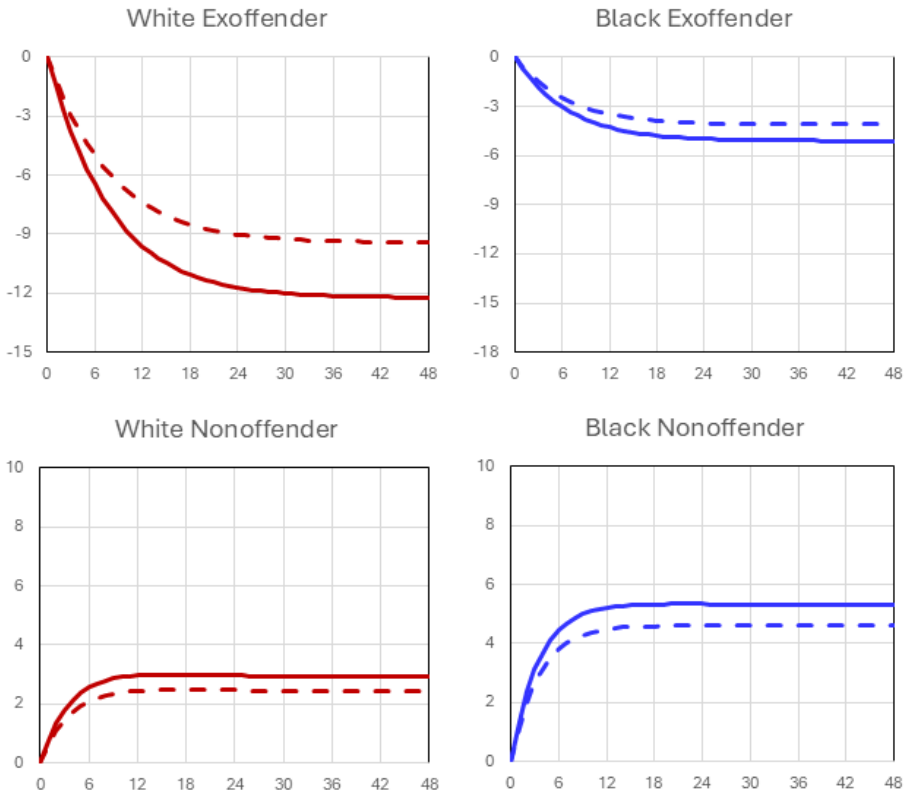
[21] The reservation values also adjust, a point that will come through more clearly when we compare these results to an alternative calibration below.

Figure 3: Effect of Ban the Box on vacancies and interviews under a higher value of $\bar{\sigma}_\eta$



Note: This plots the percent changes in vacancy and interview rates. The solid lines refer to the original calibration. The dashed lines are based on the calibration where $\bar{\sigma}_\eta = 0.8$ in present value terms.

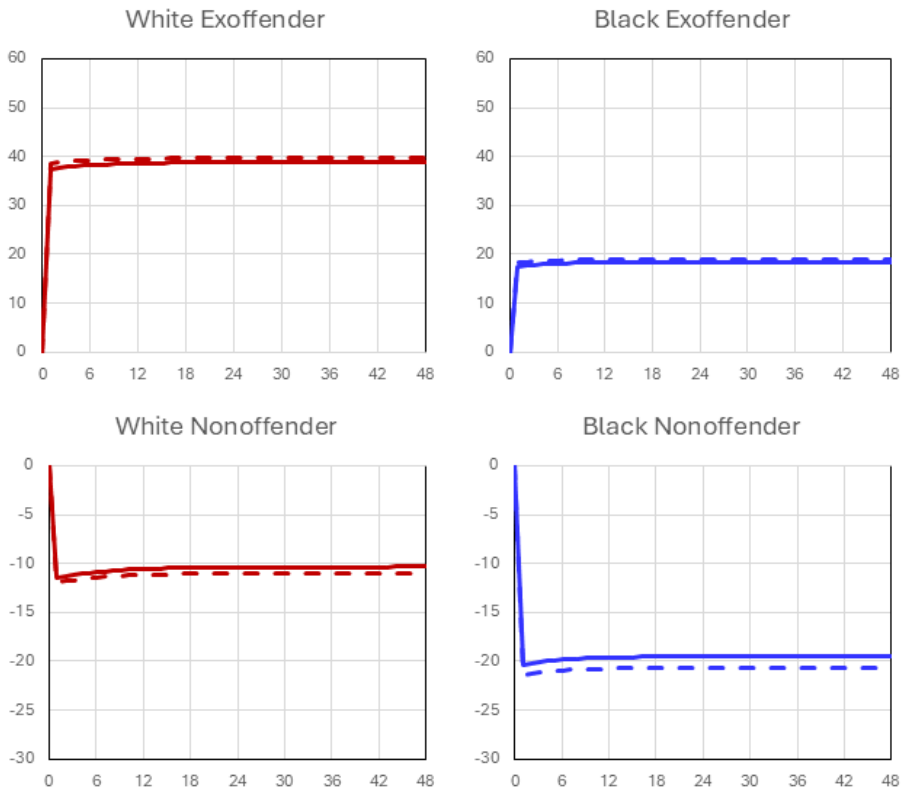
Figure 4: Effect of Ban the Box on nonemployment rates under a higher value of $\bar{\sigma}_\eta$



Note: This plots the percent changes in nonemployment rates by worker type. The solid lines refer to the original calibration. The dashed lines are based on the calibration where $\bar{\sigma}_\eta = 0.8$ in present value terms.

equilibrium is also shaped by the *indirect* effects of other parameter changes in this revised calibration. Most notably, the implied increase in the dispersion of z pushes in the opposite direction as higher dispersion of $\bar{\sigma}_\eta$. Thus, an increase in $\bar{\sigma}_\eta$ *all else equal* does attenuate the change in interviewing activity for essentially the same reason vacancy posting responds less: in this setting, selection into interviewing (and hiring) depends relatively heavily on η rather than z. However, higher dispersion of z all else equal implies a stronger incentive to interview under the Ban. Intuitively, the benefit of interviewing is increasing in the uncertainty over the final payoff. Moreover, exoffenders experience especially large increases in interviewing probabilities when z is more disperse. The reason is that the Ban virtually equalizes reservation values within race, which in turn leads to bigger increases in exoffender interviews when their z is especially low.

Figure 5: Effect of Ban the Box on interview rates under a higher value of $\bar{\sigma}_\eta$



Note: This plots the percent changes in interview rates by worker type. The solid lines refer to the original calibration. The dashed lines are based on the calibration where $\bar{\sigma}_\eta = 0.8$ in present value terms.

What we see in the figures is the net effect of increases in the dispersion in z and $\bar{\sigma}_\eta$. Our finding that interviewing activity increases somewhat suggests that this decision is especially sensitive to the increase in dispersion in z. However, the increase in exoffender interviewing

does not translate into that much hiring precisely because the interview reveals their true type. Rather, the hiring decision reflects the full distribution of (true) payoffs drawn during the interview. Accordingly, hiring outcomes more clearly reflect the change in the baseline dispersion, $\bar{\sigma}_\eta$.

Although we present just one alternative calibration here, we have confirmed that continuing to raise $\bar{\sigma}_\eta$ further attenuates job creation and employment outcomes, while having relatively little impact on interviewing probabilities. Thus, our original calibration appears to yield an upper bound on the size of labor market responses to the Ban.

Finally, we want to return briefly to another parameter that is especially hard to pin down, namely, the marginal cost of interviewing, c. Much of the available evidence on hiring costs does not distinguish between its different components. As a result, in our initial calibration, we very likely overstated the size of interviewing costs.

To develop a smaller alternative for c, we return to Barron et al. (1997). These authors highlight that the 1980 Equal Opportunity Pilot Project survey does appear to elicit estimates of time spent, specifically, on interviewing. Whereas other surveys covered by Barron et al asked firms about the total number of applicants for a position, the 1980 EOPP asked only for the number of interviews. As a result, it is plausible that, in responding to follow-up questions about the time spent in the hiring process, firms reported only on the time allocated to interviewing.

To facilitate comparison with our original estimate, we incorporate the 1980 EOPP results in the following way. First, we combine the 1980 EOPP data with estimates of *total* recruiting costs from surveys with similar samples, including the 1982 EOPP survey. Pooling these estimates, we can infer the interviewing cost *share* of total recruiting costs. We then multiple this by our earlier estimate of recruiting costs based in part on Villena Roldan (2012). This yields an estimate of interviewing costs of a hire equal to 1.3 percent of market value per worker. The model is recalibrated to match this target.

As previewed earlier, we find that reducing the marginal cost of interviewing leads to smaller employment responses to Ban the Box. Intuitively, a low cost of interviewing incentivizes firms to screen more intensively. This works to partially undo the effect of the Ban. Accordingly, exoffender nonemployment falls less, and nonoffender nonemployment rises less. Thus, we again find that our original calibration is likely to yield an upper bound on the labor market response to the Ban.

6 Conclusion

This paper introduces an equilibrium search and matching model with screening. The model formalizes screening as a decision problem over whom to interview, adapting the seminal work of Weitzman (1979). This problem is integrated within a macroeconomic model of job creation and unemployment. In addition, at the microeconomic level, the model incorporates multiple sources of heterogeneity and a flexible information structure, both of which are crucial for policy applications. Even with a rich microeconomic environment, the model retains analytical tractability: it admits closed form solutions for the *distributions* of interviewing and job finding rates over worker types in *and* out of steady state. Therefore, the model allows for a complete characterization of the dynamic equilibrium effects of policy interventions.

The model is applied to study the widely adopted Ban the Box policy. Within the model, the Ban is implemented by concealing an element of the worker's type (criminal record) from the firm during the application stage of the recruiting process. The firm can observe all the candidate's attributes only by paying the cost to interview.

We find that the policy has decidely mixed success. It does substantially lift interview rates for exoffenders. However, these gains come at the expense of nonoffenders, whose interview rates fall. Notably, among black applicants, the gains for exoffenders are attenuated, and the losses for nonoffenders are amplified. With a certain exception (discussed below), these predictions accord well with the empirical estimates. Even with a large increase in exoffender interviews, though, the model predicts a relatively modest increase in exoffender employment. These two results can coexist in the calibrated model because there are many applicants per vacancy. As a result, even if an exoffender is admitted for an interview, he often must compete against more productive nonoffenders. The modest employment gains predicted by the model are also consistent with the empirical evidence on the long-run impact of Ban the Box.

The quantitative application also serves to highlight the crucial interplay in the model between interviewing, job creation, and nonemployment. By introducing uncertainty over a worker's type, the Ban provides a strong motive to increase interviewing. This feeds back onto the job creation decision: by raising the cost of filling a position, the Ban induces less entry, and the market slackens. Importantly, the resulting rise in nonemployment includes many jobless nonoffenders, who are crowded out by exoffender applicants under the Ban. The larger nonoffender share of job searchers then entices more firms to (re)open vacancies, and the market tightens. Since each vacancy now receives fewer applications, the number of interviews per vacancy recedes.

We see several opportunities to further develop the model, each of which will sharpen the quantitative application to Ban the Box. First, in ongoing research, we are integrating wage setting into the model. This extension allows, realistically, that differentially productive workers—exoffenders and nonoffenders—are paid different wages. The introduction of a wage bargain along these lines will also enable us to leverage data on wages to learn about the dispersion of productivity and, thus, close the calibration of the model. Second, in modeling imperfect information over worker types, we are exploring departures from full Bayesian updating. One result from Agan and Starr stands out as being especially difficult to reconcile with the model, namely, the *increase* in interviewing rates among white nonoffenders. Agan and Starr suggest that this finding may point to some form of misperception or bias among firms.

References

Agan, A., Garin, A., Koustas, D., Mas, A., and Yang, C. (2024). The labor market impacts of reducing felony convictions. *American Economic Review: Insights*, 6(3):341–358.

Agan, A. and Starr, S. (2018). Ban the box, criminal records, and racial discrimination: A field experiment. *The Quarterly Journal of Economics*, 133(1):191–235.

Avery, B. and Lu, H. (2021). Ban the Box. Technical report, National Employment Law Project.

Barron, J., Berger, M., and Black, D. (1997). *On the Job Training*. Kalamazoo, MI: Upjohn Institute.

Cai, X., Gautier, P., and Wolthoff, R. (2017). Search frictions, competing mechanisms and optimal market segmentation. *Journal of Economic Theory*, 169:453–473.

Cortés, K. R., Glover, A., and Tasci, M. (2022). The unintended consequences of employer credit check bans for labor markets. *Review of Economics and Statistics*, 104(5):997–1009.

Davis, J., A1 Ouimet, P., and A1 Wang, X. (2022). Hidden performance: Salary history bans and the gender pay gap. *The Review of Corporate Finance Studies*, 11(3):511–553.

Diamond, P. A. (1971). A model of price adjustment. *Journal of Economic theory*, 3(2):156–168.

Doleac, J. L. and Hansen, B. (2020). The unintended consequences of "Ban the Box": Statistical discrimination and employment outcomes when criminal histories are hidden. *Journal of Labor Economics*, 38(2):321–374.

Fang, H. and Moro, A. (2011). Theories of statistical discrimination and affirmative action: A survey. In Benhabib, J., Jackson, M. O., and Bisin, A., editors, *Handbook of Social Economics Volume 1A*, pages 133–200. The Netherlands: North-Holland.

Jackson, O. and Zhao, B. (2017). The effect of changing employers' access to criminal histories on ex-offenders' labor market outcomes: Evidence from the 2010–2012 Massachusetts CORI reform. Working Paper 16-30, Federal Reserve Bank of Boston.

Jarosch, G. and Pilossoph, L. (2019). Statistical discrimination and duration dependence in the job finding rate. *The Review of Economic Studies*, 86(4):1631–1665.

Kaestner, R. and Wang, X. (2024). Ban-the-box laws: Fair and effective? Working paper 32273, NBER.

Lester, B., Visschers, L., and Wolthoff, R. (2015). Meeting technologies and optimal trading mechanisms in competitive search markets. *Journal of Economic Theory*, 155:1–15.

Lippman, S. A. and McCall, J. J. (1976). The economics of job search: A survey. *Economic inquiry*, 14(2):155–189.

Lofstrom, M., Martin, B., and Premkumar, D. (2023). Are younger generations committing less crime? Technical report, Public Policy Institute of California.

Lundquist, J. H., Pager, D., and Strader, E. (2018). Does a criminal past predict worker performance? Evidence from one of America's largest employers. *Social Forces*, 96(3):1039–1068.

Meier, M. and Obermeier, T. (2025). Employer screening and optimal unemployment insurance. *The Economic Journal*, 135(667):671–711.

Menzio, G. and Shi, S. (2011). Efficient search on the job and the business cycle. *Journal of Political Economy*, 119(3):468–510.

Minor, D., Persico, N., and Weiss, D. M. (2018). Criminal background and job performance. *IZA Journal of Labor Policy*, 7.

Pyle, B. (2024). A second look: Local labor markets and the impact of ban the box policies after criminal legal involvement. Mimeo, Boston University School of Law.

Rocque, M., Posick, C., and Hoyle, J. (2016). Age and crime. In Jennings, W. G., editor, *The Encyclopedia of Crime and Punishment*. John Wiley Sons.

Rose, E. K. (2021). Does banning the box help ex-offenders get jobs? Evaluating the effects of a prominent example. *Journal of Labor Economics*, 39(1):79–113.

Shannon, S. K., Uggen, C., Schnittker, J., Thompson, M., Wakefield, S., , and Massoglia, M. (2017). The growth, scope, and spatial distribution of people with felony records in the United States, 1948–2010. *Demography*, 54:1795–1818.

Shoag, D. and Veuger, S. (2021). Ban-the-Box measures help high-crime neighborhoods. *Journal of Law and Economics*, 64(1):85–105.

Villena Roldan, B. (2012). Aggregate implications of employer search and recruiting selection. Mimeo, Institute for Economic Policies, Universidad Andres Bello.

Weitzman, M. (1979). Optimal search for the best alternative. *Econometrica*, 47(3):641–654.

Wozniak, A. (2015). Discrimination and the effects of drug testing on black employment. *Review of Economics and Statistics*, 97(3):548–566.

Appendix A Proofs

A.1 Proof of Lemma 2

For an applicant of type $\mathbf{x}_k$ who draws η_k to be interviewed, two things need to occur.

First, any applicants with $y^\star > Y_k^\star$ need to have a productivity below $Y_k^\star$. The probability of this is

$$e^{-\sum_{j=1}^{k-1}\lambda_j\left(1-H_j\left(Y_k^\star\right)\right)} \tag{22}$$

by the invariance property of the Poisson distribution (Lester et al., 2015; Cai et al., 2017).

Second, the applicant needs to be among the interviewees within the set of applicants with the same $y^\star = Y_k^\star$. If there were an infinite supply of such applicants, then under the optimal stopping rule the number of interviews the firm would conduct within this pool would be geometric with success probability $1 - H_k(Y_k^\star)$, and thus have mean $1/\left(1 - H_k(Y_k^\star)\right)$.

With a finite pool, the expected number of interviews is smaller. The only overcount in the infinite-supply calculation arises when all applicants with $y^\star = Y_k^\star$ have realized productivity below $Y_k^\star$. Because the number of such applicants is Poisson(λ_k) and each applicant independently exceeds the threshold with probability $1 - H_k(Y_k^\star)$, that overcounting event occurs with probability $e^{-\lambda_k(1-H_k(Y_k^\star))}$. Conditional on that event, the expected number of interviews we counted in error (by assuming an infinite pool) is again $1/\left(1 - H_k(Y_k^\star)\right)$. Therefore, the expected number of interviews the firm actually conducts among applicants with applicants with $y^\star = Y_k^\star$ equals

$$\frac{1 - e^{-\lambda_k(1-H_k(Y_k^\star))}}{1 - H_k(Y_k^\star)}. \tag{23}$$

Multiplying (22) by (23) and dividing by the expected number of applicants λ_k gives the desired expression.

A.2 Proof of Lemma 3

A necessary condition for a firm to hire an applicant with $y^\star = Y_k^\star$ is that the firm interviews such applicants. Hence, as in the proof of Lemma 2, it needs to be the case that there are no applicants with $y^\star = Y_j^\star$ for $j \in \{1, \ldots, k-1\}$ that have a productivity greater than $Y_k^\star$. The probability of this event is given by (22).

Conditional on this event, there are several possibilities. First, with probability $1 - e^{-\lambda_k\left(1-H_k(Y_k^\star)\right)}$, the firm interviews an applicant with $y^\star = Y_k^\star$ that has a productivity strictly

above $Y_k^\star$. In that case, the firm stops interviewing and hires the applicant. This yields the first term in (7).

Second, the firm may interview all applicants with $y^\star = Y_k^\star$ but the highest productivity among them is below $Y_k^\star$. Denote this highest productivity by y' which then has density

$$e^{-\lambda_k(1-H_k(y'))}\lambda_k h_k(y') \tag{24}$$

To derive the probability that this applicant is hired, suppose that y' is in the interval $\left(Y_{\ell+1}^\star, Y_\ell^\star\right]$ for some $\ell \in \{k, \ldots, K\}$. In that case, the applicant is hired if and only if all applicants with $y^\star \in \{Y_1^\star, \ldots, Y_{k-1}^\star, Y_{k+1}^\star, \ldots, Y_l^\star\}$ have a lower productivity. The unconditional probability of this event is

$$e^{-\sum_{j=1}^{k-1}\lambda_j(1-H_j(y'))-\sum_{j=k+1}^{\ell}\lambda_j(1-H_j(y'))}. \tag{25}$$

Multiplying (24) by (25), integrating over y' and summing over ℓ gives the second term in (7).

A.3 Proof of Lemma 4

As in Lemma 3, the probability that the firm did not receive any applications from workers with $y^\star \in \left\{Y_1^\star, \ldots, Y_{k-1}^\star\right\}$ and a realization $y^i \geq Y_k^\star$, but that it *did* receive at least one application with $y^\star = Y_k^\star$ and a realization $y^i \geq Y_k^\star$ is equal to

$$e^{-\sum_{j=1}^{k-1}\lambda_j\left(1-H_j\left(Y_k^\star\right)\right)}\left(1-e^{-\lambda_k\left(1-H_k\left(Y_k^\star\right)\right)}\right).$$

Conditional on this event, the firm stops interviewing after identifying the first of such applicants. Hence, the conditional probability that this interviewee has a productivity greater than y is

$$\frac{1-H_k\left(\max\left\{y, Y_k^\star\right\}\right)}{1-H_k\left(Y_k^\star\right)}.$$

When $y < Y_k^\star$, one needs to additionally account for the possibility that the firm may end up hiring a worker with $y^\star = Y_k^\star$ and productivity realization $y^i \in \left[y, Y_k^\star\right)$. In that case, the firm will continue interviewing after exhausting all applicants $y^\star = Y_k^\star$, possibly up to applicants with $y^\star = Y_{k(y)}^\star$. The integral term represents this event in a way that is analogous to Lemma 3.