Testing for Racial Prejudice in the Parole Board Release Process: Theory and Evidence

Shamena Anwar and Hanming Fang

ABSTRACT

We develop a model of a parole board contemplating whether to grant parole release to a prisoner who has finished serving his minimum sentence. The model implies a simple outcome test for racial prejudice that is based on the released inmate's rate of recidivism and is robust to the inframarginality problem. Our model has several testable implications for which we show empirical support. Applying our test to data on all prison releases in Pennsylvania between 1999 and 2003, we find no evidence of racial prejudice.

1. INTRODUCTION

It has been widely documented that blacks compose a disproportionate share of the U.S. prison population. According to the U.S. Bureau of Justice Statistics, a total of 2,297,500 inmates were held in custody in state or federal prisons or in local jails as of June 30, 2009. Whites accounted for 34 percent of the incarcerated population, blacks 39 percent, and Hispanics 20 percent (West 2010, table 16). In contrast, the fractions of whites, blacks, and Hispanics in the U.S. population are 64, 12, and 16 percent, respectively (Humes, Jones, and Ramirez 2011).

It is an important policy concern to understand the causes of these

SHAMENA ANWAR is an Assistant Professor of Economics and Public Policy at Carnegie Mellon University. HANMING FANG is a Professor of Economics at University of Pennsylvania and a Fellow at the National Bureau of Economic Research. We would like to thank William Hubbard, Prasad Krishnamurthy, Nicolas Sahuguet, participants at the 2012 Conference on Empirical Legal Studies, and an anonymous referee for many helpful comments and suggestions. We are extremely grateful to Kristofer Bucklen at the Pennsylvania Department of Corrections for his assistance with the data. Fang gratefully acknowledges financial support from National Science Foundation grant SES-1122902. We are responsible for all remaining errors.

[*Journal of Legal Studies,* vol. 44 (January 2015)]

racial disparities in incarceration rates. Although these disparities can potentially be caused by racial differences in crime prevalence, a growing literature has investigated the extent to which racial discrimination at various stages of the criminal justice system is responsible. Studies that examine the role of prejudice in motor vehicle searches can reveal whether minorities are more likely to be caught for a given commission of crime (see Knowles, Persico, and Todd 2001; Anwar and Fang 2006; Antonovics and Knight 2009). Rehavi and Starr (2014) find evidence that, conditional on being arrested for the same crime, prosecutors use their discretion to charge black defendants with more severe crimes. Anwar, Bayer, and Hjalmarsson (2012) show that the racial makeup of the jury can have a large effect on the black-white conviction ratio. Other studies have looked at potential discrimination at the sentencing stage. Abrams, Bertrand, and Mullainathan (2012) estimate whether judges differ from each other in how they sentence minorities. In Pennsylvania, Steffensmeier and Demuth (2001) find that Hispanic defendants receive the harshest sentences, while Muhlhausen (2004) finds that black judges sentence black offenders to longer prison terms than white judges give to white offenders. Ayres and Waldfogel (1994) and Bushway and Gelbach (2010) study the role of prejudice in bail setting, while Alesina and La Ferrara (2014) study the role of racial prejudice in death penalty sentences.

This paper studies whether discrimination plays a role in the last stage of the criminal justice system—the prison release process. Sixteen U.S. states still allow parole boards to have complete discretion over the release of prisoners, subject to the constraints of the prisoner's prescribed minimum and maximum sentences.[1] Given that parole boards have complete authority over how much of the prescribed sentence range a prisoner will serve, they are in the position to either remedy or exacerbate the biases that may be present in earlier stages of the criminal justice system. In this paper we examine whether the parole boards' release decisions reflect racial prejudice using data on all Pennsylvania prison releases between 1999 and 2003.

Many previous studies have examined this issue using action-based tests, which essentially compare whether minorities serve a greater proportion of their sentence or are less likely to be paroled at a given point

1. Over the past 30 years many states have transitioned from the parole release system to truth-in-sentencing schemes, which require prisoners to serve a fixed proportion of their sentence (Kuziemko 2013).

in time than their white counterparts.[2] However, it is well known that simple racial disparities in these action-based measures, even after controlling for observable characteristics of inmates, are not necessarily evidence that the parole board is racially prejudiced. These disparities may result from an omitted-variables problem, which occurs when there are systematic differences across races in the inmates' characteristics that are observable to and used by the parole board in its release decision but that are unobserved by researchers. These disparities can also arise from statistical discrimination, which occurs when there is crucial information that is unobservable to the parole board but that is correlated with inmate race. To deal with these issues, we use an outcome-based test because, if applied properly, such tests can identify racial prejudice even in the presence of omitted variables and statistical discrimination, as racial prejudice will have a different impact on the outcome in question.

The outcome test we use is based on a simple model of the parole board's release decisions.[3] We consider a parole board that is contemplating whether to grant parole release to a prisoner who has just finished serving his minimum sentence and is thus eligible for parole. The parole board faces a trade-off. On one hand, releasing the prisoner on parole saves the imprisonment cost; on the other hand, it imposes a social cost if the prisoner has not been rehabilitated and commits crimes after release. We show that the parole board will choose to grant the prisoner parole if and only if his perceived rate of recidivism is at or below a certain threshold, where the rate is defined as the product of the perceived probability that the inmate is not rehabilitated and the rate of recidivism for nonrehabilitated inmates.[4] The parole board will use a lower threshold for minorities if it is prejudiced against minorities. If the inmate's perceived rate of recidivism at the completion of his minimum sentence is too high, the parole board will keep him incarcerated; each successive time period the inmate completes with good behavior increases the parole board's perception that he is rehabilitated and thus lowers his perceived rate of recidivism. The moment the inmate's perceived recidivism rate is lowered enough to meet the parole board's threshold, the inmate will be released. If the inmate's perceived rate never falls enough to meet the threshold, he will be released only on the completion of his maximum sentence. Impor-

2. The results of these studies are discussed in Section 2.

3. Our model is related to Bernhardt, Mongrain, and Roberts (2010), although their goal is to show the efficiency of the parole board release system rather than test for prejudice.

4. We assume that rehabilitated inmates do not recidivate.

tantly, this implies that every prisoner granted parole release between his minimum and maximum sentences has an assessed recidivism rate exactly equal to the aforementioned race-specific threshold.

To implement our outcome test for racial prejudice, we need to identify the release thresholds being used for each race and compare them. As is well known, the main difficulty that arises when implementing outcome tests is the inframarginality problem, which refers to the difference between the comparisons of the average and marginal outcomes across racial or gender groups.[5] In order to identify the threshold being used, we need to identify the recidivism rate for the marginal person who is released (that is, the person released whose rate is exactly at the threshold). Generally, however, without having access to all of the information the parole board has, the marginal person cannot be identified. This typically results in only the average recidivism rate being identified. We deal with the inframarginality problem in this paper by noting that in our model, every prisoner released by the parole board between his minimum and maximum sentence, regardless of his characteristics, has a recidivism rate exactly equal to the threshold set by the parole board. Therefore, within this subsample, the marginal prisoner released is the same as the average prisoner released, and thus our application of the outcome test is not subject to the inframarginality critique.[6]

Although this parole release setting is optimal for dealing with the inframarginality problem, the downside is that the outcome measure we use is not completely objective, as whether an individual is convicted of a new crime depends on the behavior of other agents that might harbor prejudice. While this potential downstream racial prejudice should not impact our ability to uncover racial prejudice on the part of the parole board, as we will discuss, it can result in a situation in which the parole board statistically discriminates simply because of downstream racial bias.

Our test for racial prejudice is based on our model of parole board behavior. As such, evidence for or against racial prejudice using our test is only as credible as our proposed model. Fortunately, our model has three auxiliary implications that can be tested using our data set. We find

5. See Knowles, Persico, and Todd (2001) and Anwar and Fang (2006) for descriptions of this problem, and see Persico (2009) for a comprehensive review of the recent racial profiling literature.

6. Similar ideas about dealing with the inframarginality problem in the outcome test for racial prejudice can be found in Ayres and Waldfogel (1994) and Anwar and Fang (2012).

supportive evidence for all three of these predictions. Applying our test to the data, we find no evidence that racial prejudice plays a role in Pennsylvania's parole board release process.

The remainder of the paper is structured as follows. In Section 2 we review the related literature and discuss how our test fits in with previous approaches. Section 3 describes the sentencing and parole system in Pennsylvania. In Section 4 we present a model of how the parole board makes parole release decisions and, on the basis of the implications of the model, derive an estimation equation that indicates whether racial prejudice plays a role in the parole board's decision making. Section 5 describes our data set and presents the descriptive statistics, while Section 6 presents our main empirical results regarding the role of prejudice in parole board decisions and additional evidence supportive of the auxiliary predictions of the model. Finally, Section 7 concludes.

2. RELATED LITERATURE

Research that specifically examines racial prejudice in the parole release process has been rather scarce. The majority of the previous literature falls into two main categories, both of which use what we consider to be action-based tests (since they are based on the actions of the parole board). The first type of study essentially examines whether minorities serve a greater proportion of their sentence before being paroled than do their white counterparts (Morgan and Smith 2008). The findings are mixed. Petersilia (1985) found that minorities in Texas served a higher proportion of their sentences relative to whites, but the reverse was true in Michigan. More recently, Huebner and Bynum (2006) found that race had no effect on sentence served among a sample of men incarcerated for sexual offenses.

The second strand of literature uses data from parole decisions and explicitly examines whether race has an effect on parole being granted. Carroll and Mondrick (1976) examined the cases of 243 prisoners who appeared before a parole board between 1970 and 1971 and found that race had no impact on the decision to grant parole. In a more recent study, Morgan and Smith (2008) also found that race had no effect on parole release decisions using a sample of 762 inmates in Alabama who were eligible for parole between 1993 and 1994.

While none of the studies cited above can control for the personal interactions between the parole board and the individual, the majority do control for other key factors the parole board takes into account when it makes its release decisions. These include severity and type of offense, length of sentence, institutional misconduct and program participation, and various indices of the likely risk of recidivism on release. As such, these studies are effective in determining whether minority and white individuals with similar characteristics are treated the same in the parole release process.

There are, however, a few downsides to the action-based tests that were implemented in the above studies. First, to the extent that the parole board gleans relevant information at the parole hearing that varies systematically by inmates' race, the results from these studies can suffer from an omitted-variables bias. Second, these studies require an extensive set of controls that are often difficult for researchers to obtain; this is reflected in the limited amount of previous work in this area and the relatively small sample sizes that were used. Finally, while these tests reveal all forms of discriminatory behavior if fine enough controls are used, they do not allow us to determine which form of discrimination (namely, racial prejudice or statistical discrimination) is causing the disparity. While both types of discrimination are illegal, to eliminate disparities it is useful to know why they arise in the first place.

The outcome test we develop in this paper is in response to the shortcomings of these action-based tests. As detailed in Section 4, our test requires us to observe only limited information about individuals that is readily available in most data sets kept by a state's department of corrections. As such, it is an easier test to implement. Further, ours is a test specifically designed to pick up racial prejudice and can thus help reveal why racial differences are occurring. Note, however, that because our test is for racial prejudice, it makes no statement as to whether statistical discrimination is occurring. In this sense, one can think of outcome-based and action-based tests as complementary.

Our paper is most closely related to Mechoulan and Sahuguet (2015), which also uses an outcome test to test for racial prejudice in parole release decisions.[7] Our paper differs from that of Mechoulan and Sahuguet in both modeling and data, which leads us to conduct different empirical tests. Neither paper, however, finds evidence of racial prejudice against blacks in the parole release process. In the appendix we detail the differences between the approaches and show empirical evidence that, in our data set, only our model is supported.

7. Our paper and Mechoulan and Sahuguet (2015) were developed simultaneously.

3. CRIMINAL SENTENCING AND PAROLE RELEASE IN PENNSYLVANIA

All individuals in Pennsylvania convicted of a crime are sentenced by a judge who determines their minimum and maximum sentences.[8,9] Offenders with a maximum sentence of less than 2 years are sent to jail. The sentencing judge has discretion over whether to send offenders with a maximum sentence between 2 and 5 years to jail or prison.[10] Those with a maximum sentence greater than 5 years are automatically sent to prison. Individuals sent to prison must serve at least their minimum sentence. Once they have completed it, Pennsylvania's parole board, which consists of nine members appointed by the governor, has complete discretion over when to release them, until they reach their maximum sentence.[11]

Approximately 4 months before the inmate completes his minimum sentence, board members and hearing examiners review his file. The board uses this information to fill out the Parole Decisional Instrument form, which serves as a guideline for release. The instrument takes into account the type of conviction offense (nonviolent or violent), the level of risk (low, medium, or high) of returning to prison for a new offense or violation according to the Level of Service Inventory–Revised,[12] institutional programming completion, and institutional behavior. The inmate receives scores for each of these four critical dimensions, which are summed to calculate an overall score.[13] Scores from 2 to 6 "suggest parole," while scores of 7 or greater "suggest parole refusal" (Goldkamp et al. 2010, p. 18).

The board is not bound by these guidelines when casting its vote, however, and can take into account other factors such as the recommendations of the sentencing judge, prosecuting attorney, and warden and the board members' general impression of the inmate during the parole interview. The decision makers for each case depend on the type of offense committed. For nonviolent offenses, a hearing examiner and one board member

8. Unless noted otherwise, the information regarding the parole release process described in this section is from Pennsylvania Board of Probation and Parole, Working toward Safer Communities (http://www.pbpp.pa.gov/Pages/default.aspx#.VQIILGNTdD0).

9. The judge is aided in his or her decision by the sentencing guidelines, which are presented in the form of a grid containing a range of suggested minimum sentences, where the offender's offense gravity score of his current offense is on one axis and his prior-record score (measuring his prior criminal activity) is on the other axis. Judges are not required to conform to these guidelines.

10. The sentencing judge has discretion over when individuals sent to jail are released.

11. Once this is reached, they max out and must be released.

12. The Level of Service Inventory–Revised (LSI–R) is a quantitative survey of an offender's attributes and situation and is designed to help predict recidivism.

13. For example, an offender receives 3 points for a sentence for a violent offense, 3 points for unacceptable program compliance, 3 points if the LSI–R considers him to be high risk, and 5 points for a record of serious misconduct in prison (Goldkamp et al. 2010).

vote on the case. For violent offenses (except sex crimes and murder), two board members vote on the case. The inmate must receive two affirmative votes to be granted parole. For murder and sex offenses, the full board reviews the case, and the majority of the board must approve the inmate's release. Approximately 70 percent of the final case decisions follow the recommendation of the Parole Decisional Instrument.

Individuals who are not granted parole are given a list of requirements to be fulfilled by the time of the next parole review, which is usually within 6 months to a year. Individuals who are granted parole are released and monitored by parole officers. They can be returned to prison if they have a technical violation of their parole requirements or if they commit a new crime.[14]

4. A MODEL OF THE PAROLE BOARD'S BEHAVIOR

In this section we propose a simple continuous-time learning model of the parole board that is adapted from the model developed in Bernhardt, Mongrain, and Roberts (2010). We derive several implications and use these to test whether the parole board exhibits racial prejudice in its release decisions.

4.1. The Model

We model the parole board's behavior from the first moment inmate i becomes eligible for parole release, which occurs after he has served his minimum sentence $\underline{T}^i$. At that time the parole board observes information that it uses in its parole release decision. Some of the information is available to researchers, while some of it is not. For example, information regarding the inmate's conviction (type of crime committed and the sentencing terms) and his basic demographics (gender, race, and age) are observed in our data set; however, we do not observe any information that is likely contained in an inmate's prison dossier, including his behavior and incidents while in prison and his general demeanor. We denote the information available to the parole board about inmate i at time $\underline{T}^i$ by $(r, c_{\underline{T}^i})$, where r is the race of the inmate and $c_{\underline{T}^i}$ includes all other information. For simplicity, we assume that the race of a prisoner is either white, denoted W, or minority, denoted M; that is, $r \in \{W, M\}$.

14. Common reasons for technical parole violations include failing to report to a parole officer, carrying a weapon, traveling too far from home, not maintaining employment, and failing drug and alcohol tests (Petersilia 2003). Parolees receive an average of five violations before being returned to prison.

4.1.1. Rehabilitated or Nonrehabilitated. We assume that once the inmate completes his minimum sentence he is either rehabilitated or nonrehabilitated and that his type does not change from that point on. In our model, there are two major differences between a rehabilitated and a nonrehabilitated inmate. First, a rehabilitated inmate will not recidivate, while a race-r nonrehabilitated inmate recidivates at Poisson arrival rate $g_r > 0$ if granted parole release.[15] Second, when imprisoned, a race-r nonrehabilitated inmate is involved in prison incidents at a Poisson arrival rate $\lambda_r > 0$; however, rehabilitated inmates are not involved in prison incidents. Prison incidents can be thought of as any event in prison that negatively affects the probability of parole, such as misconduct with other prisoners or guards or not completing required programming. Note that we allow both the recidivism rate for nonrehabilitated parolees and the incident arrival rate for nonrehabilitated inmates to depend on their races. The former is especially important, as minorities and whites are likely to be released on parole into very different environments, which can have an effect on their future criminal behavior.

4.1.2. The Parole Board's Payoffs, Belief Evolutions, and Release Decisions. At any time after inmate i has served his minimum sentence $\underline{T}^i$ and before his maximum sentence $\overline{T}^i$, the parole board needs to decide whether to keep him imprisoned or grant parole release. Suppose that the cost of holding a prisoner for a particular time period is B, regardless of the race of the prisoner and whether he is rehabilitated.[16] The cost of releasing a nonrehabilitated inmate of race r for a particular time period is g_rC, where g_r is the rate at which a nonrehabilitated race-r inmate will recidivate during that time period and C is the cost to the parole board that results from the inmate recidivating.

The parole board can also obtain a psychological benefit $D_r \geq 0$ from keeping a race-r prisoner imprisoned. If the parole board is prejudiced against a particular race of inmates, it is likely to feel a higher psychological benefit from keeping a prisoner imprisoned.[17] This idea is summarized in the following definition:

15. In Section 4.4.1 we discuss how recidivism is measured.

16. In Section 6 this time period corresponds to 1 month.

17. If the parole board is prejudiced against minority inmates, in effect $D_W = 0$ and $D_M > 0$, so the cost of imprisoning minorities for an extra period is less than that for whites because of this extra psychological benefit. This is similar to the racial profiling literature, whereby prejudiced police officers enjoy searching minorities more and thus the cost to search them is less (see Knowles, Persico, and Todd 2001; Anwar and Fang 2006).

Table 1. The Parole Board's Flow Costs from Race-r Inmates

	Nonrehabilitated	Rehabilitated
In	$B - D_r$	$B - D_r$
Out	$g_r C$	0

Definition 1. We say that the parole board is prejudiced against race-r inmates if $D_r > D_{r'}$ for $r \neq r'$.

The costs associated with releasing a race-r prisoner for a particular time period is summarized in Table 1. We assume that $0 < B - D_r < g_r C$. These parameter restrictions imply that, if the parole board knows for certain that an inmate is rehabilitated, it prefers that he be released; on the other hand, if the parole board knows for certain that an inmate is nonrehabilitated, it prefers that he remain imprisoned.

Because a parolee will return to prison if he recidivates, the parole board's decision of whether to release a prisoner on parole is simply to compare the cost of keeping him incarcerated for the next period with the cost of releasing him for the next period.[18] For the purpose of this comparison, we denote by π_t^i the parole board's belief at time t that inmate i is of a rehabilitated type for some $t \geq \underline{T}^i$. The board will release the prisoner on parole if and only if the cost of releasing him is lower than the cost of keeping him incarcerated at time t:

$$(1 - \pi_t^i) \times g_r C \leq B - D_r \Leftrightarrow \pi_t^i \geq 1 - \frac{B - D_r}{g_r C} \equiv \pi_r^*. \tag{1}$$

Thus, the parole board will grant parole release to inmate i only if it is sufficiently confident that i has been rehabilitated. Importantly, the threshold π_r^* defined in equation (1) is increasing in D_r. This means that if the parole board is prejudiced against race-r inmates, it needs to be more certain (probabilistically) that race-r inmates are rehabilitated before granting parole release. Intuitively, if the parole board is prejudiced against race-r inmates, then the cost of incarcerating race-r inmates is lower, and consequently at the indifference point the cost of releasing

18. If the parolee does not recidivate at time t, he will remain on parole for time $t' = t + \varepsilon$, where $\varepsilon > 0$ is small. When we discuss the evolution of the parole board's beliefs later in this section, it will be evident that if it is profitable for the board to release the inmate at time t and the inmate does not recidivate at time t, it will be even more profitable for the board to release the inmate at time t'. Thus, when deciding to release an inmate, comparing the costs and benefits at time t is all that matters.

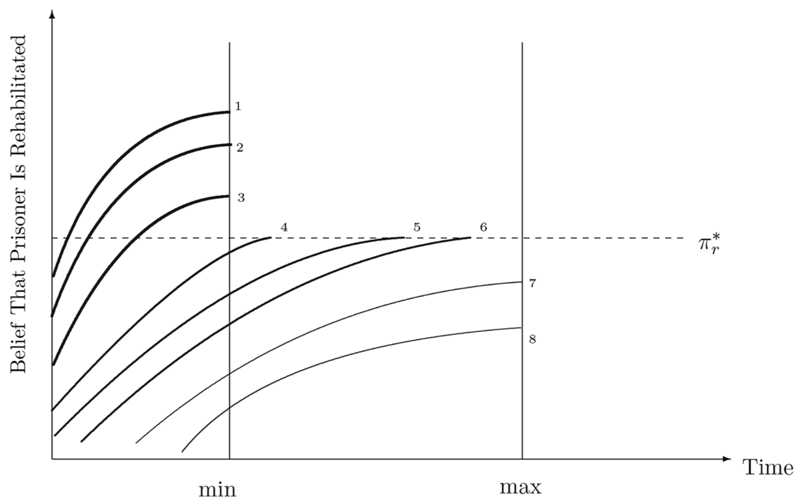

Figure 1. Parole board's release decisions for eight inmates

them must also be lower. We summarize the above discussion by the following proposition:

Proposition 1. The parole board will grant parole release to a race-r inmate at the first point in time between $\underline{T}^i$ and $\overline{T}^i$ when its belief about the inmate being rehabilitated exceeds π_r^*.

The parole board cannot perfectly know if the prisoner is rehabilitated at the time of the parole decision. Instead it forms beliefs based on available information, beginning at the time period when prisoner i has just completed his minimum sentence $\underline{T}^i$. We denote $\pi^i_{\underline{T}} \equiv \pi(r^i, c^i_{\underline{T}})$ as the parole board's belief given information $(r^i, c^i_{\underline{T}})$ that the prisoner has been rehabilitated at time $\underline{T}^i$. Figure 1 shows the relationship between the evolution of the parole board's beliefs and its release decisions for eight race-r inmates who have no incidents while in prison but enter the prison with different characteristics.[19] The vertical axis measures the evolution over time of the parole board's beliefs that these prisoners are rehabilitated. In particular, for prisoners 1–3, on completion of their minimum sentence $\underline{T}$, the parole board's belief that they are rehabilitated, $\pi^i_{\underline{T}}$, already exceeds the threshold π_r^*, and thus the parole board will release them immediately.

19. To keep Figure 1 simple, all eight prisoners have the same sentence, but this assumption has no impact on the results presented in this section.

For prisoners 4–8, however, $\pi^i_{\underline{T}} < \pi^*_r$, and thus these inmates will not be released immediately at $\underline{T}^i$. Their time of release will then depend on the evolution of the parole board's belief of π^i. Recall that the parole board will use all available information about prisoner *i* at time *t* to form its belief π^i_t. Although most of this information is static (such as prisoners' demographics and crime committed), the one component that will change over time is whether or not they are involved in prison incidents. We now derive the differential equation that governs how π^i_t changes over time.

Consider a small interval of time Δ between t and $t + \Delta$. Because we assume that prison incidents occur according to a Poisson process, we know that when Δ is small, there are two possible outcomes between t and $t + \Delta$: the first outcome is that an incident occurs in this time interval and the second outcome is that no incident occurs. If an incident occurs, then the belief of the parole board will immediately decrease to 0 and remain there through $\overline{T}$ because only nonrehabilitated inmates will be involved in an incident. If no incident occurs, then the parole board will update its beliefs using Bayes's rule. Noting that a race-*r* nonrehabilitated inmate will have no incidents during time interval Δ with probability $e^{-\lambda_{r_i}\Delta}$, we have

$$\pi^i_{t+\Delta} = \frac{\pi^i_t}{\pi^i_t + (1 - \pi^i_t)e^{-\lambda_{r_i}\Delta}}.$$

Thus, the evolution of the parole board's posterior belief if no incidents have occurred through $t + \Delta$ is governed by the following differential equation:

$$\dot{\pi}^i_t = \lim_{\Delta\to 0}\frac{\pi^i_{t+\Delta} - \pi^i_t}{\Delta} = \lim_{\Delta\to 0}\frac{\pi^i_t(1 - \pi^i_t)}{\pi^i_t + (1 - \pi^i_t)e^{-\lambda_{r_i}\Delta}}\,\frac{(1 - e^{-\lambda_{r_i}\Delta})}{\Delta} = \lambda_{r_i}\pi^i_t(1 - \pi^i_t). \quad (2)$$

Note that $\dot{\pi}^i_t > 0$, which means that each time period in which prisoner *i* does not have an incident increases the parole board's probability assessment that he is rehabilitated. This corresponds to the evolution of beliefs for prisoners 4–8 being drawn as upward sloping.

If inmate *i* has not been involved in any incident from $\underline{T}$ through time *t*, then we solve the differential equation (2) to find an expression for π^i_t:

$$\pi^i_t = \frac{1}{1 + \left[\left(1 - \pi^i_{\underline{T}}\right) / \pi^i_{\underline{T}}\right]e^{-\lambda_{r_i}t}}. \quad (3)$$

Proposition 2: Parole Board's Belief Evolution. If the parole board's initial belief that inmate *i* is of the rehabilitated type is $\pi^i_{\underline{T}}$, and the inmate is not involved in any incident from time $\underline{T}$ to time *t*, then the parole board's posterior belief at *t* is given by equation (3).

As stated in proposition 1, the parole board will want to release prisoner i the moment π_t^i reaches π_r^*. We can find this optimal release time, denoted t_i^*, by equating π_t^i with π_r^* and solving for t:

$$t_i^*(\pi_{\underline{T}}^i) = \frac{\ln\left[\left(1-\pi_{\underline{T}}^i\right) / \pi_{\underline{T}}^i\right] - \ln\left[\left(1-\pi_{r_i}^*\right) / \pi_{r_i}^*\right]}{\lambda_{r_i}}. \quad (4)$$

If this point in time occurs after the prisoner's maximum sentence $\overline{T}$, the parole board will be constrained to release him on completion of $\overline{T}$. This is the case for prisoners 7 and 8 in Figure 1. If this occurs between the prisoner's minimum and maximum sentences, as is the case for prisoners 4–6, the parole board will release him exactly at $t_i^*\left(\pi_{\underline{T}}^i\right)$. The following proposition summarizes the parole board's release decisions:

Proposition 3: Characterization of the Release Time. Let the parole board's initial belief about race-r inmate i being of rehabilitated type be $\pi_{\underline{T}}^i$. Assuming that inmate i has no incidents in prison after $\underline{T}$, the parole board's release schedule is as follows:

i) if $\pi_{\underline{T}}^i \geq \pi_r^*$, inmate i is released at $\underline{T}$;
ii) if $\pi_{\underline{T}}^i < \pi_r^*$ and $\underline{T} < t_i^*\left(\pi_{\underline{T}}^i\right) < \overline{T}$, inmate i is released at $t_i^*\left(\pi_{\underline{T}}^i\right)$; and
iii) if $\pi_{\underline{T}}^i < \pi_r^*$ and $t_i^*\left(\pi_{\underline{T}}^i\right) \geq \overline{T}$, inmate i is released at $\overline{T}$.

An important implication of the model is that all race-r prisoners released between $\underline{T}$ and $\overline{T}$ will have a probability of being rehabilitated that is exactly π_r^*. As is evident from Figure 1, this is not the case for those released at either $\underline{T}$ or $\overline{T}$. Among race-r prisoners released at the completion of their minimum sentence, there is a substantial amount of heterogeneity in their probability of being rehabilitated at the time of their release. The same is true among race-r prisoners released at the completion of their maximum sentence.

For simplicity we have assumed that rehabilitated inmates are not involved in prison incidents. However, in Section A2 we show that all of the key model implications continue to hold even if we allow rehabilitated inmates to be involved in prison incidents, as long as they are involved in them at a rate lower than that of their nonrehabilitated counterparts.

4.1.3. Reasons for Racial Differences in Sentence Served. The framework of the model also allows us to see the various reasons that inmates of different races serve different proportions of their sentences. The first case we consider is a racially prejudiced parole board. For ease of exposition, assume for now that λ and g are the same across races. As discussed earlier, if the parole board is racially prejudiced against minority inmates, it will require them to have a higher probability of being rehabilitated than white inmates (that is, $\pi^*_{\mathrm{M}} > \pi^*_{\mathrm{W}}$). Figure 2*A* shows the effects of this on two inmates, one white and one minority, who have exactly the same characteristics (and are thus represented by the same π^i_t curve). As is evident from the figure, $t^*_{\mathrm{W}} < t^*_{\mathrm{M}}$, and thus the minority inmate will be forced to serve more of his sentence than the identical white inmate.

Disparities in time served can also arise from statistical discrimination, which occurs when there is crucial information that is unobservable to the parole board and is correlated with inmate race. It will be efficient for the parole board to proxy for this unobservable information by taking an inmate's race into account. If, on average, minorities are known to rate worse with respect to this unobservable information, statistical discrimination will result in the parole board assuming a lower initial probability of a minority inmate being rehabilitated than an observationally equivalent white inmate. As shown in Figure 2*B*, the parole board will then require more incident-free time of minorities before they reach the release threshold.

If the parameters λ and g differ across races, it will be efficient for the parole board to take these differences into account, which will again lead to observationally equivalent individuals serving different amounts of their sentences. Because the parole board uses this information for efficiency (and not racial prejudice), this is another manifestation of statistical discrimination.

Finally, an omitted-variables problem will also result in racial differences in time served. This is different than racial prejudice or statistical discrimination, because in both of those cases, researchers have access to the same information as the parole board but will still find racial differences in time served using a regression framework. With an omitted-variables problem, we can find racial differences in time served simply because we cannot control for all of the factors that the parole board takes into account in making release decisions.

The above analysis makes clear that a valid test for racial prejudice cannot rely on time served, because racial prejudice, statistical discrimination, and an omitted-variables problem all have similar implications for time served. In the next section we will develop a test that does not have these problems.

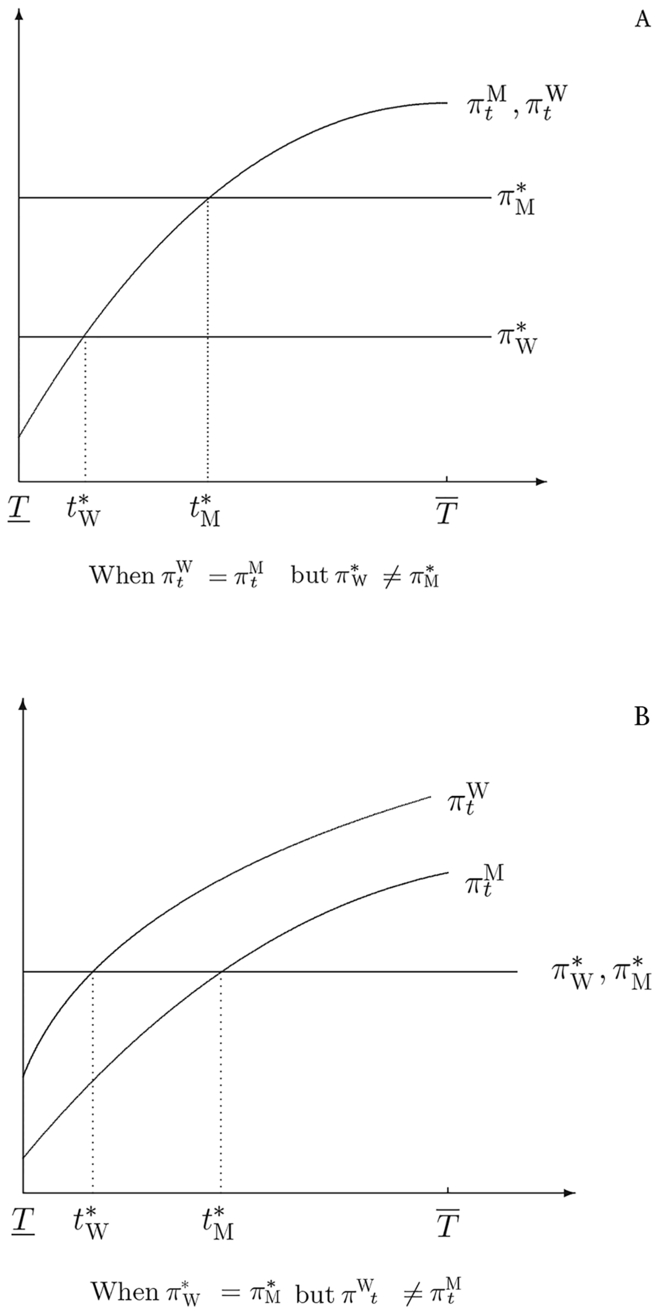

Figure 2. Differences in sentence served due to a racially prejudiced parole board (*A*) and statistical discrimination (*B*).

4.2. Test for Racial Prejudice

Given the difficulty of using the racial disparities in parole release time as indications of racial prejudice, we instead use an outcome test. This test requires an outcome that can be identified with available data and where racial prejudice and statistical discrimination have different impacts. If we restrict ourselves to examining inmates who are released between their minimum and maximum sentences, the outcome that satisfies both of these requirements is the inmate's expected rate of recidivism on parole, which is defined as $(1-\pi_r^*)g_r$. Recall that this rate is the same for everyone within a race because the parole board strategically releases every race-r inmate at the time his probability of being rehabilitated is exactly π_r^*.

Inserting the expression of π_r^* from equation (1), we have

$$(1-\pi_r^*)g_r = \frac{B-D_r}{C}. \tag{5}$$

Thus, the expected rate of recidivism differs across races only if D differs across races, which occurs if racial prejudice is present. Intuitively, with statistical discrimination, the parole board will take race into account when making release decisions, but it does so in a way that the expected rate of recidivism across all races is the same. When the parole board is racially prejudiced against race-r inmates, it will require them to serve longer than is optimal, and they will thus recidivate at a lower expected rate.

Note that the probability that parolees are rehabilitated, π_r^*, does not satisfy the outcome test requirement because both racial prejudice and statistical discrimination can lead to π_r^* differing by race. From equation (5) we see that the optimal π_r^* depends not just on prejudice but on g_r as well. Even if there is no racial prejudice, race-r members would be required to have a higher probability of being rehabilitated if their nonrehabilitated members have a higher rate of recidivating on release. While the fact that our test relies on comparing recidivism rates rather than rehabilitation rates is a direct implication of the model, this result is in line with the likely incentives of the parole board. From the parole board's perspective, two individuals who are released with the same rehabilitation probability but with different rates g_r impact it in different ways. Even though both individuals have the same eventual likelihood of recidivating, the one with the higher g_r (and thus the higher recidivism rate) will likely recidivate much sooner. While the recidivism cost of both individuals is the same, the parole board is harmed much more by the individual who recidivates right away because it essentially gets no benefit from releasing him because it avoids the cost of incarcerating him for only a short time period. Thus, a parole board that is not racially prejudiced will work to ensure that all released individuals recidivate at the same rate, because this measure takes into account not just whether an individual will recidivate but when.

We cannot explicitly identify the recidivism rate in our data because doing so would require averaging the number of crimes an individual commits in a given release period across all members of his race. In our data we observe only whether an individual was returned to prison for committing at least one crime but do not observe how many crimes he committed during that period. We can, however, indirectly estimate the rate of recidivism by exploiting the fact that this rate will positively affect the probability that an individual will recidivate at least once within a given release period. Because only nonrehabilitated types recidivate, and they do so at Poisson arrival rate $g_{r_i} > 0$, the probability that inmate i will recidivate at least once within a given amount of time t_i, is

$$\left(1 - \pi^*_{r_i}\right)\left[1 - e^{-g_{r_i} t_i}\right]. \tag{6}$$

where t_i is inmate i's exposure time.

Figure 3 illustrates this expression for members of a given race who have varying exposure times t_i. The probability that an individual recidivates at least once within his exposure time is positively related to his exposure time and asymptotically approaches the proportion of race-r individuals who are not rehabilitated. Because an individual is expected to commit a certain number of crimes in a given period (defined by his rate), the longer we observe him, the more likely it is that he has committed at least one crime. Once we observe them for a long enough time frame, we would expect all nonrehabilitated individuals to have recidivated at least once; rehabilitated individuals will never recidivate, and thus the curve approaches $(1 - \pi^*_{r_i})$.

Taking a second-order Taylor series approximation of the above curve explicitly shows how the recidivism rate of race-r members affects the shape of the curve:

$$\begin{aligned} \left(1 - \pi^*_{r_i}\right)\left[1 - e^{-g_{r_i} t_i}\right] &\approx \left(1 - \pi^*_{r}\right)\left\{1 - \left[1 - g_{r_i} t_i + g^2_{r_i} t^2_i\right]\right\} \\ &= \left(1 - \pi^*_{r}\right) g_{r_i} t_i - \left(1 - \pi^*_{r}\right) g^2_{r_i} t^2_i \qquad (7) \\ &= \left(\frac{B - D_r}{C}\right) \times t_i - \left(\frac{B - D_r}{C} g_{r_i}\right) \times t^2_i, \end{aligned}$$

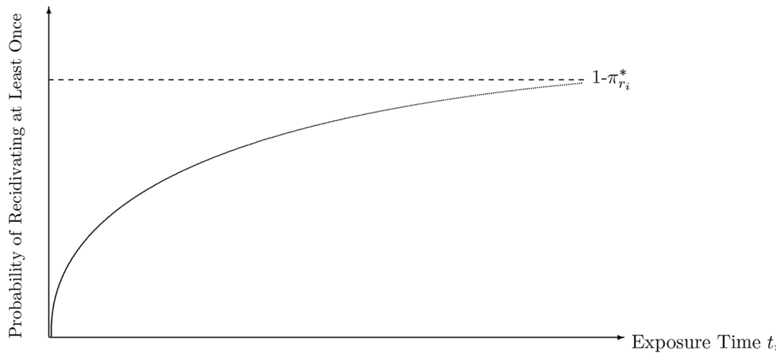

Figure 3. Probability of race-*r* parolee recidivating at least once as a function of exposure time.

where the last equality follows from inserting the expression π^*_r from equation (1).

Note that the coefficient on exposure time t_i, which measures the slope of the curve at the origin, exactly corresponds to the recidivism rate of race-*r* members whom we are trying to identify. Intuitively, it makes sense that the recidivism rate defines the slope at the origin, as this slope closely corresponds to the proportion of people expected to have recidivated at least once at the end of the first release period. All else equal, if a race of individuals is released at a higher recidivism rate, either because its members are rehabilitated at a lower rate or because its members have a higher recidivism rate conditional on being nonrehabilitated, one would expect a higher fraction of those individuals to recidivate right away.[20]

We can thus estimate the recidivism rate of each race by estimating equation (8) in a standard regression framework. To estimate the rate for each race separately, we perform the following regression:

$$\begin{aligned}\text{Recidivate}_i &= \alpha_1 \times t_i + \beta_1 \times \text{Minority}_i \times t_i + \alpha_2 \times t_i^2 \\ &\quad + \beta_2 \times \text{Minority}_i \times t_i^2 + \varepsilon,\end{aligned} \tag{8}$$

where Recidivate_i is an indicator for whether the parolee recidivates within his exposure time and Minority_i is an indicator for whether inmate

20. Note that although at the origin the curve shape is affected only by the product of the two components of the recidivism rate $(1-\pi^*_{r_i})$ and g_{r_i}, as we move farther from the origin, the individual components of the rate have distinct effects on the shape of the curve. At higher exposure times, the race with the higher nonrehabilitated rate will have a curve that lies above that of the race with a lower nonrehabilitated rate. Thus, because our test involves comparing the overall recidivism rate, we focus on the slope near the origin.

i is a minority.[21] The coefficient α_1 is our estimate of the expected recidivism rate for whites, while $(\alpha_1 + \beta_1)$ is the expected recidivism rate for minorities. Recall that if the parole board is racially prejudiced against minorities, it will result in minority inmates having a lower expected recidivism rate. Thus, our test for racial prejudice will be whether $\beta_1 < 0$. Note that our test does not require us to have any observable information about the inmate except his race, exposure time, and whether he recidivated during this exposure time.

In estimating equation (9), we are essentially estimating a curve similar to that shown in Figure 3 for each race. We thus need to have members of a given race spread out among different exposure times. To do this, we define exposure time as the number of months from an individual's release date from prison until January 1, 2004. As discussed in Section 5, our data include all individuals released from prison in Pennsylvania between 1999 and 2003. This results in the individuals in our sample having exposure times that vary from 1 day to 5 years.[22]

It is important to note that the only reason we can use equation (9) to estimate the recidivism rate for race-r inmates is because all inmates of race r are released with the same π_r. This ensures that the coefficient on t_i is race specific and thus can be estimated. If instead individuals within a race were released at different values of π, the coefficient would be individual specific and could not be estimated. This issue is more generally known as the inframarginality problem and is a common problem for outcome tests. We avoid this issue because in our context the parole board can perfectly adjust the treatment variable (time served) to ensure that everyone has the same rate. This point highlights why our test is not valid for those who are released at their minimum or maximum sentence, since those inmates are released with various rates of rehabilitation.

Our solution to the inframarginality problem also highlights our core assumption: that the parole board is able to release individuals at exactly t_i^*. This assumption is consistent with the way in which the parole process

21. In Section 6 we estimate this separately for blacks and Hispanics.

22. Note that even if there are systematic differences between individuals with different exposure times (resulting from them being released at different times), the parole board can still ensure that everyone within a race is released with the same predicted rate of recidivism. For example, suppose that the environment into which these individuals are released changes over time, which would affect the rate g_r at which nonrehabilitated inmates recidivate. Our model predicts that the board will respond by changing the probability threshold π_r at which these race-r members need to be rehabilitated so that the resulting predicted recidivism rate at which they are released remains unchanged (and independent of exposure time).

works. If individuals are turned down for parole on completion of their minimum sentence, they are given a list of things to do and the time when they will next come up for parole. This time window is variable from inmate to inmate and is at the complete discretion of the parole board. We can think of the time window the parole board sets for inmate i as the extra incident-free time that is needed for its perception that inmate i is rehabilitated to reach the race-specific threshold π_r^*. Thus, for all intents and purposes, the parole board is able to release every inmate at the optimum time.[23]

4.3. Testable Implications of the Model

Because our test for prejudice comes directly from our model, it is important to conduct some validity checks. In this section, we delineate three implications of the model that can be directly tested. These tests are performed in Section 6.1.

The first testable implication is based on equation (8), which estimates the curve shown in Figure 3 separately for each race. This equation implies that the probability that a race-r inmate recidivates within his exposure time is positively related to the length of his exposure time and is negatively related to the square of his exposure time. (The intuition behind the positive coefficient on exposure time was explained in Section 4.2.) The intuition behind the negative coefficient on t_i^2 is as follows: At an exposure time of 0, the full $(1-\pi_r^*)$ fraction of race-r individuals is eligible to recidivate. As exposure time increases and more of them have recidivated, there is an increasingly smaller stock of individuals who can transition to the state of recidivating at least once. Thus, we would expect the probability of recidivating at least once to increase but at a decreasing rate. Note that this equation applies only to individuals released between $\underline{T}$ and $\overline{T}$, and thus we restrict the analysis to these individuals when we perform the test.

The second testable implication is based on Figure 1, which implies that, within a race, all individuals released between $\underline{T}$ and $\overline{T}$ will recidivate at exactly the same rate. As Figure 1 shows, this implies that their re-

23. If individuals are turned down for parole at their next opportunity, it is usually because they have had some type of misconduct during that time window (that is, they have been involved in prison incidents) and/or they did not complete their required programming. With our current model, this implies that these individuals would be considered nonrehabilitated and would never be released. However, in Section A2 we expand our model to show that we can easily allow rehabilitated individuals to be involved in some incidents while in prison, and thus it is possible for these individuals to eventually be released on parole.

cidivism rate will be independent of the fraction of the assigned sentence they serve.

The third testable implication is based on the probability that a race-r inmate released between $\underline{T}$ and $\overline{T}$ recidivates within a certain time window P:

$$\left(1-\pi_r^*\right)\left[1-e^{-g_r(P)}\right]. \tag{9}$$

This expression is similar to equation (6), except now the time window over which we examine the probability of recidivism is not inmate specific. Recall that any race-r inmate released exactly at $\underline{T}$ should have a rehabilitation probability π_i that is at or above π_r^*. This means that the average probability of recidivating within P among all race-r inmates released at $\underline{T}$ should be lower than the average among race-r inmates released between $\underline{T}$ and $\overline{T}$ (which is given by equation [10]). Likewise, among race-r inmates released at $\overline{T}$, $\pi_i \leq \pi_r^*$, which implies that their average probability of recidivating within P should be greater than for the group released between $\underline{T}$ and $\overline{T}$. Thus, we should find that the probability of recidivating within P should be in ascending order for those inmates released immediately after serving their minimum sentences, those released between their minimum and maximum sentences, and those released after serving their maximum sentences.

4.4. Notes about the Model

4.4.1. Objectivity of the Outcome Measure. One of the key requirements of using an outcome test is that the outcome must be objectively measured.[24] Our outcome is based on recidivism, which we measure in the two ways most commonly used in the literature: a return to prison for the commission of a new crime and a return to prison because the individual was convicted of a new crime or committed a technical parole violation. For the outcome to be truly objective, all individuals who committed a new crime or technical parole violation would need to have the same probability of being returned to prison regardless of race. The fact that there is a body of literature showing racial disparities at various stages of the criminal justice process (referenced in the introduction) suggests that

24. An example of an objective outcome is the success rate of a motor vehicle search, which is used in the outcome tests examining racial prejudice in such searches. Once a vehicle is searched, whether or not contraband is found is completely based on whether or not the individual is carrying it and is independent of the officer's behavior.

this may not be true. This section discusses the potential implications of having a nonobjective outcome measure.

Since our test is designed to detect racial prejudice on the part of the parole board, the biggest potential hindrance to the validity of the test would be if the parole board could affect the outcome. The board would then be able to cover up racial prejudice at the parole release stage with further prejudice at the outcome stage. Fortunately, the parole board does not have much impact on either of the recidivism measures used. The parole board has no direct involvement in whether an inmate is returned to prison for a new crime; the parole board has some input into whether an inmate will be returned to prison for a technical parole violation, but ultimately it is up to an independent hearing examiner to decide.

The recidivism measures we use can, however, be impacted by any racial prejudice on the part of other downstream agents such as police officers, prosecutors, judges, and parole agents. This downstream prejudice would result in minorities having a higher g_r not because their nonrehabilitated members are more likely to commit crimes but simply because (conditional on committing a crime) they are more likely to be convicted. Recall that a nonprejudiced parole board will respond to differences in g_r across races by requiring minorities to serve more time so that they have a higher probability of being rehabilitated before release. This means that minorities might be serving more time simply because of downstream racial prejudice, which goes beyond the traditional definition of statistical discrimination. While our test will still correctly conclude that there is no racial prejudice on the part of the parole board, it has no power to pick up the latter suboptimal result.

4.4.2. Crime Controls. Until now we have assumed that the cost to the parole board that results from an individual recidivating, denoted C, is the same across all prisoners. However, the cost to the parole board is likely to be affected by the type of offender: the parole board is likely to view recidivism by a violent offender as more costly than recidivism by a drug offender. From equation (1), one can see that if the cost to the parole board from the individual recidivating is higher, it will respond by making those types of offenders have a higher probability of being rehabilitated, which results in a lower expected rate of recidivism. Because offender type is strongly related to race, it is important to control for this when estimating equation (9). We separate crimes into three groups: murder or sex crimes, violent crimes other than murder or sex crimes, and nonviolent crimes such as drug or property crimes. As noted in Section 3, the parole board has more stringent release procedures for the release of the first two groups of offenders, which implies that there is a higher cost when these groups recidivate. Note that in order to allow different types of offenders to have different recidivism rates, we need to estimate the curve shown in Figure 3 separately for these groups. This requires us to interact indicators for crime group with both exposure time and the square of exposure time.

5. DATA

We use data from the Pennsylvania Department of Corrections for all individuals who were released from prison between January 1, 1999, and December 31, 2003. The data include individuals who were released before the completion of their maximum sentence and were thus on parole from their release date until their maximum sentence expired; it also includes individuals who were released at the completion of their maximum sentence and thus spent no time on parole. We restrict the data set to individuals who were new court admissions when they first were included in our data set and who were white, black, or Hispanic males. We are left with a total of 26,343 individuals. We observe their sentence length (minimum and maximum) prescribed by the judge, admission date, release date, the completion dates of their minimum and maximum sentences, name, state identification number, date of birth, and main offense committed.[25]

We also observe each prisoner's subsequent returns to prison before March 31, 2009 (if any). We have information on the date and the reason for the return: a new crime or a technical parole violation.[26,27] We use this information to code the two key dependent variables in the regression

25. The date the minimum sentence is completed is often different than just the sum of the prison admission date and the minimum assigned sentence. Many individuals who cannot afford bail (or are deemed too risky) spend time in jail while they are awaiting formal sentencing and get credit for this time served once the formal sentence is handed down. Having the date the minimum sentence is completed allows us to accurately identify individuals who are released right after serving their minimum sentence (by comparing the prison release date with the minimum sentence completion date). It also allows us to accurately calculate time served as the minimum assigned sentence plus the difference between the prison release date and the minimum sentence completion date.

26. Note that our data pick up whether an offender commits a new crime only if it results in him going back to prison. This should not pose too much of a problem, as individuals will be on parole for the majority of the time that we need to observe their recidivism behavior; while on parole any new conviction should automatically send the offender back to prison.

27. We do not observe the specific reason for the technical parole violation, nor do we observe whether the parolee committed multiple violations.

form of our test for racial prejudice specified in equation (9): whether an individual returns to prison because of the commission of a new crime within his exposure time and whether an individual returns to prison because of the commission of a new crime or a technical parole violation within his exposure time. For each of the individuals released between his minimum and maximum sentence, we calculate his exposure time as the number of months between his date of release and January 1, 2004.[28] We then code an indicator variable for whether he recidivates within this exposure time by examining whether his return to prison is before January 1, 2004.[29]

Table 2 provides some descriptive statistics for our sample. The majority of our sample consists of blacks and Hispanics. The type of crimes individuals commit varies significantly by race, with whites more likely to commit crimes in the most serious category (murder or sex crimes). Blacks tend to be assigned and serve longer sentences than both whites and Hispanics. However, whites are more likely than blacks or Hispanics to be required to serve their full sentence.

6. EMPIRICAL ANALYSIS

6.1. Testing Model Implications

This section presents the results of the three model tests that were outlined in Section 4.3. These model checks should all hold within a race, and so we conduct the tests separately by race and show the results for all races pooled for completeness.

The first test estimates equation (8) by regressing whether an individual recidivates at least once within his exposure time on Exposure Time and Exposure Time^2 (without a constant). As noted in Section 4.3, the coefficients on the variables should be positive and negative, respectively, when the sample is restricted to those individuals who are released be-

28. Note that we choose January 1, 2004, as our cutoff (even though we observe recidivism behavior up through early 2009) because when estimating Figure 3 it is important that we have some individuals with exposure times that are very short (that is, close to the origin). Since our sample includes people released from 1999 to 2003, this will give us individuals with exposure times that are reasonably spread out from 1 day to 5 years. If we had instead used early 2009 as the cutoff, we would not have any individuals with exposure times close to the origin.

29. When recidivism is measured by the first measure, we compare the date the individual returns for committing a new crime with January 1, 2004. When recidivism is measured by the second measure, we compare the earliest date of return (either for a technical parole violation or new crime) with January 1, 2004.

Table 2. Descriptive Statistics

Variable	All	Whites	Blacks	Hispanics
Race:				
White	.356	1	0	0
Black	.515	0	1	0
Hispanic	.129	0	0	1
Age at release:				
18–25	.247	.200	.276	.259
26–35	.370	.328	.388	.410
36–45	.257	.301	.231	.239
46–55	.098	.126	.085	.076
56+	.029	.045	.020	.017
Crime type:				
Murder or sex	.103	.144	.084	.066
Other violent	.286	.260	.332	.176
Property	.168	.248	.134	.084
Drug	.297	.126	.348	.563
Other	.146	.223	.102	.110
Sentence length (months):				
Minimum	33.8	31.0	36.2	32.1
Maximum	82.5	78.3	86.7	77.7
Served	47.6	46.1	50.0	42.4
Released:				
At minimum sentence	.313	.313	.302	.360
Between minimum and maximum sentences	.491	.455	.513	.500
At maximum sentence	.197	.233	.189	.141
Recidivism measures:				
New crime within exposure time	.119	.106	.129	.112
New crime or technical parole violation within exposure time	.384	.358	.410	.345
N	26,343	9,384	13,571	3,388

Note. Estimates for recidivism measures include only inmates who were released between their minimum and maximum sentences.

tween their minimum and maximum sentences. Table 3 shows the results of these regressions for both recidivism measures. All of the coefficients have the predicted sign, and all except one are statistically significant.

The second test checks whether the amount of an assigned sentence an offender serves is unrelated to his recidivism rate. As Figure 1 indicates, all individuals within a race released between their minimum and maximum sentences should recidivate at exactly the same rate regardless of sentence. To determine this, we regress whether an individual recidivates at least once within his exposure time on Exposure Time and Exposure

Table 3. Relationship between Exposure Time and Recidivism

	New Crime within Exposure Time		New Crime or Parole Violation within Exposure Time	
All (N = 12,921):				
Exposure Time	.00477**	(.00034)	.02322**	(.00051)
Exposure Time^2	–.00002**	(.00001)	–.00026**	(.00001)
Whites (N = 4,268):				
Exposure Time	.00425**	(.00056)	.02210**	(.00088)
Exposure Time^2	–.00002	(.00001)	–.00025**	(.00002)
Blacks (N = 6,960):				
Exposure Time	.00499**	(.00047)	.02410**	(.00070)
Exposure Time^2	–.00002$^+$	(.00001)	–.00027**	(.00001)
Hispanics (N = 1,693):				
Exposure Time	.00519**	(.00093)	.02227**	(.00141)
Exposure Time^2	–.00003$^+$	(.00002)	–.00026**	(.00003)

Note. Estimates are from ordinary least squares regressions without a constant and include only inmates who were released between their minimum and maximum sentences. Exposure time is measured in months. Heteroskedasticity-robust standard errors are in parentheses.

$^+$ $p < .10$.

** $p < .01$.

Time^2 and on these two variables interacted with actual sentence length, the assigned minimum and maximum sentence lengths, and indicators for whether the crime was a murder or sex crime or another violent crime.[30] This specification essentially estimates a curve similar to that in Figure 3 for everyone who commits the same type of crime and has the same assigned sentence.[31] The coefficient on the interaction between exposure time and sentence length then reveals whether the recidivism rate (and thus the resulting curve) depends on the actual sentence served. Table 4 shows the results of these regressions; for brevity we show only the key coefficients. In all but one instance, the coefficient on Exposure Time × Sentence Length is statistically insignificant, which implies that the fraction of sentence served is unrelated to an individual's recidivism rate.

The third test compares the probability of recidivating at least once within a given time period for individuals released at different points of

30. Note that crime controls are necessary here because sentence length is likely correlated with crime type. The more severe the crime, the lower the predicted recidivism rate at which an individual will be released; on average, the parole board must keep more severe offenders incarcerated longer to reach the lower rate. It is still the case, however, that within a crime type, sentence length should have no effect on the recidivism rate.

31. Note that in order to do this, any control variable must be fully interacted with both Exposure Time and Exposure Time^2.

Table 4. Relationship between Sentence Served and Recidivism

	New Crime within Exposure Time		New Crime or Parole Violation within Exposure Time	
All (N = 12,921):				
Exposure Time	.00579**	(.00060)	.02485**	(.00090)
Exposure Time × Sentence Length	−.00002	(.00002)	−.00001	(.00004)
Exposure Time^2	−.00003$^+$	(.00001)	−.00031**	(.0002)
Exposure Time^2 × Sentence Length	.00000	(.00000)	.00000	(.00000)
Whites (N = 4,268):				
Exposure Time	.00512**	(.00101)	.02496**	(.00156)
Exposure Time × Sentence Length	.00003	(.00004)	−.00004	(.00007)
Exposure Time^2	−.00002	(.00002)	−.00032**	(.00003)
Exposure Time^2 × Sentence Length	−.00000	(.00000)	.00000	(.00000)
Blacks (N = 6,960):				
Exposure Time	.00591**	(.00084)	.02520**	(.00125)
Exposure Time × Sentence Length	−.00003	(.00003)	.00004	(.00005)
Exposure Time^2	−.00002	(.00002)	−.00030**	(.00003)
Exposure Time^2 × Sentence Length	.00000	(.00000)	−.00000	(.00000)
Hispanics (N = 1,693):				
Exposure Time	.00725**	(.00171)	.02611**	(.00254)
Exposure Time × Sentence Length	−.00006	(.00008)	−.00022$^+$	(.00013)
Exposure Time^2	−.00005	(.00004)	−.00033**	(.00006)
Exposure Time^2 × Sentence Length	.00000	(.00000)	.00001*	(.00000)

Note. Estimates are from ordinary least squares regressions without a constant and include only inmates who were released between their minimum and maximum sentences. Exposure time is measured in months. Heteroskedasticity-robust standard errors are in parentheses. All regressions include Exposure Time and Exposure Time^2 interacted with the assigned minimum and maximum sentence length and indicators for whether the crime was a murder or sex crime or another violent crime.

$^+$ $p < .10$.

* $p < .05$.

** $p < .01$.

their assigned sentences. As discussed in Section 4.3, we would expect this probability to be the lowest for individuals released on the completion of their minimum sentence and the highest for individuals released when their maximum sentence is completed. The recidivism probability for individuals released between their minimum and maximum sentences should be between these two extremes. Table 5 presents the results from regressing the likelihood of recidivating within a certain time period on indicators for when the individual was released, including crime controls (whose coefficients are not shown for brevity). For robustness, we use two different time frames for each of our recidivism measures, namely, 3 years after release and 5 years after release. Because individuals released on completion of their maximum sentence cannot return to prison for a parole violation (since they are not released on parole), we exclude them and compare only recidivism probabilities between the other two groups. The constant coefficient shows the corresponding recidivism probability for those individuals released between their minimum and maximum sentences. We would expect the coefficients on the indicator for being released at the minimum sentence to be negative and on the indicator for being released at the maximum to be positive (when used). This is precisely what we find, although not all coefficients are statistically significant.

6.2. Main Result: Test for Racial Prejudice

In this section we implement the test implied by the model to determine whether there is evidence that racial prejudice plays a role in the parole board's discretionary parole release decisions. The regression results reported in Table 6 correspond to the test outlined in equation (9) and use only inmates who are released between their minimum and maximum sentences. The coefficient on Exposure Time corresponds to the recidivism rate for whites (that is, the expected number of times a white individual would be expected to recidivate within 1 month). The coefficients on Exposure Time × Black and Exposure Time × Hispanic reveal whether blacks and Hispanics, respectively, recidivate at a different rate than whites. Regardless of the recidivism measure used, these coefficients are always small and statistically insignificant, which implies that all racial groups are released at the same recidivism rate threshold, and thus we conclude that the parole board is not racially prejudiced in its parole release decisions.

7. CONCLUSION

In this paper we develop a model of a parole board contemplating whether to grant parole to a prisoner who has finished serving his minimum sentence. In our model the parole board chooses to grant the pris-

Table 5. The Relationship between Recidivism and Release Time

	New Crime				New Crime or Parole Violation			
	Within 3 Years		Within 5 Years		Within 3 Years		Within 5 Years	
All (N = 26,343):								
Released at Minimum	–.024**	(.005)	–.034**	(.006)	–.040**	(.007)	–.040**	(.007)
Released at Maximum	.031**	(.006)	.034**	(.007)				
Constant	.181**	(.004)	.272**	(.005)	.487**	(.005)	.546**	(.005)
Whites (N = 9,384):								
Released at Minimum	–.035**	(.008)	–.042**	(.010)	–.073**	(.012)	–.074**	(.012)
Released at Maximum	.018+	(.009)	.017	(.011)				
Constant	.167**	(.007)	.252**	(.008)	.474**	(.000)	.525**	(.009)
Blacks (N = 13,571):								
Released at Minimum	–.009	(.007)	–.020*	(.009)	–.015	(.010)	–.018*	(.010)
Released at Maximum	.030**	(.009)	.037**	(.010)				
Constant	.189**	(.006)	.284**	(.007)	.506**	(.008)	.572**	(.008)
Hispanics (N = 3,388):								
Released at Minimum	–.042**	(.013)	–.050**	(.016)	–.024	(.019)	–.017	(.019)
Released at Maximum	.096**	(.022)	.105**	(.024)				
Constant	.180**	(.010)	.269**	(.012)	.452**	(.014)	.505**	(.014)

Note. Estimates are from ordinary least squares regressions with a constant (inmates released between their minimum and maximum sentences are the benchmark group) and include the full sample. Heteroskedasticity-robust standard errors are in parentheses. All regressions include indicators for whether the crime was a murder or sex crime or another violent crime.

+ p <.10.

* p <.05.

** p <.01.

Table 6. Relationship between Race and Recidivism

	New Crime within Exposure Time		New Crime or Parole Violation within Exposure Time	
Exposure Time	.00505**	(.000650)	.0237**	(.000982)
Exposure Time × Black	.000429	(.000730)	.00154	(.00113)
Exposure Time × Hispanic	.000477	(.00108)	−.000623	(.00166)
Exposure Time × Violent Crime	−.00269**	(.000871)	−.00799**	(.00161)
Exposure Time × Other Violent Crime	−.000929	(.000743)	−.00117	(.00114)
Exposure Time^2	−.0000177	(.0000142)	−.000277**	(.0000208)
Exposure Time^2 × Black	.00000544	(.0000160)	−.00000554	(.0000239)
Exposure Time^2 × Hispanic	−.0000118	(.0000234)	.000000972	(.0000350)
Exposure Time^2 × Violent Crime	−.00000385	(.0000190)	.000101**	(.0000343)

Note. Estimates are from ordinary least squares regressions without a constant and include only inmates who were released between their minimum and maximum sentences. Exposure time is measured in months. Heteroskedasticity-robust standard errors are in parentheses. N = 12,921.

** $p < .01$.

oner parole if and only if the assessed recidivism rate is at or below a threshold, with the threshold being lower for minority prisoners if the parole board is prejudiced against minorities. We show that when inmates complete incident-free time periods in prison, the parole board responds by revising downward its perception of the inmate's rate of recidivism on release. Because the parole board has complete discretion over when to release prisoners within the constraints of their minimum and maximum sentences, this results in all prisoners released between these bounds being released at exactly the point at which their rate of recidivism reaches the optimum race-specific threshold.

Our model implies that we can identify the race-specific thresholds used by simply identifying the race-specific average rate of recidivism for those individuals released between their minimum and maximum sentences. This approach is immune to the inframarginality problem because, within these bounds, the marginal prisoner released is exactly the same as the average prisoner released. Using data on all prisoners released in Pennsylvania from 1999 to 2003, we find no evidence of racial prejudice on the part of the parole board.

Finally, it is important to point out that while our test shows no evidence of racial prejudice, it does not necessarily mean that parole release decisions do not reflect discrimination. The parole board can also engage in statistical discrimination, a practice that is illegal but that our test does not detect. This highlights the use of tests that are action based as a complement to our outcome-based approach, since the former will detect whether observationally equivalent minority and white inmates are treated differently for any reason. Because our test shows no evidence of racial prejudice, any difference detected is likely due to statistical discrimination. We do not conduct these tests here, as we do not have the requisite controls to rule out omitted-variables bias. However, obtaining a detailed record of the information available to the Pennsylvania parole board that would enable these action-based tests to be conducted would be an important focus of future research, especially since Pennsylvania has a large prison population and parole release continues to be the sole way in which prisoners can obtain an early release.

APPENDIX: COMPARISON AND ROBUSTNESS OF THE MODEL

A1. Comparison with Mechoulan and Sahuguet's Study

In Section 2 we mentioned that Mechoulan and Sahuguet (2015) is the only other work that uses an outcome test to study prejudice in the parole board release process. While their paper comes to a conclusion similar to ours—namely, there is no evidence of racial prejudice against blacks in the parole release process—the papers use reasonably different tests for discrimination on different data sets.[32] This section details both the theoretical and empirical differences between the papers.

The key difference between our paper and Mechoulan and Sahuguet (2015) is that we model the incentives of the parole board differently and thus end up conducting a different test for racial prejudice. In their model the parole board aims to minimize the total number of individuals who recidivate while on parole, and thus, in the absence of prejudice, all prisoners are released such that their probability of recidivating while on parole is the same. The parole board can accomplish this because of the mechanical relationship between recidivating on parole and time on parole—the less time an individual spends on parole, the lower the probability he will recidivate on parole. Thus, when we examine the individuals who are released, they will have different recidivism rates; however, those with high recidivism rates will be released only when they have a short amount of time left on parole so that the overall likelihood of recidivating on parole is equal-

32. Mechoulan and Sahuguet do not include Hispanics in their sample.

ized across all inmates. In contrast, in our model the parole board compares the benefits and costs of keeping an inmate for the next period, which results in it releasing all inmates with the same rate of recidivism in the absence of prejudice. Realistically, the core difference between the models is that Mechoulan and Sahuguet assume that the parole board views recidivism at any time during parole to bear the same net cost; in contrast, our model assumes that the parole board finds recidivism that comes closer to the release time to have a higher net cost than recidivism long after release. This results in Mechoulan and Sahuguet comparing probabilities of whether individuals ever recidivate, while we compare rates that incorporate not just if an individual recidivates, but when.

The papers also differ in the data used. We use a detailed data set that is specific to prison releases in Pennsylvania. Mechoulan and Sahuguet (2015) uses data from the National Corrections Reporting Program (NCRP), which is less detailed but covers many states; they do not report results from Pennsylvania in their study.

In order to assess the practical differences between these tests, it is useful to run both tests on the same data. Table A1 presents the results from performing Mechoulan and Sahuguet's test on our data. We regress the probability that an inmate recidivates while on parole on an indicator for whether he is black or Hispanic.[33] Mechoulan and Sahuguet (2015) uses two different definitions of recidivism: whether an inmate returns to prison because of the commission of a new crime and whether the inmate returns to prison for any reason (columns 5 and 6). While we use the first measure as well, we do not find the latter measure to be a good definition of recidivism. Our data reveal that approximately 14 percent of returns include individuals who are charged with committing either a technical parole violation or a new crime but for whom the charges were later dropped and they were subsequently released. It thus does not seem accurate to code these individuals as recidivating.[34] To see how robust the results are to the recidivism definition used, we include in columns 3 and 4 our other definition of recidivism, which is a return to prison for either a new crime or technical parole violation. Like ours, Mechoulan and Sahuguet's test is designed to be performed on only the individuals whom the parole board releases between their minimum and maximum sentences. However, the NCRP data include all individuals released on parole and do not identify which individuals were released exactly at their minimum sentence and which were released between their minimum and maximum sentences; Mechoulan and Sahuguet thus must include both groups of individuals in their sample. To mirror this sample specification, columns 1, 3, and 5 present the results of the analysis on all individuals released on parole; columns 2, 4, and 6 restrict the sample to those on whom the test was designed to be performed.

33. We exclude individuals for whom we do not observe the full time period during which they are on parole.

34. There are large racial differences in who is affected by these mistaken returns: 16 percent of the returns by blacks fall into this category, in contrast to only 9 percent of the returns by whites.

Table A1. Relationship between Race and the Likelihood of Recidivating while on Parole

	New Crime		New Crime or Parole Violation		Return to Prison	
	(1)	(2)	(3)	(4)	(5)	(6)
Black	.0290**	.0183**	.0717**	.0512**	.111**	.0971**
	(.00523)	(.00666)	(.00798)	(.0103)	(.00798)	(.0103)
Hispanic	.0124⁺	.0146	.0291*	.00675	.0131	−.00968
	(.00752)	(.00973)	(.0115)	(.0150)	(.0115)	(.0150)
Constant	.104**	.106**	.417**	.434**	.449**	.461**
	(.00405)	(.00521)	(.00617)	(.00804)	(.00617)	(.00803)
R^2	.002	.001	.004	.002	.012	.010

Note. Estimates are from ordinary least squares regressions with a constant and include only inmates for whom the entire time on parole is observed (parole expires before March 31, 2009). Exposure time is measured in months. Standard errors are in parentheses. The sample in columns 1, 3, and 5 comprises all parolees, which includes inmates released either after serving their minimum sentence or between their minimum and maximum sentences (N = 18,805); the sample in columns 2, 4, and 6 comprises inmates released between their minimum and maximum sentences (N = 11,412).

⁺ $p < .10$.
* $p < .05$.
** $p < .01$.

The results imply that regardless of the sample specification or recidivism measure used, the coefficient on blacks is always positive and statistically significant. This is the same result that that Mechoulan and Sahuguet find in their paper. It is reassuring that these different approaches both reach the conclusion that there is no racial bias against blacks in the parole release process.[35]

As a final exercise we test whether the data are more consistent with the predictions of our model or Mechoulan and Sahuguet's model. One key testable implication in which our models differ is the effect of parole time on the probability of recidivism on parole. We detail in Section 4.3 that our model predicts that the longer an individual is observed, the more likely we are to observe him recidivating. Applied to Mechoulan and Sahuguet's setting, this would mean that the longer an individual is on parole, the more likely he is to recidivate while on parole. Their model, however, says that all individuals should have the same probability of recidivating on parole and thus time on parole should have no effect on the probability of recidivism. In fact, in their model this is the strategic variable the parole board uses to ensure that recidivism probabilities are equal—inmates who have a higher recidivism rate will be released with less time on parole so that they

35. Although technically Mechoulan and Sahuguet's model would interpret this positive coefficient as the parole release process favoring blacks, they realistically suggest that it is likely to reflect the parole board's trade-off between equalizing recidivism outcomes and equalizing timing of release.

will (mechanically) have less chance of recidivating while on parole. Table A2 presents the specifications from Table A1 with a control for months on parole to explicitly determine the relationship between parole time and the probability of recidivating on parole. Regardless of the specification used, the coefficient on parole time is always positive and strongly statistically significant, which supports our model.

A2. Robustness of the Model to Rehabilitated Prisoners' Involvement in Prison Incidents

In the basic model we assume that only nonrehabilitated inmates are involved in prison incidents. Now suppose that rehabilitated prisoners can also be involved in such incidents but at a lower rate than nonrehabilitated types. In particular, suppose that race-r inmates are involved in prison incidents with Poisson arrival rate λ_{1r} if they are nonrehabilitated and λ_{0r} if they are rehabilitated, with $\lambda_{1r} > \lambda_{0r} \geq 0$. Here we demonstrate that all of the key implications of the model still hold.

Similarly to the derivation of the belief evolution equation (2), we can show that if there is no occurrence of incident at time t, then the parole board's belief that i is rehabilitated evolves according to

$$\dot{\pi}_t^i = \left(\lambda_{1r_i} - \lambda_{0r_i}\right)\pi_t^i\left(1 - \pi_t^i\right). \tag{A1}$$

Note that equation (A1) coincides with equation (2) when $\lambda_{0r_i} = 0$.

On the other hand, if an incident occurs at time t, the parole board's revision of its belief about the inmate will not decrease to 0 as in the basic case in which $\lambda_{0r} = 0$. However, the posterior will still exhibit a discrete drop whose magnitude is derived as follows: Consider a short time interval between t and $t + \Delta$. If an incident occurs in the interval, then $\pi_{t+\Delta}^i$ can be obtained using Bayes's rule as

$$\pi_{t+\Delta}^i = \frac{\pi_t^i\left[1 - e^{-\lambda_{0r_i}\Delta}\right]}{\pi_t^i\left[1 - e^{-\lambda_{0r_i}\Delta}\right] + \left(1 - \pi_t^i\right)\left[1 - e^{-\lambda_{0r_i}\Delta}\right]}.$$

Thus,

$$\begin{aligned} \lim_{\Delta\to 0}\left(\pi_{t+\Delta}^i - \pi_t^i\right) &= \lim_{\Delta\to 0}\frac{\pi_t^i\left(1 - \pi_t^i\right)\left(e^{-\lambda_{1r_i}\Delta} - e^{-\lambda_{0r_i}\Delta}\right)}{\pi_t^i\left[1 - e^{-\lambda_{0r_i}\Delta}\right] + \left(1 - \pi_t^i\right)\left[1 - e^{-\lambda_{1r_i}\Delta}\right]} \\ &= \frac{\left(\lambda_{0r_i} - \lambda_{1r_i}\right)\pi_t^i\left(1 - \pi_t^i\right)}{\pi_t^i\lambda_{0r_i} + \left(1 - \pi_t^i\right)\lambda_{1r_i}}. \end{aligned} \tag{A2}$$

Note that expression (A2) implies that $\lim_{\Delta\to 0}\left(\pi_{t+\Delta}^i - \pi_t^i\right) = -\pi_t^i$ when $\lambda_{0r_i} = 0$, which coincides with our basic case in which following an incident the parole board's posterior belief decreases to 0.

Table A2. Relationship between Time on Parole and the Likelihood of Recidivating while on Parole

	New Crime		New Crime or Parole Violation		Return to Prison	
	(1)	(2)	(3)	(4)	(5)	(6)
Black	.0279**	.0170*	.0691**	.0479**	.108**	.0929**
	(.00520)	(.00662)	(.00788)	(.0101)	(.00783)	(.00997)
Hispanic	.0126+	.0156	.0296**	.00936	.0137	–.00645
	(.00749)	(.00967)	(.0113)	(.0147)	(.0113)	(.0146)
Months on parole	.00149**	.00156**	.00348**	.00412**	.00438**	.00512**
	(.000108)	(.000129)	(.000164)	(.000197)	(.000162)	(.000195)
Constant	.0568**	.0580**	.307**	.307**	.310**	.303**
	(.00529)	(.00653)	(.00802)	(.00996)	(.00796)	(.00985)
R^2	.002	.001	.004	.002	.012	.010

Note. Estimates are from ordinary least squares regressions with a constant and include only inmates for whom the entire time on parole is observed (parole expires before March 31, 2009). Exposure time is measured in months. Standard errors are in parentheses. The sample in columns 1, 3, and 5 comprises all parolees, which includes inmates released either after serving their minimum sentence or between their minimum and maximum sentences (N = 18,805); the sample in columns 2, 4, and 6 comprises inmates released between their minimum and maximum sentences (N = 11,412).

+ $p < .10$.

* $p < .05$.

** $p < .01$.

Therefore, in this extended environment where rehabilitated inmates may also be involved in prison incidents, the parole board's evolution of beliefs becomes more complicated, as it increases continuously with episodes of no incidents but exhibits a discrete drop following any incident. The more complicated belief evolution makes it impossible to provide an analytical expression of the release time $t_i^*\left(\pi_{\underline{T}}^i\right)$ as we provided in equation (4); in the extended model, the release time t_i^* will depend not only on the parole board's initial belief $\pi_{\underline{T}}^i$ about inmate i but also on the complete incident history of inmate i. Nonetheless, at whatever time t_i^* inmate i is released (if he is released at all between $\underline{T}^i$ and $\overline{T}^i$), it must satisfy

$$\pi_{t_i^*}^i = \pi_{r_i}^*,$$

where $\pi_{r_i}^*$ is characterized in equation (1). The effect of this generalization on Figure 1 is that the time paths for belief evolutions will stochastically exhibit discrete declines. However, the key feature for our test—that prisoners who are released between their minimum and maximum sentences are all released at the rehabilitation belief threshold π_r^*—remains valid.

REFERENCES

Abrams, David, Marianne Bertrand, and Sendhil Mullainathan. 2012. Do Judges Vary in Their Treatment of Race? *Journal of Legal Studies* 41:347–83.

Alesina, Alberto, and Eliana La Ferrara. 2014. A Test for Racial Bias in Capital Punishment. *American Economic Review* 104:3397–3433.

Antonovics, Kate L., and Brian G. Knight. 2009. A New Look at Racial Profiling: Evidence from the Boston Police Department. *Review of Economics and Statistics* 91:163–77.

Anwar, Shamena, Patrick Bayer, and Randi Hjalmarsson. 2012. The Impact of Jury Race in Criminal Trials. *Quarterly Journal of Economics* 127:1017–55.

Anwar, Shamena, and Hanming Fang. 2006. An Alternative Test of Racial Profiling in Motor Vehicle Searches: Theory and Evidence. *American Economic Review* 96:127–51.

———. 2012. Testing for the Role of Prejudice in Emergency Departments Using Bounceback Rates. *B.E. Journal of Economic Analysis and Policy (Advances)* 12(3), art. 4, pp. 1–47.

Ayres, Ian, and Joel Waldfogel. 1994. A Market Test for Race Discrimination in Bail Setting. *Stanford Law Review* 46:987–1047.

Bernhardt, Dan, Steve Mongrain, and Joanne Roberts. 2010. Rehabilitated or Not: An Informational Theory of Parole Decisions. *Journal of Law, Economics, and Organization* 28(2):186–210.

Bushway, Shawn D., and Jonah B. Gelbach. 2010. Testing for Racial Discrimination in Bail Setting Using Nonparametric Estimation of a Parametric Model. Unpublished manuscript. Yale Law School, New Haven, CT.

Carroll, Leo, and Margaret E. Mondrick. 1976. Racial Bias in the Decision to Grant Parole. *Law and Society Review* 11(1):93–107.

Goldkamp, John S., E. Rely Vîlcică, M. Kay Harris, and Doris Weiland. 2010. Parole and Public Safety in Pennsylvania: A Report to Governor Edward G. Rendell. Temple University, Department of Criminal Justice, Philadelphia.

Huebner, Beth M., and Timothy S. Bynum. 2006. An Analysis of Parole Decision Making Using a Sample of Sex Offenders: A Focal Concerns Perspective. *Criminology* 44:961–91.

Humes, Karen R., Nicholas A. Jones, and Roberto R. Ramirez. 2011. *Overview of Race and Hispanic Origin: 2010*. 2010 Census Briefs. Washington, DC: US Department of Commerce, Bureau of the Census. http://www.census.gov/prod/cen2010/briefs/c2010br-02.pdf.

Knowles, John, Nicola Persico, and Petra Todd. 2001. Racial Bias in Motor Vehicle Searches: Theory and Evidence. *Journal of Political Economy* 109:203–28.

Kuziemko, Ilyana. 2013. How Should Inmates Be Released from Prison? An Assessment of Parole versus Fixed Sentence Regimes. *Quarterly Journal of Economics* 128(1):371–424.

Mechoulan, Stephane, and Nicolas Sahuguet. 2015. Assessing Racial Disparities in Parole Release. *Journal of Legal Studies* 44:39–74.

Morgan, Kathryn D., and Brett Smith. 2008. The Impact of Race on Parole Decision-Making. *Justice Quarterly* 25(2):411–35.

Muhlhausen, David. 2004. The Determinants of Sentencing in Pennsylvania: Do the Characteristics of Judges Matter? Center for Data Analysis Report No. 04-02. Heritage Foundation, Washington, DC.

Persico, Nicola. 2009. Racial Profiling? Detecting Bias Using Statistical Evidence. *Annual Review of Economics* 1:229–54.

Petersilia, Joan. 1985. Racial Disparities in the Criminal Justice System: A Summary. *Crime and Delinquency* 3(1):15–34.

———. 2003. *When Prisoners Come Home: Parole and Prisoner Reentry*. New York: Oxford University Press.

Rehavi, M. Marit, and Sonja Starr. 2014. Racial Disparities in Federal Criminal Sentences. *Journal of Political Economy* 122:1320–54.

Steffensmeier, Darrell, and Stephen Demuth. 2001. Ethnicity and Judges' Sentencing Decisions: Hispanic-Black-White Comparisons. *Criminology* 39(1):145–78.

West, Heather C. 2010. Prison Inmates at Midyear 2009—Statistical Tables. Washington, DC: US Department of Justice, Bureau of Justice Statistics.