How Important is Labor Market Signaling? New Evidence from High School Exit Exams

Paco Martorell*
RAND

Damon Clark
Princeton University, University of Florida, NBER, IZA

PRELIMINARY AND INCOMPLETE: PLEASE DO NOT CITE WITHOUT AUTHORS' PERMISSION

Abstract

The most direct way to test whether employers use education as a signal of unobserved productivity is to test whether firms statistically discriminate in favor of more-educated workers. Previously-used versions of these tests suggest that they do. In particular, many studies have shown that workers who hold educational credentials such as a GED, high school diploma, or college degree earn more than uncredentialed workers that have completed the same amount of schooling. Yet if credentialed workers are more productive than uncredentialed workers, and if employers have productivity information that is unobserved by the researcher, these estimates will be biased. This paper presents new evidence on the role of education-based signaling using quasi-experimental variation in the likelihood of earning a high school diploma, the most commonly-held education credential in the U.S. Specifically, we use a regression discontinuity approach based on a comparison of individuals who barely pass and barely fail high school exit exams, tests that students must pass in order to earn a high school diploma. Using administrative data on education and earnings from the states of Texas and Florida, we estimate that the high school diploma premium is less than 5 percent, around one quarter of the size of previous estimates. We examine several possible threats to our conclusions, notably the possibility that receipt of a high school degree affects "downstream" outcomes such as postsecondary schooling, and conclude that these do not undermine our central conclusions. Our analysis suggests that the signaling role of education may be smaller than was previously thought.

*We thank David Lee and various seminar participants for helpful comments and suggestions.

1 Introduction

The concept of market signaling, first proposed by Spence (1973), has added greatly to our understanding of markets characterized by incomplete information. Applied to the labor market, the idea is that firms have incomplete information about worker productivity, hence base productivity expectations, hence wages, on alterable "signals" such as education and unalterable "indices" such as race and sex. The idea that education may signal labor market productivity has important education policy implications. In particular, since the signaling role of education could account for some of the private returns to education, signaling implies that these might exceed the social returns to education, the returns relevant to policy-makers. Yet the practical importance of education-based signaling is unclear. Many methods have been used to address this question, but conclusions vary across and often within the methods used.[1]

The most direct way to test whether education signals labor market productivity is to test whether firms statistically discriminate in favor of more-educated workers. That is because statistical discrimination given uncertainty is the key property of the signaling hypothesis. The difficulty with such tests, as Altonji and Blank (1999) note in relation to tests of race- and sex-based statistical discrimination, is that while "one would like to compare outcomes of individuals in the same job who are identical in all respects that are relevant to performance" (p.3192), these "equal productivity" comparisons are difficult to approximate with observational data. It would be almost impossible to make equal-productivity comparisons between workers that have spent different amounts of time in education and no such tests have been attempted. It may be possible to make equal-productivity comparisons between workers with and without certain educational

[1] These approaches include tests of the returns to education among self-employed workers (Wolpin (1977)), comparisons of the returns to education across occupations (Riley (1979)), tests of whether education policy changes impact the educational decisions of students they do not directly affect (as they would if these students wished to differentiate themselves from the directly-affected students) (Lang and Kropp (1986), Bedard (2001)), tests of whether some years of education generate higher returns than others (Layard and Pschacharopoulos (1974), Hungerford and Solon (1987)), estimates of the speed of employer learning (Altonji and Pierret (1997), Altonji and Pierret (2001)), Lange (2007)) and estimates of the difference between the private and social returns to education (Acemoglu and Angrist (2000)).

credentials and a large literature has asked whether, conditional on completed years of education, there is a wage premium paid to workers with credentials such as a GED, a high school diploma and a college degree. Nearly all of these studies suggest that there is, and these "sheepskin effects" are often cited as providing some of the clearest evidence in support of a signaling role for education. Yet as Riley (2001) notes in a critique of this literature, it is not clear why we would expect workers with credentials to be as productive as workers without them.

To address these problems, this paper provides quasi-experimental estimates of the wage premium associated with a particular education credential, a high school diploma. We begin with a theoretical analysis of the conditions under which the credential wage premium can be identified. These require that there is some uncertainty in the credential acquisition process and that the econometrician can harness this uncertainty to implement equal-productivity comparisons of workers with and without credentials. We then argue that high school exit exams provide a compelling natural experiment that meets them. High school exit exams are statewide tests of academic skills that students must pass in order to graduate high school and receive a high school diploma. Our estimates of the high school diploma wage premium are based on a comparison of the wages of workers that narrowly pass and narrowly fail these exams. We argue that these two sets of workers are likely identical in all respects relevant to labor market performance, and differ only in that the group that passes is more likely to have received a high school diploma.

We use this approach to estimate high school diploma premiums in Florida and Texas. Our dataset links high-quality administrative earnings data to detailed education records for large samples of students that were enrolled in Florida and Texas high schools. We observe these workers for up to seven years after they leave high school, roughly until age 25. Our estimates point to small diploma premiums. For example, our main estimates suggest that a high school diploma increases earnings by less than 5 percent and we can rule out premiums larger than 10 percent in either direction. An analysis of these premiums by sex and race suggests that these small effects apply to all of these subgroups. When we consider the experience profile of these estimated premiums we find no consistent pattern.

Although we think that our approach is highly credible, two caveats must be applied to our estimates. First, because students can retake exit exams, our estimates are based on the subsample of students that sit the last administration of the exam (in high school), a group for whom the exam outcome is found to exert a strong influence on diploma status. If diploma premiums are heterogeneous (e.g., because they vary with the other productivity signals observed by firms), our estimates will be specific to the students on the margin of passing this "last chance exam". As such, these estimates may not identify the high school diploma premiums enjoyed by other types of students (for whom the other productivity signals observed by firms may differ). They should however identify high school diploma premiums relevant to students affected by policies designed to increase high school graduation rates, since these are likely to affect similarly marginal groups. The second caveat is that passing the exit exam and receiving a high school diploma may affect downstream decisions such as whether to attend college. In that case, while our estimates would still identify the causal effect of obtaining a diploma, these estimates would be harder to interpret. Although we obviously cannot examine all possible downstream outcomes, we estimate diploma impacts on a variety of post-secondary education outcomes and show that these are unlikely to account for any of our results. In contrast, previous estimates of the high school diploma premium (Jaeger and Page (1996)) were made conditional on not pursuing postsecondary education, an outcome possibly endogenous to obtaining a diploma.

Our estimates are much smaller than previous estimates of the high school diploma premium, which range from 10 to 20 percent. The obvious explanation is that those estimates are biased upwards. Indeed, when we apply those methods to our data, we obtain similarly large effects. Another interesting comparison is between our estimates and quasi-experimental estimates of the GED premium, some of which suggest effects as high as 20 percent (Tyler, Murnane, and Willett (2000), Tyler (2004)) and some of which suggest no effects (Lofstrum and Tyler (2007)). One possible explanation for the difference between our estimates and these is that these use methods that rely on stronger identifying assumptions. As such, some of them may also be biased upwards. Alternatively, some educational credentials may send stronger signals than others hence may carry larger wage premiums than others.

We would not, in general, expect different educational credentials to command the same wage premium, and we would not expect the same educational credentials to command the same premium in different labor markets. As we show in our theoretical framework, the credential wage premium will depend, among other things, on the underlying productivities of workers with and without credentials and on how informative are the other productivity signals observed by firms. Nevertheless, we think it is significant that we find only a small premium associated with the most commonly held educational credential in the US, in two of the largest states in the US, and across all of the major demographic groups in those states. The strong correlation that we find between earnings and high school diploma status, even among students in the "last chance exam" sample, suggests that high school diplomas predict productivity. The absence of a high school diploma premium suggests that diplomas do not predict productivity conditional on the other information observed by firms (although the source of this other information remains unclear). This suggests that the signaling role of education may be less important than was previously thought.

2 Theoretical Framework and Related Literature

In this section we lay out a simple theoretical framework that we use to analyze the identification and interpretation of the diploma premium. We use this framework to show why previous estimates of the high school diploma premium may be biased upwards and why exit exams can help to identify this premium. The framework also aids interpretation of the estimated diploma premium.

2.1 Diploma Acquisition Under Certainty

We begin with a baseline model of diploma acquisition under certainty. The model assumes that individuals live for two periods. In the first period they attend school, in the second period they work. We assume individuals are perfect substitutes in production and we assume there are a large number of risk-neutral firms competing for their services. This implies that the second-period wage equals expected productivity. We characterize individuals in terms of a one-dimensional index of ability denoted a drawn from a distribution with density $f(a)$. We assume that ability is known to individuals at the start of the first period but is never observed by firms.

We assume that students graduate and earn a diploma when a one-dimensional measure of high school performance p exceeds some threshold level P. The one-dimensional nature of p is unrestrictive: this is a general formulation that encompasses the requirements that exist in the absence and presence of an exit exam. In the absence of exit exams, p could refer to the minimum of the number of credits accumulated and the number of days attended (assuming these two criteria determined whether a student graduated); in the presence of exit exams, p could refer to the minimum of the math, reading and writing scores. We assume that p is determined by the following equation:

$$p = \gamma_0 a + \gamma_1 s$$

where s is the amount of study effort exerted in the schooling (i.e., first) period. We think of study effort as encompassing various behaviors and activities that make high school graduation more likely given ability. In the absence of exit exams, studying could correspond to attendance; in the presence of exit exams, it could correspond to exam preparation. We assume that study effort is costly for all individuals and is, potentially, a function of ability. We can parameterize this dependence as:

$$C(s) = (c + \gamma_2 a)g(s); g(0) = 0, g'(s) > 0$$

so that $C(\overline{s}(a))$ represents the "study expenditures" required to obtain a diploma, where:

$$\overline{s} = \frac{P - \gamma_0 a}{\gamma_1} \text{ if } a \leq \frac{P}{\gamma_0}; \overline{s} = 0 \text{ if } a > \frac{P}{\gamma_0}$$

Let the indicator D denote whether a worker obtained the diploma:

$$D = 1 \text{ iff } p \geq P$$

We assume that second-period productivity (denoted π) is determined by a simple but general function of ability and first-period study effort:

$$\pi = a + \gamma_3 s$$

This captures the idea that productivity is related to underlying ability (since $\frac{d\pi}{da} = 1$). It also implies that the studying might be valued by firms (if $\gamma_3 > 0$).

The framework described by these equations encompass a number of special cases. For example, if $\gamma_0 = 0, \gamma_1 = 1, \gamma_2 < 0$ and $\gamma_3 = 0$ we have the classic signaling idea that students obtain a diploma by engaging in unproductive studying, the costs of which are decreasing in ability. If $\gamma_0 = \gamma_1 = 1, \gamma_2 = 0, \gamma_3 = 1$, then second-period productivity equals high school performance and diploma status is a direct signal of productivity. This captures the idea that exit exams can induce students to engage in productivity-enhancing studying that they would not exert in the absence of these exams (c.f., Costrell (1994), Betts (1998)). Indeed, this special case can be seen as a greatly simplified version of the Betts (1998) model.

There is no uncertainty in this model: students know exactly how much studying is required in order to graduate. In this case, it can be shown that the credential wage premium cannot be identified if firms observe more productivity information than the econometrician. To illustrate this result, we consider this model under the assumption that firms observe only diploma status and under the assumption that firms observe diploma status and a noisy signal of productivity. We show that once firms observe an additional signal of productivity, the credential wage premium can no longer be identified.

Firms Observe Only Diploma Status

Assuming that firms only observe D (i.e., whether a worker obtained a diploma), it follows that there can be at most two wage levels on the labor market: a wage for those with the diploma and a wage for those without the diploma. This implies that in this simple baseline model, the diploma premium is a constant equal to $\Delta W = W(D=1) - W(D=0)$.

Consider the study responses to a given diploma premium $\Delta W_0 > 0$. Assuming $\gamma_0 > 0$, the study level required to attain the diploma $\overline{s}(a)$ is decreasing in ability (over the ability range $[0, \frac{P}{\gamma_0}]$). Assuming $\gamma_2 \leq 0$, this implies that the cost of attaining the diploma is decreasing in ability. This implies that there will be an ability cutoff a_0 such that workers with ability $a < a_0$ will not study and will not obtain the diploma while workers with ability $a \geq a_0$ will engage in study $\overline{s}(a)$ and obtain the diploma. Study levels other than $s = 0$ and $s = \overline{s}$ will never be observed, since the wage depends solely on whether workers have a high school diploma.

Given these effort strategies, a Nash equilbrium in this game is an equilibrium credential wage premium ΔW_* that induces a cutoff $a_*(\Delta W_*)$ that ensures that this premium equals the expected productivity difference between those with and without the credential. In equilibrium, such a wage premium must satisfy:

$$\begin{aligned}
\Delta W_* &= E[\pi|D=1] - E[\pi|D=0] \\
&= \{E[a|a \geq a_*(\Delta W_*)] - E[a|a < a_*(\Delta W_*)]\} + \gamma_3 E[\overline{s}(a)|a \geq a_*(\Delta W_*)]
\end{aligned}$$

where the first term captures the difference in average ability between those with and without the diploma and the second term captures the average additional effort exerted by those that obtained the diploma. These equilbrium effort and wage levels are illustrated in Appendix Figure 1.

Firms Observe Diploma Status and A Noisy Signal of Productivity

In practice, firms are likely to observe other signals of productivity. Suppose that firms observe a noisy signal of productivity $\pi^s = \pi + \varepsilon$, where ε is a random variable assumed to be independent of π with mean zero and variance σ^2_ε. Since firms will predict productivity using both π^s and D, there will be two wage functions offered on the labor market: $W(D = 1, \pi^s)$ and $W(D = 0, \pi^s)$. Provided σ^2_ε is finite, π^s will provide some information about π, hence $\frac{d}{d\pi^s}W(D, \pi^s) > 0$ hence $\frac{d}{ds}E[W(D, s) > 0]$ provided $\gamma_3 > 0$. This implies that there will exist incentives to study even for students for whom high school performance is too low to obtain a diploma and for students that can obtain a diploma with zero study effort. Provided $\sigma^2_\varepsilon > 0$, π^s will not be completely predictive of π and D will provide additional productivity information. In this case, the diploma premium - the premium paid to workers with $D = 1$ conditional on π^s - will be positive. That is, $W(D = 1, \pi^s) > W(D = 0, \pi^s) = \Delta W(\pi^s) > 0$. This implies that the expected diploma premium for students with performance close to P is positive: $E[\Delta W(p = P)] = E[W(D = 1, p = P)] - E[W(D = 0, p = P] > 0$. In turn, this implies that for students with high school performance close to P, there is an extra return to obtaining the diploma. This implies that students will exert the additional study effort required to obtain it. This implies that for a given wage structure as characterized by $E[W(D, p)]$, study levels will jump up to $\overline{s}$ at some ability level a_2 below $\frac{P}{\gamma_1}$. *Equilibrium* study responses must imply productivity differences that ensure the wage functions generate zero profits. The study and wage functions that characterize this type of equilibrium are illustrated in Appendix Figure 2.

Interpretation

This model suggests that the diploma premium will reflect a number of factors. To analyze these, start with the case in which firms observe only diploma status. Even in this simple case, the diploma premium will depend on a number of factors. First, for a given ability distribution, it will depend on the cutoff P. Second, for a given cutoff P, it will depend on the relevant ability distribution. Third, provided $\gamma_3 > 0$, it will include a component that reflects the additional study effort exerted by those that received the diploma. The presence of this component implies that the diploma premium will depend on the extent to which study effort improves productivity. That is likely to reflect factors such as school quality. In addition, the role of study effort suggests that the premium will depend on the correlation between diploma requirements and productivity. Importantly, different requirements will be associated with different premia. This point was made by Dee and Jacob (2006) in their discussion of the effects of expanding high school graduation requirements to include high school exit exams.

The interpretation of the diploma premium is even more involved when firms observe other productivity signals, as they surely do in practice. Suppose, as above, that they observe a one-dimensional productivity signal π^s and assume that this variable is discrete, taking values $\pi_1^s, \pi_2^s, ..\pi_n^s$. In that case we can think of workers as belonging to one of n sub-markets. One obvious implication is that as firms acquire more detailed productivitiy information (i.e., can assign workers to more and narrower submarkets), we would expect the premium in any sub-market to become smaller. That is, as firms acquire other productivity information, diploma status becomes less predictive of productivity conditional on this other productivity information. A less obvious implication is that the premium will likely vary across submarkets. If productivity in every market was determined by the same function of ability and study effort, this variation will be driven by differences in how informative is D given π^s. For example, we would expect the premium to be highest in high-π^s markets, since in those markets knowledge of D will have the largest effects on firms beliefs about true productivity π. If productivity is market-specific, there might be other forces at work. For example, suppose that productivity depends on cognitive skills in the highest-π^s market but depends on perserverance in the lowest-π^s market. If the diploma requirements relate mainly to cognitive skills, we would expect the diploma premium to be larger in the high-π^s markets. We return to these issues below.

Identification

In both of these cases we have a "no overlap" result. That is, there is no pair of workers that have different values of D and the same value of $E[\pi^s]$ (c.f. Appendix Figures 1 and 2). In the second case in which firms have a second productivity signal, and in the general case in which firms observe more than the econometrician, this implies that we cannot identify the diploma premium when there is no uncertainty in the credential acquisition process. In the first case in which firms only observe D, we can identify the premium, by comparing the average wages of those with and without the diploma. More generally, which the econometrician observes as much as the firm, we can identify the premium by comparing average wages conditional on the other information available. Intuitively, firms cannot reward other characteristics associated with whether or not workers have the credential if no such characteristics exist. The assumption that the econometrician observes as much as the firm is an extremely strong one. For example, it is plausible to assume that firms can infer additional productivity information by simply observing workers (e.g., in a job interview).

2.2 Diploma Acquisition under Uncertainty

We now wish to contrast these cases with the case in which there is some uncertainty in the diploma acquisition process. There are at least three ways to introduce uncertainty into this process. First, students may be imperfectly informed about ability a. Then, even if high school performance is a deterministic function of ability and study effort, students will not be able to forecast performance precisely. Second, students might not know how productive will be their study effort (i.e., γ_1 could be a random variable whose distribution varied across students). Third, performance itself could be stochastic - a function of ability, study and a random component, where the random component could be driven by unanticipated factors such as the quality of a student's teacher.

To take the simplest of these cases (the other cases have the same implications), suppose that high school performance is stochastic, such that students obtain a diploma if $p = \widehat{p} + \eta > P$, where

$\widehat{p} = \gamma_0 a + \gamma_1 s$ and η is a random variable assumed to be independent of $\widehat{p}$ with mean zero and variance σ_η^2. Provided σ_η^2 is sufficiently high, all students will obtain a diploma with probability between zero and one. To see how this affects study choices, return to the case in which the diploma is a firm's only productivity signal and assume the diploma premium is ΔW_0. In that case, the expected return to the non-stochastic component of performance is $B(\widehat{p}) = P(D = 1|\widehat{p})\Delta W_0$. This implies that the marginal return to study effort is $\gamma_1 \frac{dP}{d\widehat{p}} \Delta W_0 = -\gamma_1 f_\eta(P - \widehat{p})\Delta W_0$, where $-f_\eta()$ is the impact of an additional unit of effort on the probability of obtaining the diploma. If η is normally distributed, then for an individual with ability a, $-f_\eta()$ will be an inverted U shape with a maximum at $s = \frac{P - \gamma_0 a}{\gamma_1}$ (i.e., $\widehat{p} = P$ and $P(D = 1) = 0.5$). Intuitively, since small shocks are more frequent than large shocks, additional effort will have the biggest impact on the probability of obtaining a diploma when performance is most sensitive to small shocks (i.e., when $\widehat{p} = P$). This implies that if study costs are unrelated to ability ($\gamma_2 = 0$), the individual with ability a_3 will exert most effort, where $C'(\widehat{p}(a_3)) = C'(P) = \gamma_1 f(0)\Delta W_0$. All others exert less effort, but effort will be a smooth function of ability. Because uncertainty affects study choices, it will also affect firms' productivity inferences. Again, an *equilibrium* premium is one that induces effort choices such that the premium equals the difference in expected productivity between workers with and without the credential. This type of equilibrium will feature effort choices and wage levels similar to those depicted in Appendix Figure 3.

Interpretation

This type of uncertainty complicates the firms' calculation of expected productivity conditional on high school diploma status and the other productivity signals that it observes. In the extreme case in which the diploma outcome is completely random, diploma status would be uninformative. Provided that diploma status is determined mainly by factors correlated with productivity (i.e., ability and study effort), diploma status will still predict productivity and all of the points previously made about interpretation will continue to hold.

Identification

The important implications of uncertainty relate to identification. In partiuclar, under uncertainty (Appendix Figure 3), there now exist two types of overlap. First, we can find workers who differ in terms of diploma status but have the same level of ability (hence study effort and expected productivity). Assuming that the econometrician does not observe ability a, this does not help to identify the diploma premium. Second, we can find workers who differ in terms of diploma status but have (approximately) the same level of performance and (approximately) the same level of expected productivity. Since high school performance p is potentially observable, this overlap can be used as the basis for identifying the diploma premium. In particular, the diploma premium can be identified via a threshold comparison of those with performance slightly below P (who do not obtain a diploma) and those with performance slightly above P (who do). Notice that this overlap exists regardless of the productivity information held by firms.

2.3 Relation to the Previous Literature

The diploma premium can be identified when the diploma acquisition process is characterized by uncertainty and when the econometrician can observe the factors determining diploma status. When states operate high school exit exams, we show that these assumptions are satisfied when exit exam scores cannot be forecast by students and when these scores can be observed by the econometrician.

Previous estimates of the high school diploma premium were obtained from settings in which traditional U.S. high school graduation requirements applied. Typically, these involve the accumulation of a certain number of credits, where these credits depend on passing courses, and where the passing criteria for these courses are course- and school-specific. Some courses might, for example, require only that students attend class. Others might require that students complete course papers and submit homework. These requirements ensure there is little uncertainty in the diploma acquisition process. Students are generally aware of the graduation requirements and the course requirements on which they are based. Even if a student inadvertently fails to meet one of these course requirements (e.g., forgets to complete a course paper), she can typically rectify this by, for example, completing a course paper later in the course.

The absence of uncertainty implies that students with and without high school diplomas are unlikely to be equally productive.[2] As such, the validity of these previous estimates rests on the econometrician observing all of the productivity information observed by firms. Unfortunately, these estimates are based on datasets (notably the Current Population Survey) that contain very little information beyond basic demographics, years of education completed and whether or not a diploma was received. For this reason we suspect that the diploma premia estimated in this literature are biased upwards. Below, we apply the methods used in the previous literature to our own data. We obtain estimates comparable to those obtained in the previous literature and much larger than those generated by our preferred approach.

Previous estimates of the GED premium were obtained from settings in which the credential acquisition process was characterized by uncertainty. Specifically, as with high school diplomas and high school exit exams, workers wishing to acquire a GED must first pass a series of GED examinations. Three studies have used GED exam scores to estimate the GED premium, although none have employed the regression discontinuity approach used here. Using data from a single state (Florida), Tyler (2004) estimates the GED premium using worker fixed effects and various types of regression adjustment. He finds GED effects of between 10 and 20 percent. Tyler, Murnane, and Willett (2000), and Lofstrum and Tyler (2007) estimate the GED premium by using difference-in-difference methods to exploit variation in GED passing standards. Both papers compare groups of workers that have the same GED exam scores but who do not all have the GED. In the first paper that is because some groups live in states with higher GED passing standards. In the second paper it is because some groups took the exam after the standard had been raised. Both papers

[2] This is the essence of the Riley (2001) critique of these estimates: it is hard to conceive of reasons why these two sets of workers will be equally productive.

adjust this difference by the difference in earnings among workers whose scores met the standard under both regimes. The first study suggests that the GED increases earnings by between 11% and 20%. The second study suggests that the GED has no effect on earnings. Lofstrum and Tyler (2007) conclude that the most plausible explanation for these differences is that the low-premium state (Texas) has especially low passing standards.

3 High School Exit Exams

In the previous section we argued that exit exams can be used to provide credible estimates of the high school diploma premium. In this section we discuss these exams in more detail. We begin with a stylized description of these exams that illustrates the idea behind the strategy that we use to identify the diploma premium. In the second part of this section we discuss the practical operation of these exams. We discuss how our strategy can take account of various features of these exams and discuss how these affect the interpretation of our estimates.

3.1 Exit Exams: A Stylized Description

Exit exams were introduced to combat the idea that high school diplomas were awarded for "time-serving". They do so by making diploma receipt conditional on passing statewide exams of basic skills. To see how these exams facilitate identification of the diploma premium, and to see how this premium should be interpreted, consider a stylized version of these exams in which each student that remains in school until the end of twelfth grade takes the exam. Those that pass receive a high school diploma, those that fail do not.

In such a setting, the requirements for identification of the high school diploma premium would be met. First, these exams would introduce an element of uncertainty into the diploma acquisition process. In particular, it is reasonable to suppose that students cannot forecast their performance on this exam, hence reasonable to suppose that by chance, there exist workers with and without these diplomas that are, to firms, observably equivalent. Second, performance on this exam is potentially observed by the econometrician. In that case, the diploma premium can be identified via an estimate of the earnings discontinuity associated with narrowly passing the exam. The intuition is that workers with scores either side of the passing threshold should be observably equivalent to firms *irrespective of any other productivity signals that they observe.* Any earnings difference between these groups must reflect the signalling value of the high school diploma.

The interpretation of the estimated premium depends what else firms observe. If diploma status is the only productivity signal observed by firms, the diploma premium will reflect the difference in expected productivity between those that obtain the diploma and those that do not obtain the diploma, where this difference consists of the between-group difference in underlying ability and any differences in expected study effort between the two groups. If firms observe additional productivity signals, such that workers can be placed in particular sub-markets, then the estimated premium can be thought of as a weighted average of submarket-specific premia (Lee (2008)). As noted above, the correlation between the submarket-specific productivity signal and the submarket-specific diploma premium is unclear.

3.2 Exit Exams in Practice

There are several differences between this stylized description and the practical operation of these exams. In what follows we focus on the operation of these exams in Texas, since most of the results in this version of the paper relate to Texas. Although exit exams operate slightly differently in all of the states that employ them, the key features are the same.

Multiple test subjects

In practice, students must pass exams in several subjects. For example, in Texas, since the exit exam was reformed in 1990, it has consistent of three sections (reading, math and writing), each of which must be passed before a diploma is received. This is of little consequence for identification of the diploma premium, since we can normalize exam scores in relation to the relevant passing thresholds and define a minimum normalized score that determines whether or not a student passes all three sections.

Multiple opportunities to take the exam

A more important consideration is that in practice, students have multiple opportunities to take the exit exam. In Texas, the number of retake opportunities changed over our study period. Starting in 1993 students first took the exit exam in spring 10th grade and prior to that, students first took the test in fall of 11th grade. Throughout our study period, the test was administered three times a year (fall, spring and summer) and starting in 1994, there was a final retest opportunity offered in late spring to 12th graders who had not yet passed the test. Thus, students who first took the test in fall 1991 had 5 chances to retake the test prior to the end of 12th grade (assuming no grade repetition) while 10th graders in 1993 and later had up to 8 chances. To consider the implications of these multiple opportunities, suppose the exam is offered L times over the course of a student's high school career, where the $L'th$ test is the "last-chance" test. Suppose a student that passes on the $i'th$ administration of the test receives a diploma, while a student that fails on the $i'th$ administration can retake on the $i+1'th$ administration (provided $i \leq L-1$). These retake opportunities have implications for both the interpretation and identification of the credential wage premium.

With regard to identification, assuming firms cannot observe the number of times the test was taken, a comparison based on the last-chance sample will still identify the diploma premium.[3] With regard to interpretation, multiple test-taking and a focus on the last-chance sample has two

[3] Without this assumption, a threshold comparison can still be used, but the estimated premium will be specific to this (observable) group of students. Since this group would, presumably, be less heterogeneous than the wider group of high school students, we might expect this credential wage premium to be small. We are reasonably confident that firms cannot observe this information and, in the discussion and interpretation section, we provide evidence in support of this claim. In particular, we estimate the wage impact of narrowly failing the first test. We show that this has no impact on the probability of graduating high school but does impact the number of times the test is taken. If firms could observe this, they would use it as a signal of productivity and pay workers that passed at the first attempt more than workers that took two or more attempts. Since we find no wage effect of passing the first test, we find no evidence of this type of behavior.

implications. First, assuming that firms observe additional productivity signals, then a focus on the last-chance sample implies that we will be identifying the diploma premium for a particular subsample of workers, those that failed the first $L-1$ administrations of the test. Second, even if firms do not observe additional productivity signals, multiple testing will change the equilbrium wage premium because it will change the incentives to exert effort and, therefore, the inferences made by firms given the productivity signals they observe.[4] While these effort strategies complicate the calculation of expected productivity given credential information, the equilibrium with multiple testing will be qualitatively similar to the equilibrium with a single test. In particular, it will consist of a credential premium that induces effort choices that ensure this premium equals the expected productivity difference between those with and without the credential. The intuition extends to the cases in which firms observe additional productivity signals.

Imperfect compliance

In practice, there is some slippage between passing the exam and obtaining a diploma, even among students in the last-chance sample. Some of this "imperfect compliance" is the result of students failing the last-chance exam but being exempt from the passing requirement. Some is due to to students re-taking and passing the exam later, typically by returning to high school for a "thirteenth" year. Some is due to students that pass the last-chance exam being denied a diploma because they have not fulfilled the other requirements for high school graduation. Again, this type of imperfect compliance has implications for the identifcation and interpretation and identification of the diploma premium.

With regard to identification, imperfect compliance can be handled via a modified version of

[4] For example, suppose educational credentials are the only productivity signal that firms observe. Then, for a given credential wage premium, an individual's effort strategy will consist of an optimal effort level e_1 for the first exam, an optimal effort level e_2 for the second exam (assuming the first exam was failed), up to e_L for the last exam. If the test outcome is certain, students do not gain from additional exams, since they only pass if they exert effort $e \geq T-a$ on at least one exam. With uncertain exam outcomes, students can gain from the possibility that they may, by chance, pass an early exam with minimal effort. As such, we would expect effort choices to satisfy $e_1 < e_2 < ..e_L$. Depending on the type of uncertainty (e.g., over ability), workers effort decisions may depend on the score obtained on the previous tests.

the threshold comparison on the last chance sample. In particular, we can scale up the earnings difference at the last chance passing threshold with the difference in the probability of obtaining a diploma as observed at the last chance passing threshold. For example, if the probability of obtaining a diploma is one for those that pass the last-chance exam and one half for those that fail (because some pass on a later attempt and some are exempt), we simply double the estimated last-chance wage premium to get a consistent estimate of obtaining the diploma. The intuition is that the threshold comparison is still a good one (the individuals either side should have equal levels of expected productivity), but the imperfect compliance, unless corrected, causes us to misclassify students and underestimate the credential wage premium.

With regard to interpretation, imperfect compliance again implies that we are identifying the wage premium for a specific subgroup: the "compliers" for whom the last-chance exam outcome determines whether or not they receive the diploma (Imbens and Angrist (1994)). We might also expect exemptions to affect the equilibrium diploma premium although in practice only a small number of exemptions are granted.[5] The same can be said for the practice of denying diplomas to students that pass the exit exam but do not complete other requirements. The existence of future retake opportunities has no implications beyond those discussed in the context of multiple testing.

High school dropout and further education enrollment

The stylized description assumed that students complete high school, take the exam then work. In practice, regardless of high school graduation requirements, students can drop out of education before the end of high school or continue their education beyond it.

The assumption of no dropout is easy to relax. If firms can observe years of schooling, hence

[5] To see this, recall the equilibrium in the model in which high school performance is certain and firms observe only diploma status (Appendix Figure 1) and assume that a fraction of those that fail the test are re-classified as exempt and given a credential. This will reduce the difference in expected productivity between those with and without a credential which must reduce the diploma premium. This will reduce optimal study levels which will further reduce expected productivity differences and the credential wage premium. This process stops when the study choices induced by the credential premium ensure this premium equals the expected productivity difference between those with and without the credential.

can observe which students drop out, the main consequence will be to reduce the diploma premium by effectively truncating the left tail of the productivity distribution. The threshold approach will continue to provide consistent estimates of this premium, since any other productivity signals observed by firms will be the same (on average) for the group that narrowly passes and the group that narrowly fails. If firms cannot observe years of schooling, hence cannot observe which students drop out, the high school diploma signals exam performance and high school completion. Proper consideration of this case would require an extended analysis that modeled student dropout behavior. In turn, this would require assumptions about the costs and productivity effects of additional time spent in school. Without specifying such a model, it seems reasonable to expect that these considerations would generate larger credential wage premia and that these will be consistently estimated by threshold comparisons.[6] We assume that firms can observe years of schooling hence can observe which students drop out. This assumption was also made in the previous literature.

The assumption of no further education is more problematic. With regard to identication, strategies that ignore further education could generate biased estimates of the diploma premium. Suppose for example that students with diplomas are more likely to enroll in college. Then ignoring college enrollment will cause us to load some of the returns to college onto the diploma premium. Strategies that deal with further education by excluding students that enroll in further education could also generate biased estimates of the diploma premium. To see this, suppose that in addition to observing worker's educational credentials, firms receive a second productivity signal that is either good or bad with probability related to productivity. The groups on either side of the exit exam threshold have similar productivity hence will contain similar proportions of good-signal workers. If, however, students with diplomas are more likely to attend college if they receive a bad

[6] The intuition for larger premia is that students that did not complete high school would choose to drop out at the earliest opportunity. Assuming human capital increases with time in high school (even if only slightly), this will increase the expected productivity difference between those with and without the credential. The intuition for the consistency of the threshold comparison is the same as before: any other productivity signals observed by firms will be the same (on average) for the groups on either side of the threshold.

signal, then among the subsample of threshold students that do not go to college, the students that receive the diploma will contain a larger fraction of good-signal workers. This will cause us to over-estimate the diploma premium. If students without diplomas are more likely to pursue a GED if they receive a bad signal, then students that do not receive a diploma will contain a larger fraction of good-signal workers. This could cause us to under-estimate the diploma premium. In the next section we describe how we deal with diploma effects on these "downstream" outcomes.

4 Empirical Framework

In this section we apply the insights of our theoretical framework and our discussion of the practical operation of these exams to a more general model of wage determination. This forms the basis of our empirical strategy. Without loss of generality, we assume that firms observe two productivity signals – whether workers hold a high school diploma (D) and a noisy measure of productivity (π^s). As discussed above, the weights placed on these signals will depend on the accuracy of the productivity signal. We also assume that firms observe years of completed schooling, such that the following equation can be considered as a linear approximation to the wages of those with at least 12 years of schooling:

$$W_i = \beta_0 + \beta_1 D_i + \beta_2 \pi_i^s \tag{1}$$

The high school premium is β_1, the difference between the expected productivity of a worker with a high school diploma and a worker without a high school diploma conditional on the other information observed by firms. For ease of exposition, we are assuming that this premium is the same for all workers. As noted above, a less restrictive assumption is that it is heterogenous, a function of π_i^s.

Previous estimates of the high school diploma premium have been based on regressions of earnings on diploma status among workers that have completed twelve years of school. To see what this regression identifies, we can project π^s onto D and substitute into equation (1):

$$W_i = \beta_0' + (\beta_1 + \beta_3 r_1)D_i + \varepsilon_i \tag{2}$$

where r_1 is the projection coefficient on D and ε is the projection error. If $\beta_2 > 0$ (i.e., firms observe other productivity signals) and if these are positively correlated with D (i.e., these are, on average, better among workers with a diploma), this procedure will generate estimates of β_1 that are biased upwards. The theoretical framework considered above showed that unless the diploma acquisition process is characterized by uncertainty, both conditions are likely to hold.

In contrast, we seek to identify the effects of D by comparing workers with exit exam scores close to the passing threshold. Intuitively, among this group, we expect that π_i^s will be approximately constant. More precisely, we estimate the diploma premium using data away from the passing threshold and assuming that π_i^s is a smooth function of the exit exam score p. With complete compliance, this would imply that in the projection of π^s onto D and $g(p)$, the projection coefficient on D was zero. This would allow us to rewrite (1) as:

$$W_i = \beta_0 + \beta_1 D_i + g(p_i) + \varepsilon_i \tag{3}$$

where ε is again the projection error. With imperfect compliance, this no longer follows, since among those that narrowly fail the exam, there may be differences in the productivity signals associated with those that subsequently obtain (and do not obtain) a diploma. It does however imply that in the projection of π^s onto $PASS$ and $g(p)$, the projection coefficient on $PASS$ is zero, where this is a dummy variable for scoring to the right of the threshold. Hence rather than estimate equation (3), we will instead use passing the exam as the excluded instrument in a two-stage least squares procedure. That is, we will instrument for D in the "outcome equation" (3) using the variable $PASS$ in the following "first-stage" equation:

$$D_i = \alpha_0 + \alpha_1 PASS_i + g(p_i) + \omega_i \tag{4}$$

4.1 Estimation Issues

As shown above, the regression discontinuity approach can identify the diploma wage premium with miminal assumptions regarding the other productivity signals observed by firms. In implementing this regression discontinuity approach however, various choices must be made. In this version of the paper we take a fairly standard approach to these issues. In particular, we use a "global polynomial" approach that uses a wide range of data and allows $g(.)$ to be a high-order polynomial (e.g., a fourth-order polynomial). We check the validity of our functional form assumptions by visual comparison of the fitted functional form and the raw data. We also experiment with adding covariates to equations (3)-(4) in order to improve this fit.

5 Data

The data used for the Texas portion of the analysis comes from the Texas Schools Microdata Panel (TSMP) which is maintained by the Texas Schools Project. The TSMP is a collection of administrative databases from various Texas state agencies all of which use a common identifier which makes it possible to link records across files. With these data, a longitudinal database can be created that tracks students through high school and later into the labor force or college. We draw our analysis dataset from high school records from the Texas Education Agency (TEA) that contain detailed information on enrollment, attendance, test scores and high school completion. Specifically, we analyze five cohorts of test-takers: those who first tested in fall of 11th grade in 1991 or 1992, and students who first took the test in spring of 10th grade in 1993-1995 The TEA files have scores from all test attempts so it is possible to determine a student's test taking history.

We focus on students who had not yet passed the exit exam prior to the final administration given in their senior year which we refer to as the "last chance" test. Despite our labelling these as the "last chance" test, it is important to note that students could have re-taken the test at a later date. Nonetheless, the evidence presented below clearly indicates that for many students, failing this test administration dramatically lowers the likelihood of graduating, so in this sense, it effectively many students' "last chance" to pass the test.

A central strength of the TSMP for this study is the breadth of student outcomes it covers. The first of these is whether a student graduates from high school. Our primary measure of graduation is whether we are able to find a record in the roster of high school graduates indicating that a student received a high school degree within two years of taking the "last chance" test. To help understand the dynamics involved in how the exit exam affects graduation, we also look at graduation within shorter and longer time horizons. In addition to high school graduation, we also examine whether a student attends school in the year following the last chance test (i.e., a "thirteenth year") and acquisition of a General Educational Development (GED) degree. Enrollment information comes from the TEA's roster of students attending public high schools in Texas. Information on the GED also comes from the TEA which maintains a database of Texas high school students who received a GED degree between 1992 and 2002 and all students who took the GED qualifying test (regardless of whether the credential was received) between 1995 and 2002. With these data we construct two measures - receipt of a GED degree within 5 years of the last chance test, and taking the GED test within 5 years of the last chance test.

The data on earnings comes from the Unemployment Insurance (UI) tax reports submitted to the Texas Workforce Commission by employers subject to the state's UI law. Subject employers are required to report, on a quarterly basis, the wages paid to each employee in order to determine the firm's tax liability. Data are available through 2004Q3, which permits a follow-up of at least seven years following the last chance test for all students in our study.

The use of administrative earnings data presents advantages and drawbacks relative to more commonly used survey data such as the CPS or NLSY. A key strength of administrative data is their accuracy, which constrasts will with survey data (Jacobson, Lalonde, and Sullivan, 1993). The main limitation is our inability to construct an hourly wage measure. Instead, we must estimate effects on total earnings. This will conflate diploma effects on wages with diploma effects on labor supply.

Although this may at first glance seem like a disadvantage, there are at least two reasons why it need not be. First, since we would like to compare these estimates to those in the related literature, we note that previous estimates of the GED premium were based on similar types of data. Although previous estimates of the high school diploma premium were not, we can use our data to compare discontinuity-based estimates of the diploma premium with estimates based on the approach taken in this literature. Second, since firms may base hiring decisions on productivity signals such as a high school diploma, the labor supply channel is one of independent interest. Indeed, in the related literature on race- and sex-based discrimination, hiring has typically been the focus of the audit studies and quasi-experimental approaches.

5.1 The Last Chance Sample

Since our analysis is based on the last chance sample it is important to understand how this is constructed. Figure 1 is an important first step to that end. This shows the probability of graduating high school as a function of the test obtained at each administration of the high school exit exam, starting with the first administration in grade 10 and finishing with the last chance sample at the end of grade twelve.

The first graph is based on data for all five cohorts and for all students that took the initial test (nearly all of them). The x-axis is defined as the minimum score on the three subtests, such that students with a minimum greater than zero passed at the first attempt and students with a minimum less than zero failed at least one component of the exam. Not surprisingly, there is a positive relationship between the score on this exam and the probability of graduating. Students that passed the exam graduate with probability at least 0.8: those that do not graduate will include those that dropped out before the end of grade twelve and those that completed grade twelve but failed to meet other graduation requirements. Those that narrowly failed the first test graduate with probability around 0.8. This reflects the fact that nearly all of these students retook the exam and passed on a subsequent attempt.

The next graph is based on data for the students in all five cohorts that took the exam at least twice. This will consist of a large fraction of those students that scored less than zero on the first test. The x-axis in this case is the minimum of the score obtained the subtests that had still not been passed. Again, a minimum greater than zero implies the exam was passed on this second attempt, a minimum less than zero implies that at least one subtest was failed. Again, there is a positive relationship between this minimum and the probability of ever graduating. Again, there is no discontinuity at the passing threshold, a reflection of the fact that students that fail still have multiple opportunities to retake and pass.

Only when we get to the penultimate test (the lower left-hand panel) and the final "last chance" test (the lower right-hand panel) do we observe a large discontinuity in the probability of obtaining a diploma. This last chance sample will, for the most part, contain students that have already failed the test seven times. Again, the score represented on the x-axis is the minimum of the (rescaled) score obtained on the unpassed sections. Around 90% of the students with positive scores (i.e., who passed the last chance exam) graduate. Just over 40% of the students who fail the last chance test graduate, generating a discontinuity of around 45% in the probability of graduation as a function of the last chance test score.

We return to this discontinuity later. First, we present some stylized facts associated with the last chance sample. Since the students in the last chance sample have failed the exam seven times, we should not be surprised to see that their initial scores are towards the bottom of the *initial test score* distribution. This is illustrated in Figure 2a, which plots the density of first scores obtained by the last chance sample against the density of first scores obtained by the full sample. The mean of the first score among the last chance sample is at the 11th percentile of the full sample distribution. Note that the full sample distribution is shaped by ceiling effects. That is, a lot of students score close to the maximum when they take the test for the first time.

In Figure 2b we focus on the last chance sample and plot the distribution of *last chance test scores*. In our theoretical framework we argued that exit exams ensure that high school performance is uncertain. In turn, this ensures that there will be students with scores just below and just above the last chance passing cutoff. The distribution seen in Figure 2b is consistent with this prediction. A formal test establishes that there is no discontinuity in the density of the last chance scores around the passing threshold (McCrary (2008)).

Table 1 presents some stylized facts for the last chance sample and presents more tests of the validity of the identification assumptions that we use to identify the credential wage premium. Recall that the key assumption is that other productivity signals observed by firms are smooth through the last chance passing threshold. We do not observe all of the productivity signals observed by firms, but we can assess whether the characteristics that we observe are smooth through this passing threshold.

The characteristics that we observe are listed in the rows of Table 1. The columns present the means of these variables among the last chance sample, the number of observations associated with the last chance sample (around 40,000) and the estimated discontinuity in these variables through the passing threshold. These estimated discontinuities are obtained by regressing these characteristics on a dummy variable for passing the last chance exam and a fourth-order polynomial in the last chance test score. The estimates are small and for the most part statistically insignificant. Where they are statistically significant, the associated graphs (not shown) suggest that these findings are unlikely to be robust to minor changes in the regression discontinuity specification.

6 Results

In this section we report three sets of estimates. First, we report estimates of the impact of passing the last chance exam on the probability of obtaining a high school diploma. Second, we use these estimates to generate estimates of the high school diploma premium. Third, we report the results of various robustness checks.

6.1 Estimates of the Diploma Effects of Passing the Last-Chance Exam

We have already seen that passing the last chance exam is associated with a roughly 45 percent increase in the probability of obtaining a high school diploma (Figure 1). We now investigate this relationship in more detail. Table 2 reports estimates of the diploma effects of passing the last-chance exam at various points in time relative to the last chance test. One semester after the last chance test the effect (i.e., discontinuity) is around 0.45. At longer intervals the effect gets smaller, reaching a minimum of 0.38 two years after the last chance exam.

It is not surprising that the estimated effect decreases with time. This reflects the increased opportunities that students that fail the last chance test have to retake and pass. Effectively, as more time elapses after the last chance exam, the probability of earning a diploma conditional on failing increases.

Passing at future retakes is not the only phenomenon at work here. As suggested by the relative flatness of the diploma-score relationship to the left of the threshold, many students obtain a diploma without ever passing the test, presumably because they are exempted from the graduation requirements. This can be seen in Figure 3, which plots separately the two routes that students can take to a diploma: exemption and passing. Note that while the fraction of students that obtain a diploma via exemption is large among the last chance sanple, it is small among the overall sample of students that pass the exam. Unless firms can observe at which point students passed the exam (which we assume they cannot), these exemptions should have little effect on the size of the credential wage premium.

6.2 Estimates of the Earnings Effects of Obtaining a Diploma

With these "first stage" estimates in hand, we can use the last-chance sample to estimate the earnings effects of obtaining a diploma. Figure 4 graphs the relationship between earnings and last-chance exam scores one, three, five and seven years after the last-chance exam. Three aspects of this relationship are apparent from this graph. First, average earnings and earnings conditional on any last-chance exam score increase with years since the last-chance exam. This is not surprising, since earnings are known to increase with labor market experience. Second, there is a positive correlation between last-chance exam scores on earnings and this strengthens with strengthens with years of experience. This could reflect a process by which it takes time for individuals to settle into the labor market (Topel and Ward, 2002). Third, despite this positive correaltyion, there is no apparent jump in earnings at the exit exam passing threshold.

Figure 5 presents these graphs in a way that allows for a clearer look at possible discontinuities. Only in the top-right panel, for earnings three years after the exam, is there any evidence of a positive jump, and no such jump is seen in the other panels of the graph. Table 3 reports the estimates associated with these graphs, and the estimates associated with the earnings effects of passing the exam two, four and six years after the exam. The first column presents mean earnings among the group that narrowly fails. The third and fourth columns report "reduced-form" regression discontinuity estimates based on a fourth-order polynomial in last chance scores. The third column presents estimates that do not condition on other observables (hence these correspond to the jumps seen in the graphs), the fourth column presents estimates that do condition on these. As we would expect, conditioning on these observables has almost no impact on the estimates. The estimates are also robust to changing the order of the polynomial. As seen in columns (5) and (6), a second-order polynomial generates similar estimates. In future versions of the paper we will also experiment with local linear approaches to the estimation of these effects.

These reduced-form estimates range from -$200 to $200. This is less than 2 percent of the relevant mean earnings. They show no obvious pattern by subgroup (not reported) and do not appear to vary over the seven year window. The instrumental variables estimates (reported in the right-hand panel of the Table) correspond to these reduced-form estimates scaled up by the estimated effect of passing the last chance exam on the probability of receiving a diploma. They therefore follow the same pattern but are roughly twice as large as the reduced-form estimates, ranging from -$500 to $500, around 5 percent either side of zero.

6.3 Robustness Checks

Our estimates suggest that the high school diploma premium is small. Before discussing how these estimates should be interpreted, we check they are robust to various sources of bias.

Further education biases

If diploma receipt had large effects on further education outcomes, our estimates of the diploma premium would still be unbiased, but these would be more difficult to interpret. In particular, they could not be interpreted through the lens of statistical discrimination and signaling. To assess this, we estimated the effects of passing the exam on the probability of pursuing various types of further education. These estimates are reported in Table 4. They are again based on models that include a fourth-order polynomial in the last chance test scores. To save space, estimates are presented for just one model: that including a fourth-order polynomial in the last-chance test score and other pre-determined characteristics.

The columns to the left of the vertical line refer to outcomes for which we expect to find negative effects of passing the last chance test. For example, in the first column we look at the probability of being enrolled in high school one year after the last chance exam. Since students will, typically, only return to high school if they need to retake the exit exam, we would expect to find a negative effect. In fact the high school effect is negative, but small and short-lived. The effect, around five percentage points, is only observed for one year, consistent with the pattern of first stage estimates presented in Table 2. We would also expect to find negative effects of passing the exam on the probability of attempting the GED and earning the GED. This is also consistent with our estimates, although the effect on earning a GED is a relatively small one, around five percentage points.

The columns to the right of the vertical line refer to outcomes for which we expect to find positive effects of passing the exam. That is in part because enrollment in some college courses is, formally, conditional on possessing a high school diploma. Even if this requirement can be waived, the pursuit of these options may be costly for those without a diploma, since they may have to satisfy other requirements (e.g., enroll in college remedial classes). The college enrollment estimates take the expected positive sign but are small in magnitude. It seems that a lot of students that narrowly pass the last-chance exam enroll in college for one semester but drop out shortly afterwards. While this may at first glance seem surprising, it is worth recalling from Figure 2 that the last chance sample is one with low academic ability, at least as measured by the initial exit exam score. As seen in the final columns of this Table, this difference in college attendance generates a tiny effect on college course credit accumulation and no effect on whether students earn a college degree.

To summarize, the main effect of passing the last chance sample is to reduce the probability that students pursue a GED. If pursuing - and earning - a GED is associated with a large increase in earnings, this could explain why our diploma earnings effects are smaller than those found in the preceding literature. To shed light on this possibility we estimated the impact of obtaining a GED among the subsample of those that fail the last chance test. These estimates, not reported, suggest that the GED premium is at most around 10% of the earnings of the marginal fail group. A rough calculation suggests that even these effects are far too small to explain more than a tiny part of the difference. For example, a 5% GED effect of passing multiplied by a 10% GED effect on earnings corresponds to an earnings effect of passing the exam of less than one percentage point.

Conditional-on-working biases

To this point we have used total earnings to measure the earnings premium. This has involved assigning zero earnings to individuals not observed in the earnings data. As noted above, this approach has been taken by the GED literature, although the earlier estimates of the high school diploma premiums were based on subsamples of workers in employment (e.g., Jaeger and Page, 1996).

For the purposes of identifying the diploma premium, the ideal scenario would be one in which all individuals work. In that case, we could estimate the diploma wage premium without worrying about selection into the labor force. Since around 20% of the individuals in our sample do not work (i.e., are not observed in the earnings data), an alternative approach must be taken. As noted above, the approach taken to this point has been to use the full sample of individuals to estimate diploma effects on earnings, with zero earnings assigned to non-workers.

The main advantage of estimating effects on total earnings is that these can be given a causal interpretation even when there are diploma effects on the probability of having positive earnings. A second advantage is that the total earnings effect combines the diploma effect on conditional-on-positive earnings and the diploma effect on the probability of having positive earnings. As noted already, the second of these is a parameters of interest, and is often the focus of the wide statistical discrimination literature.

The main disadvantage of this approach is that we cannot separate these two types of effects. Hence when there are no diploma effects on one, our estimates do not identify effects on the other. For example, suppose that the diploma was unrelated to the probability of reporting positive earnings and suppose that 80% of the groups with and without diplomas reported positive earnings. In that case, the diploma premium we would like to identify is the earnings difference between those with and without diplomas reporting positive earnings. The diploma premium that we actually identify is this premium multiplied by 0.8. Suppose instead that the diploma was unrelated to the diploma premium conditional on having positive earnings but affected the probability of having positive earnings. In that case, the total earnings effects would mask the effect of interest, the probability of reporting positive earnings.

Based on this discussion of the disadvantages of our approach, it is tempting to estimate these two effects separately. While it is straightforward to identify diploma effects on the first, the probability of reporting positive earnings, it is harder to identify diploma effects on conditional-on-positive earnings. Indeed, unless the probability of reporting positive earnings is unrelated to diploma status, estimates of these effects will likely be biased. Moreover, it is not clear in which direction these biases will go. In the classic selection scenario, we might expect estimates on conditional-on-positive earnings to be biased downwards. That is because positive effects will attract lower-quality diploma holders into the workforce. If diploma holders are more likely to pursue further education, the bias in the conditional-on-positive estimates will depend on which types of diploma holders are attracted to education.

To summarize, while our approach cannot answer all of the questions we might like to ask of the effects of a diploma, there is no obvious alternative. Moreover, our approach can be supplemented by estimates of diploma effects on the probability of reporting positive earnings. Under some assumptions, it can also be supplemented by bias-corrected estimates of diploma effects on conditional earnings. In particular, if we are wiling to make a monotonicity assumption, we can implement a version of the Lee (2008) trimming procedure to get upper and lower bounds on conditional earnings effects. If we are willing to make stronger functional form and distribution assumptions, we can use the Heckman (1979) selection-correction procedure to estimate diploma effects on conditional earnings.

If the main concern with our earnings estimates is that they obscure strong diploma effects on conditional-on-positive earnings, then this is not supported by any of these supplementary analyses. First. the diploma effect on the probability of observing positive earnings is small, around one percentage point. This can be seen in the third and fourth columns of Table 5, which reports diploma effects on the probability of working in each of the first seven years after the last-chance exam (the first two columns reproduces the preferred estimates for unconditional earnings). The fourth and fifth columns show the conditional-on-positive earnings estimates. These are comparable to the unconditional effects, although as discussed above, these may be biased. The Heckman (1979) selection-correction approach suggests that the bias may actually be positive, such that the true conditional-on-positive effects are smaller than these uncorrected conditional estimates and closer to the unconditional earnings estimates. Trimming for an upper bound generates larger estimates, as expected, but these are based on strong assumptions. While we cannot rule these out, the balance of evidence does not support this interpretation. Instead, it supports a simpler story under which the conditional-on-positive premium is small, the diploma effect on the probability of working is small and the unconditional earnings effect is also small.

Zero earnings biases

The discussion above assumed that workers with zero earnings really were out of the labor force. In practice there will be some slippage between true labor force participation and whether or not individuals are observed in our earnings data. Some individuals will be out of the earnings data because they work for the federal government. Individuals in the military will be the most important group in this category. Some individuals will be out of the earnings data because they are self-employed, work outside of the covered sector or work in the black economy. In our analysis, all of these individuals will be assigned zero earnings. In practice, their earnings will be larger.

This measurement error in earnings will only be a problem if there are diploma effects on the probability of being in one of these categories. With respect to the last categories - self-employed, in the uncovered sector, in the black economy - there is no strong reason to expect such an effect. In future version of the paper we analyze NLSY data to assess the correlation between these categories and student ability (as proxied by the AFQT). Military employment could be a bigger problem, because the military accepts very few individuals who do not have a regular high school degree. While there is no way of knowing how many of our sample joined the military, the military accepts very few applicants who score below the 31st percentile on the AFQT (Angrist, 1998). Since the mean score on the initial exit exam attempt was around the 10th percentile for students at the passing cutoff in the last chance sample, it is plausible that many students in the last chance sample would not be eligible for military service even if they held a high school diploma. Thus, any effect of failing the last chance test on the likelihood of enlisting in the Armed Forces may be small. Finally, the small estimates that we find for women also push against a military-based explanation, since women are less likely to enlist.

Information biases

Our working assumption has been that firms do not observe the number of times that a student takes the exit exam and do not observe workers' exact scores. As noted already, if firms can observe the number of times a worker took the exam, they can in effect observe a strong productivity signal. This will reduce the power of the additional signal sent by the diploma. We do not believe that firms observe this information and intend to test this assumption by assessing whether there is an earnings discontinuity associated with passing the first test. The workers on either side of the first test passing threshold should be equally productive. We have already shown that they earn a diploma with equal probability. If firms observe that one set of workers took the exam more often, there should be a positive earnings effect associated with passing the exam first time.

We have also assumed that firms cannot observe the exact test scores. This is consistent with studies of firms' knowledge regarding workers' high school performance (Bishop (1989)) and with anecdotal evidence from Texas and Florida. This assumption cannot be tested using our data, but we intend to test it using the information contained on AFQT scores in the NLSY. Specifically, if firms can observe this information, and if the exit exam score is correlated with the AFQT score (as we would expect), then the AFQT-wage correlation will be higher in states with exit exams.

7 Discussion and Interpretation

In this section we try to reconcile our estimates with previous estimates of the high school diploma premium. We then discuss why the high school premium might be small.

7.1 Relation to the previous literature

Previous estimates of the high school diploma premium suggest that this is large, around 15-20 percent. We have already noted that these estimates might be upward biased. The problem is that in settings in which traditional high school graduation standards apply, there are likely productivity differences between workers with and without high school diplomas. Estimates will be biased if these are observed by firms but not the econometrician. As a partial test of whether these biases can explain why our estimates are much smaller, we use these methods on our data. The results are presented in Table 6. In the first column we reproduce our preferred discontinuity-based estimates of the high school diploma premium. In the middle and final columns we report estimates of the diploma premium obtained using the previous approaches. The middle columns report estimates based on the last chance sample. The final column reports estimates based on the full sample. The last chance sample estimates are significantly larger than the discontinuity-based estimates. The full sample estimates are in line with those found in the previous literature.

Our estimates are less easily reconciled with estimates of the GED premium. As already noted, these range from small to fairly large. At least two factors might explain why some of these estimates are larger than ours. First, the difference-in-difference methods used to estimate these GED premiums may not adequately control for productivity differences between workers with the same scores in different regimes. While both sets of estimates pass various robustness checks, the estimates reported by Lofstrum and Tyler (2007) are somewhat sensitive to regression adjustment for other observables, especially pre-treatment earnings. Second, different credentials may be associated with different premiums, in part because they held by workers in different segments of the labor market. In the theoretical framework outlined above, we showed that credential premiums will depend, among other things, on the productivity differences between workers with and without the credential and on the other productivity information held by firms. To the extent that the GED premium is higher than the high school diploma premium, it may be because firms hold much less information about high school dropouts. These will, for example, have patchier employment histories, making it harder for firms to base productivity expectations on this type of information.

7.2 Why is the high school diploma premium small?

To consider why the high school diploma premium might be small, return to equation (1). This assumed that wages reflect a worker's expected productivity conditional on the information observed by firms: high school diploma status D and other productivity signals π^s. The parameter β_1 in this equation reflects the extent to which a high school diploma predicts productivity conditional on the other information observed by firms.

We have shown that this coefficient is small, which suggests that high school diploma status does not predict productivity given the other information observed by firms. This cannot be because diploma status does not predict productivity at all, since we found a positive reduced-form relationship between exit exam scores (predictors of high school diploma status) and earnings. Instead, it implies that exit exam scores do not predict productivity conditional on firms' other information. It is interesting to consider what this information might include. Among other things, it could include the outcome of productivity tests taken on the job, firms' observations of worker performance in interviews, letters of recommendation, employment histories and so on. In future versions of the paper we will investigate the other information held by firms using data from the Bay Area Labor Study (BALS).[7]

A possibility not yet considered is that high school diploma status is not verifiable and not observed. Suppose in the extreme case that all of the workers that did not receive a diploma nevertheless claim that they did. In that case, firms would observe no variation in diploma status hence there could be no diploma premium. Suppose instead that a fraction of the workers that did not receive a diploma claim that they received one. This will have two consequences. First, the diploma premium of interest will be that associated with a reported diploma, since this is what firms observe. This reported diploma premium will be smaller than the true diploma premium would be, since misreporting will narrow the difference between the expected productivity of workers that report and do not report having a diploma. Second, if we re-interpret our estimates as estimates of the reported diploma premium, they will be biased downwards. That is because our "first stage" estimate of the effect of passing on diploma status will over-estimate the impact of passing on reported diploma status. As a result, the reduced-form estimates will be scaled up by too small a factor. While misreporting is a real possibility, it seems unlikely that our approach has under-estimated a true positive reported diploma premium. First, our reduced-form estimates are already small. Even if these were multiplied by five instead of 2.5 (i.e., consistent with a misreporting rate of 50 percent), they would still be within 5 percent and in some years would be negative. Second, out estimates of the impact of diploma status on GED attempting is not consistent with the claim that diploma status is unverifiable such that all workers can report having one.

Whatever the reason for the small high school diploma premium, it is clear that this need not generalize to other credentials and other labor markets. As already noted, firms might have less information about other types of workers such as high school dropouts. Other things equal, this will lead to higher credential premiums in these markets. If high school diploma receipt is hard to

[7] This contains detailed information on the hiring and pay practices of firms employing non college-bound workers.

verify, credentials that are more easily verified by command larger signals. Despite these caveats regarding the generalizability of our findings, as noted in the Introduction, we think it is significant that we find only a small premium associated with the most commonly held educational credential in the US, in two of the largest states in the US, and across all of the major demographic groups in those states.

8 Conclusion

Estimates of the wage premiums associated with educational credentials - so-called sheepskin effects - are said to provide some of the clearest evidence in support of a signaling role for education. Yet previous estimates of these premiums may be biased upwards. This paper exploited the existence of high school exit exams to implement methods that identify the high school diploma premium under extremely weak assumptions. Our estimates suggested that the premium is small, much smaller than the premiums estimated in the previous literature.

The second section of Spence's classic article is headed "Hiring As An Investment Under Uncertainty". This is the foundation upon which the signaling hypothesis is built. Our analysis suggests that conditional on the other information held by firms, addutional information on education does not reduce the extent of this uncertainty. It is not clear where this additional information comes from - this is a "black box" that we leave for future research to investigate. Wherever it comes from, it implies that labor market signalling may be a less important phenemenon than was previously thought.

References

Acemoglu, D., and J. Angrist (2000): "How Large are Human-Capital Externalities? Evidence from Compulsory-Schooling Laws," *NBER Macroeconomics Annual*, 15, 9–74.

Altonji, J., and C. Pierret (1997): "Employer Learning and the Signaling Value of Education," in *Industrial Relations, Incentives and Employment*, ed. by I. Ohashi, and T. Tachibanaki. Macmillan, London.

——— (2001): "Employer learning and statistical discrimination," *Quarterly Journal of Economics*, 116, 313–350.

Altonji, J. G., and R. M. Blank (1999): "Race and Gender in the Labor Market," in *Handbook of Labour Economics Volume 3A*, ed. by O. Ashenfelter, and D. Card. Elsevier, Amsterdam.

Bedard, K. (2001): "Human Capital Versus Signaling Models: University Access and High School Drop-outs," *Journal of Political Economy*, 190(4), 749–775.

Betts, J. (1998): "The Impact of Educational Standards on the Level and Distribution of Earnings," *American Economic Review*, 88(1), 66–76.

Bishop, J. (1989): "Incentives for Learning: Why American High School Students Perform So Poorly Relative to Their Counterparts Overseas," CAHRS Working Paper Series, Cornell University ILR School.

Costrell, R. (1994): "A Simple Model of Educational Standards," *American Economic Review*, 84(4), 856–971.

Dee, T. S., and B. Jacob (2006): "Do High School Exit Exams Influence Educational Attainment or Labor Market Performance," NBER Working Paper 12199.

Hungerford, T., and G. Solon (1987): "Sheepskin Effects in the Returns to Education," *Review of Economics and Statistics*, 69, 175–177.

Imbens, G. W., and J. D. Angrist (1994): "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 61(2), 467–476.

Jaeger, D., and M. Page (1996): "Degrees Matter: New Evidence on Sheepskin Effects in the Returns to Education.," *Review of Economics and Statistics*, 77, 733–9.

Lang, K., and D. Kropp (1986): "Human Capital vesus Sorting: The Effects of Compulsory Attendance Laws.," *Quarterly Journal of Economics*, 101(3), 609–624.

Lange, F. (2007): "The Speed of Employer Learning," *Journal of Labor Economics*, 25(1), 1–35.

Layard, R., and G. Pschacharopoulos (1974): "The Screening Hypothesis and the Returns to Education.," *Journal of Political Economy*, 82, 989–98.

Lee, D. S. (2008): "Randomized Experiments from Non-random Selection in U.S. House Elections," *Journal of Econometrics*, 142(2), 675–697.

Lofstrum, M., and J. Tyler (2007): "Modeling the Signaling Value of the GED with an Application to an Exogenous Passing Standard Increase in Texas," IZA Discussion Paper No. 2953, IZA, Bonn.

McCrary, J. (2008): "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test," *Journal of Econometrics*, 142(2), 698–714.

Riley, J. G. (1979): "Testing the Educational Screening Hypothesis.," *Journal of Political Economy*, 87(5), s227–s252.

——— (2001): "Silver Signals: Twenty-Five Years of Screening and Signaling.," *Journal of Economic Literature*, 39(2), 432–478.

Spence, M. (1973): "Job market signaling.," *Quarterly Journal of Economics*, 83, 355–79.

Tyler, J., R. Murnane, and J. Willett (2000): "Estimating the Labor Market Signaling Value of the GED.," *Quarterly Journal of Economics*, 115, 431–468.

Tyler, J. H. (2004): "What is the Value of the GED to Dropouts Who Pursue the Credential?," *Industrial and Labor Relations Review*, 57(4), 587–598.

Wolpin, K. (1977): "Education and Screening," *American Economic Review*, 73, 949–958.

Figure 1: Exit exam administrations and high school diploma receipt

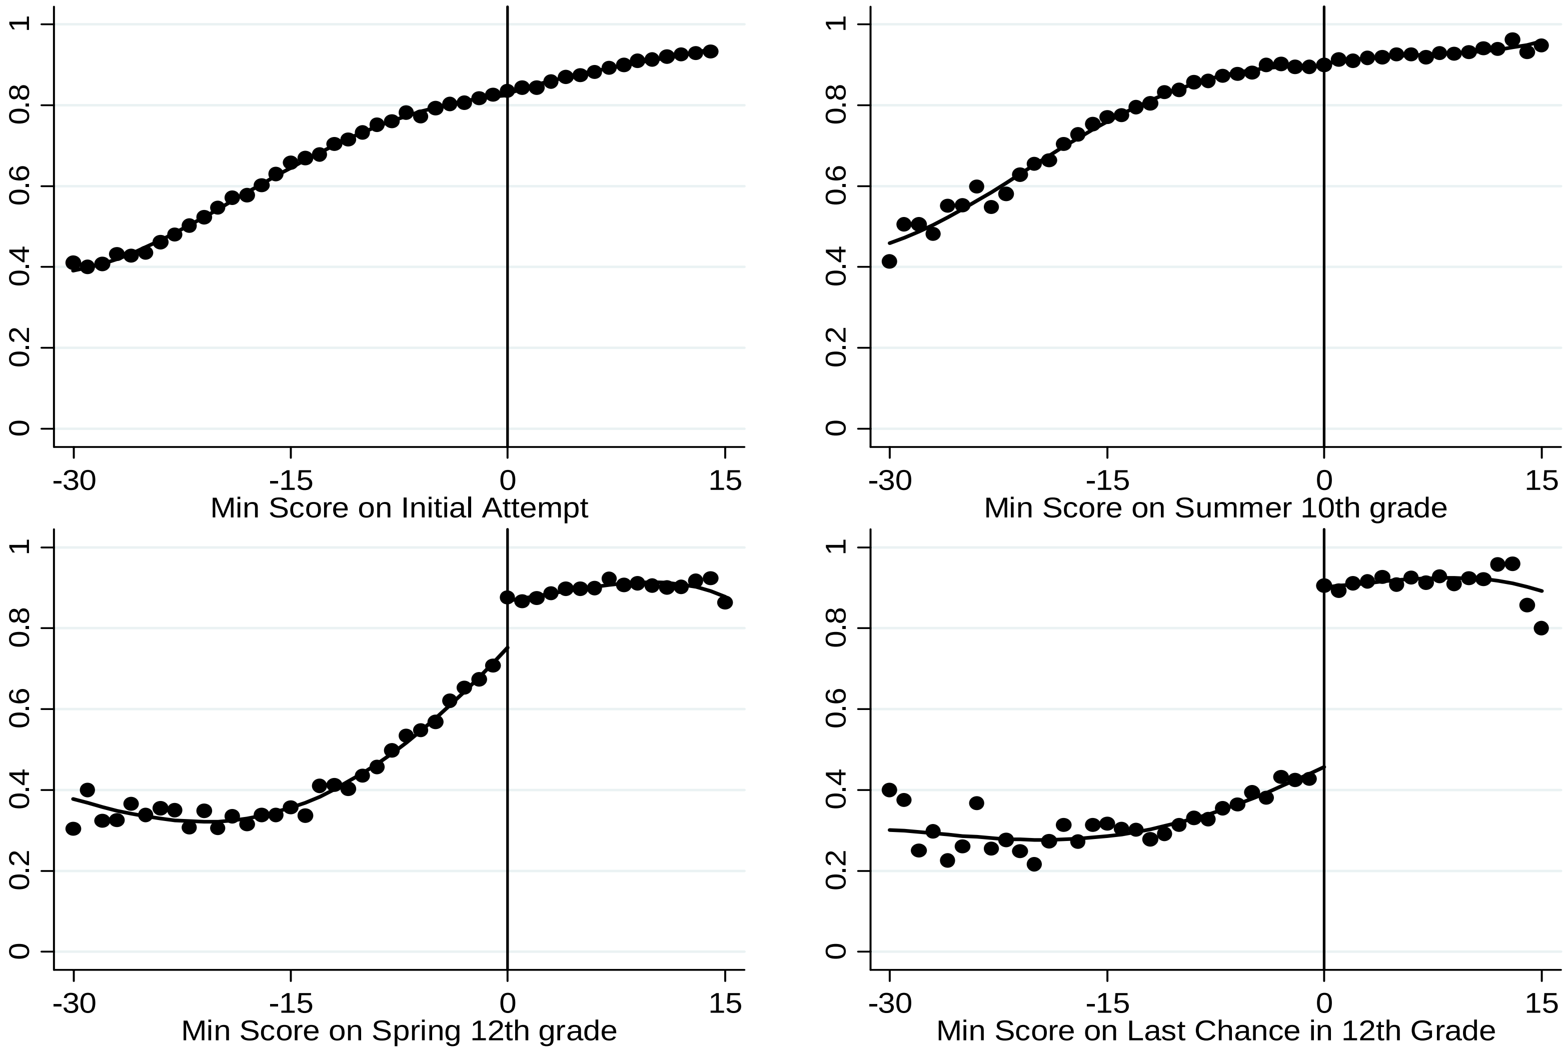

Notes: graphs based on full sample of those that take the first exit exam in grade 10 in 1991.

Figure 2a: Last chance sample (dash) versus full sample (solid)

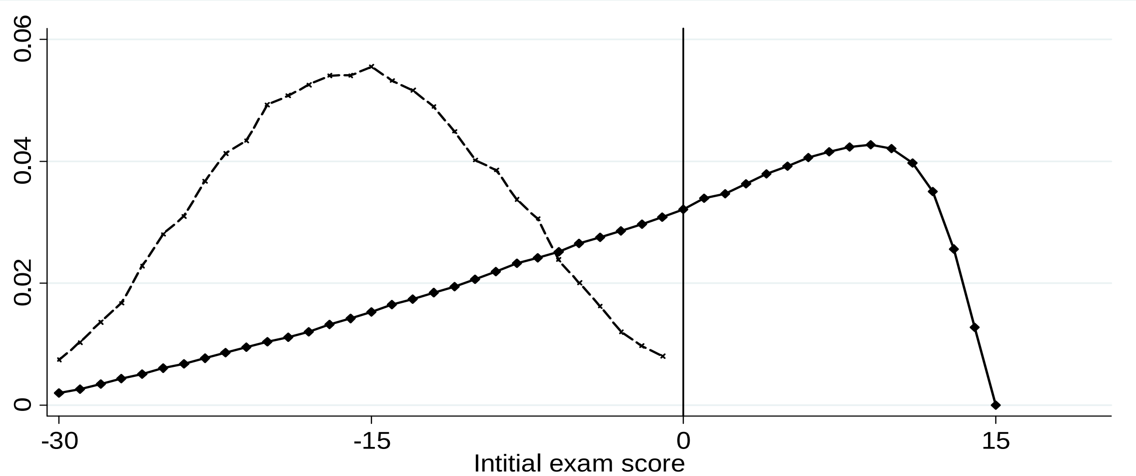

Notes: Full sample as defined in Figure 1. See the text and Table 1 for details on the last-chance sample.

Figure 2b: Last chance sample

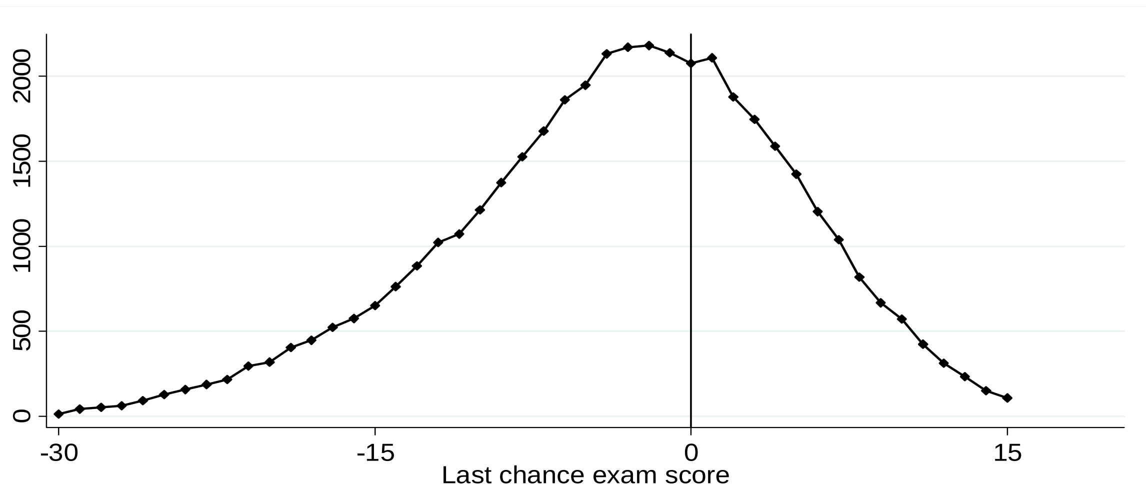

Notes: See the text and Table 1 for details on the last-chance sample.

Figure 3: Impact of passing last chance exam on probability high school diploma receipt before summer of 12th grade (solid) and within 2 years (dash)

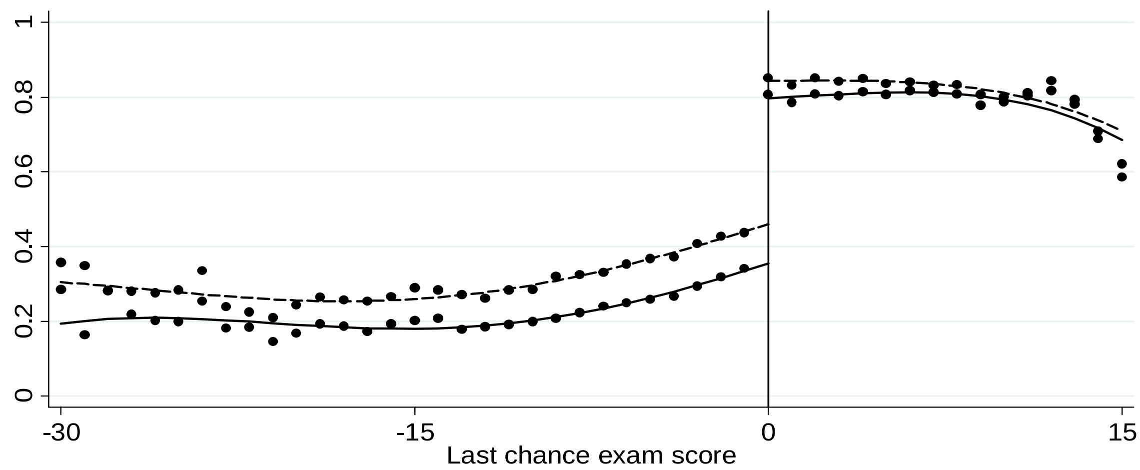

Notes: Figure based on last chance sample. See text and Table 1 for details.

Figure 4: Impact of passing last chance exam on earnings 1, 3, 5 and 7 years after last chance exam

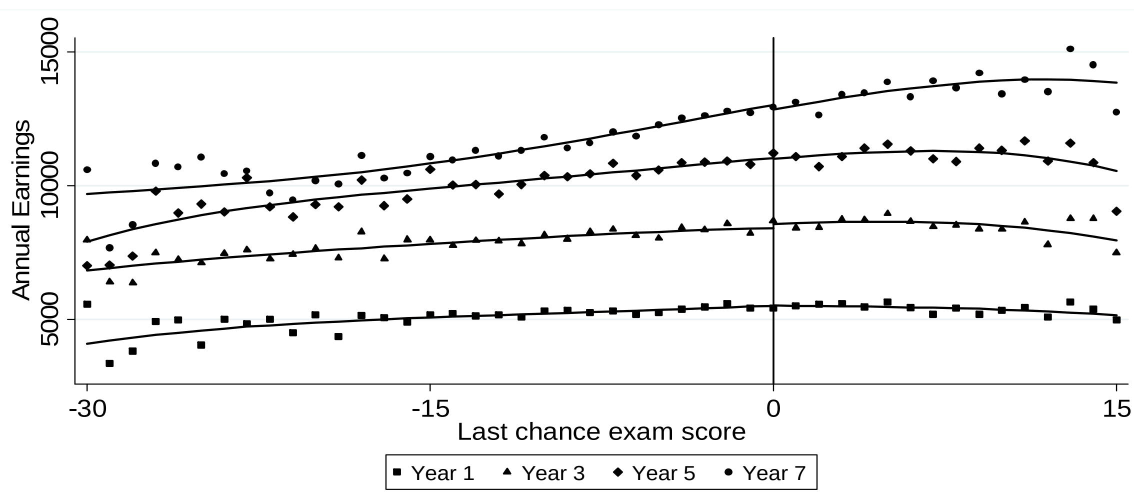

Notes: Figure based on last chance sample. See text and Table 1 for details.

Figure 5: Last chance exam score and earnings 1, 3, 5 and 7 years later

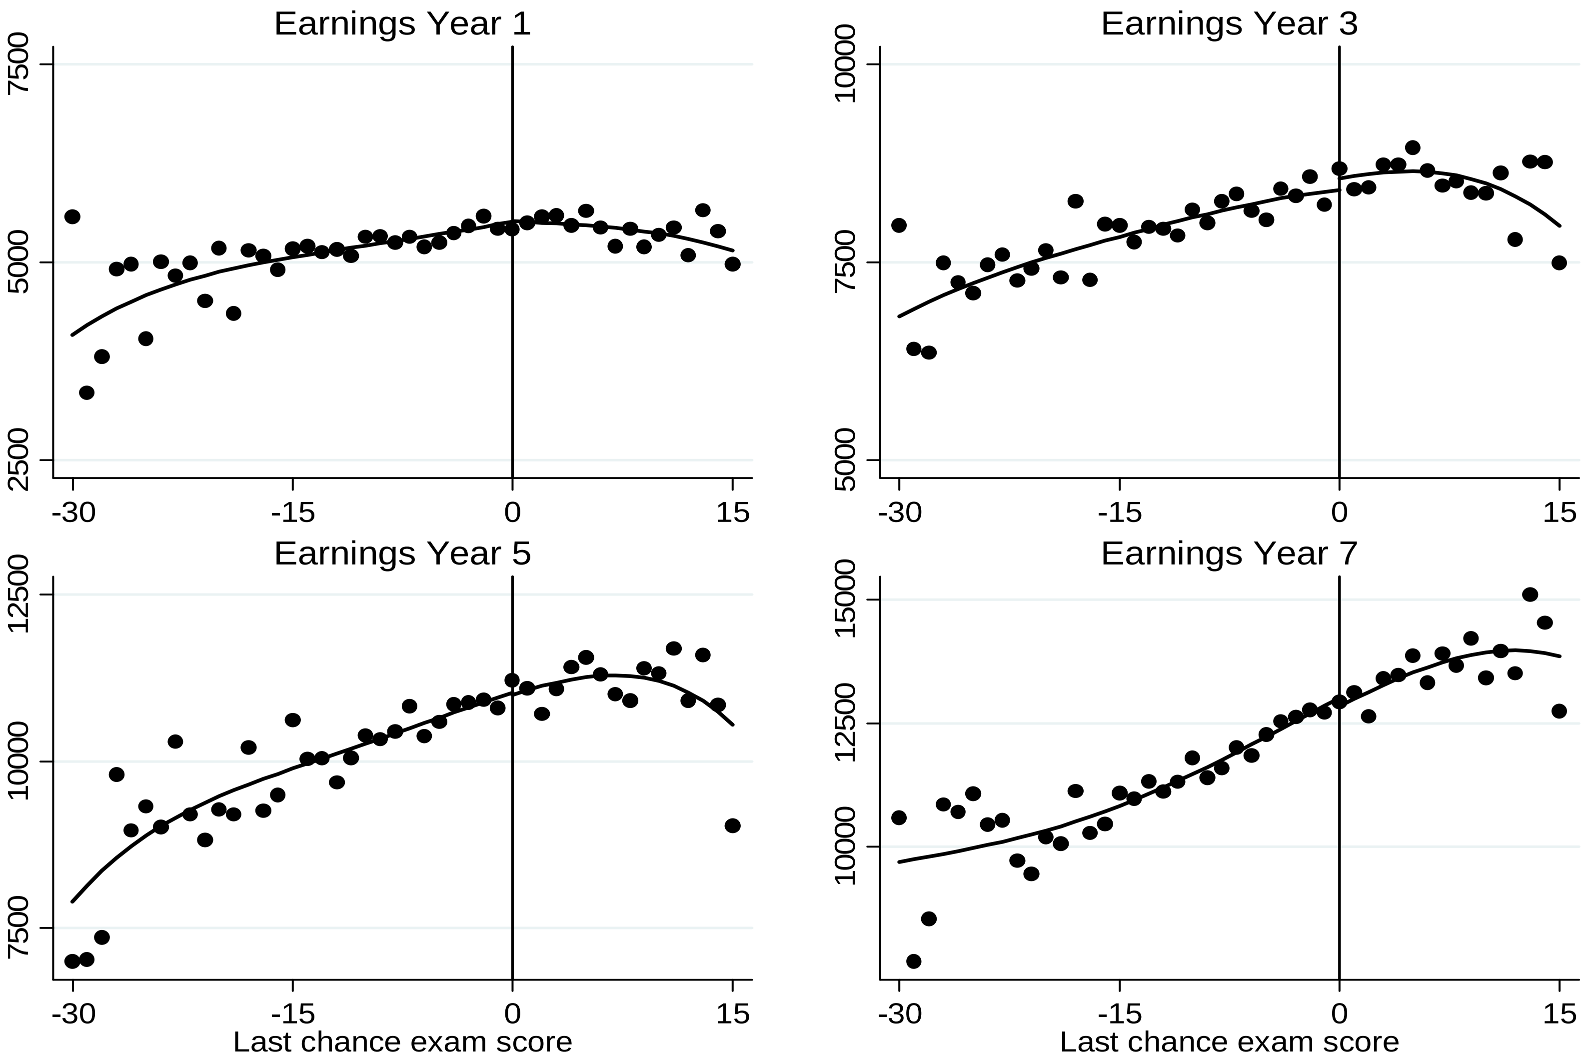

Notes: graphs based on "last chance sample" (see text and Table 1 for details). Zero earnings are assigned to those not observed in the earnings data.

Table 1: Descriptive statistics

	All			Fail	Pass	Discontinuity	
	Mean	SD	N	Mean	Mean	Coeff	Se
Initial math score	-1.711	0.988	41703	-1.921	-1.374	0.017	0.016
Initial reading score	-0.760	0.986	41743	-0.959	-0.441	0.032	0.019
Missing init math	0.019	0.136	42510	0.018	0.020	0.001	0.002
Missing init reading	0.018	0.133	42510	0.017	0.019	0.001	0.002
At grade level	0.543	0.498	42510	0.494	0.621	*-0.017*	*0.008*
Black	0.240	0.427	42510	0.252	0.222	-0.010	0.006
Hispanic	0.472	0.499	42510	0.502	0.425	0.002	0.008
Econ Disadv	0.394	0.489	42510	0.430	0.337	-0.008	0.006
LEP	0.145	0.353	42510	0.176	0.096	*-0.013*	*0.007*
Male	0.425	0.494	42510	0.418	0.436	-0.011	0.009
Special Ed	0.030	0.170	42510	0.036	0.020	0.001	0.003
1991 Cohort	0.358	0.479	42510	0.301	0.449	-0.009	0.006
1992 Cohort	0.162	0.369	42510	0.177	0.139	0.005	0.005
1993 Cohort	0.168	0.374	42510	0.174	0.160	-0.009	0.007
1994 Cohort	0.159	0.366	42510	0.182	0.123	0.008	0.007

Notes: Statistics refer to those taking the exit exam at the final administration before the end of high school.

Table 2: First-stage estimates: impact of passing last chance exam on probability of obtaining a high school diploma

	4th order polynomial		Local linear	
	w/o X	w/ X	w/o X	w/ X
Within 1 semester	0.442	0.444	0.440	0.446
	0.009	0.007	0.012	0.010
Within 2 semesters	0.421	0.424	0.420	0.426
	0.007	0.007	0.010	0.008
Within 3 semesters	0.389	0.392	0.380	0.387
	0.007	0.006	0.009	0.008
Within 1 year	0.383	0.387	0.376	0.383
	0.008	0.007	0.009	0.008
Within 2 years	0.381	0.384	0.373	0.380
	0.008	0.007	0.010	0.008
Within 3 years	0.380	0.384	0.372	0.379
	0.008	0.007	0.010	0.008
N	42510	42510	42510	42510

Notes: Table based on last chance sample. See text and Table 1 for details.

Table 3: Estimates of high school diploma earnings premium

	Reduced-form				Instrumental variables			
	4th order polynomial		Local linear		4th order polynomial		Local linear	
	w/o X	w/ X	w/o X	w/ X	w/o X	w/ X	w/o X	w/ X
Year 1	9.219	-4.704	19.581	16.135	24.052	-12.162	52.140	42.178
	89.319	78.523	83.539	72.844	232.943	203.047	222.568	190.402
	0.170	-0.087	0.361	0.298	0.444	-0.224	0.962	0.778
Year 2	120.603	115.014	121.432	133.155	314.639	297.342	323.346	348.075
	128.938	115.263	119.321	110.734	332.690	295.376	312.893	285.243
	1.745	1.664	1.757	1.926	4.552	4.302	4.678	5.035
Year 3	148.522	149.296	114.787	142.239	387.476	385.972	305.651	371.821
	159.163	156.297	159.431	162.980	411.299	401.405	419.574	421.712
	1.805	1.815	1.395	1.729	4.710	4.692	3.716	4.520
Year 4	149.817	156.879	133.032	184.064	390.855	405.576	354.233	481.156
	148.633	153.074	148.119	154.948	385.311	394.351	389.770	400.745
	1.559	1.632	1.384	1.915	4.066	4.220	3.685	5.006
Year 5	-30.704	-15.183	8.499	81.999	-80.104	-39.253	22.631	214.352
	175.866	191.373	179.993	192.914	459.175	494.878	478.992	502.332
	-0.284	-0.140	0.079	0.759	-0.741	-0.363	0.209	1.984
Year 6	-195.989	-170.232	-108.603	-22.127	-511.314	-440.097	-289.185	-57.841
	134.456	149.933	143.967	157.714	352.330	388.566	386.359	412.717
	-1.618	-1.405	-0.897	-0.183	-4.221	-3.633	-2.387	-0.478
Year 7	-167.960	-125.514	-232.171	-130.026	-438.189	-324.488	-618.217	-339.896
	176.159	185.683	172.026	179.572	459.073	479.832	460.345	470.495
	-1.321	-0.987	-1.826	-1.023	-3.447	-2.552	-4.863	-2.674
N	42510	42510	42510	42510	42510	42510	42510	42510

Notes: Table based on last chance sample. See text and Table 1 for details.

Table 4: Estimates of the impact of passing the high school exit exam on further education outcomes

	P(attend HS)	P(attempt GED)	P(earn GED)	P(enroll in coll)	College credits	Earn BA/AA
Year 1	-4.923			8.759		
	0.593			0.750		
Year 2	-0.073			0.543		
	0.212			0.717		
Year 3	-0.116			0.440		
	0.164			0.600		
Year 4				-0.669		
				0.463		
Year 5		-7.563	-5.782	0.211	0.419	-0.274
		0.521	0.277	0.480	0.322	0.219
	42510	42510	42510	42510	42510	42510

Notes: Table based on last chance sample. See text and Table 1 for details.

Table 5: Estimates of high school diploma premium conditional on positive earnings

	Reduced-form estimates									
	Unconditional		P(COP)		COP		Selection-corrected		Trimmed UB (2%)	
	GP4	LLR	GP4	LLR	GP4	LLR	GP4	LLR	GP4	LLR
Year 1	-4.704	16.135	0.217	0.217	-7.281	43.993	-20.003	47.080	334.338	367.421
	78.523	72.844	0.602	0.602	79.349	70.857	78.065	70.780	73.949	67.251
	-0.087	0.298			-0.106	0.642	-0.292	0.687	4.880	5.363
Year 2	115.014	133.155	0.542	0.542	126.809	137.589	177.436	192.899	544.545	523.250
	115.263	110.734	0.924	0.924	92.013	86.012	108.536	100.028	73.243	73.630
	1.664	1.926			1.433	1.554	2.005	2.179	6.152	5.911
Year 3	149.296	142.239	-0.499	-0.499	318.796	321.928	188.592	161.890	812.253	775.278
	156.297	162.980	0.648	0.648	204.955	184.379	193.430	182.540	183.161	173.346
	1.815	1.729			3.107	3.138	1.838	1.578	7.916	7.556
Year 4	156.879	184.064	-0.300	-0.300	277.465	304.683	234.593	275.689	786.117	787.703
	153.074	154.948	0.722	0.722	120.900	122.293	126.459	125.099	119.111	122.495
	1.632	1.915			2.273	2.496	1.922	2.258	6.440	6.453
Year 5	-15.183	81.999	-0.835	-0.835	157.228	329.332	-193.185	-136.730	708.391	842.775
	191.373	192.914	0.542	0.542	227.342	205.143	259.070	294.515	226.944	209.137
	-0.140	0.759			1.151	2.411	-1.414	-1.001	5.186	6.170
Year 6	-170.232	-22.127	-1.030	-1.030	5.432	271.504	-216.586	-49.784	667.098	893.866
	149.933	157.714	0.379	0.379	175.900	171.698	259.237	356.838	167.959	169.054
	-1.405	-0.183			0.035	1.751	-1.397	-0.321	4.302	5.765
Year 7	-125.514	-130.026	-0.900	-0.900	82.364	61.392	-33.755	-50.923	824.126	755.670
	185.683	179.572	0.620	0.620	221.744	205.708	225.926	233.833	182.178	180.398
	-0.987	-1.023			0.505	0.377	-0.207	-0.312	5.056	4.636
N	42510	42510	42510	42510	32404	32404	32404	32404	32023	32023

Notes: Table based on last chance sample. See text and Table 1 for details.

Table 6: Relation to literature: previous estimates of high school diploma premium

	Means			Last chance sample			Complete grade 12	
	LCS, fail at cutoff	LCS, no diploma	Finish G12, no diploma	IV estimates	Earnings(diploma) - Earnings(no diploma)		Earnings(diploma) - Earnings(no diploma)	
					No controls	Control for first score	No controls	Control for first score
Year 1	5422.15	5422.15	4989.811	-12.162	73.525	67.745	-340.601	-25.292
				203.047	53.957	54.024	24.753	25.042
				-0.224	1.356	1.249	-6.826	-0.507
Year 2	6912.458	6912.458	6532.346	297.342	315.051	313.082	-129.813	268.557
				295.376	66.403	66.489	30.678	31.031
				4.302	4.558	4.529	-1.987	4.111
Year 3	8226.184	8226.184	8048.953	385.972	577.983	574.337	122.318	569.221
				401.405	78.686	78.787	37.020	37.467
				4.692	7.026	6.982	1.520	7.072
Year 4	9611.93	9611.93	9749.169	405.576	783.079	775.714	411.610	783.599
				394.351	89.681	89.795	43.293	43.895
				4.220	8.147	8.070	4.222	8.038
Year 5	10806.72	10806.72	12375.080	-39.253	886.025	874.709	1560.506	1384.225
				494.878	97.020	97.140	52.535	53.347
				-0.363	8.199	8.094	12.610	11.186
Year 6	12113.2	12113.2	15016.030	-440.097	1129.815	1113.343	2909.515	2207.010
				388.566	106.537	106.664	62.335	63.117
				-3.633	9.327	9.191	19.376	14.698
Year 7	12712.85	12712.85	16685.510	-324.488	1323.734	1300.234	3824.286	2857.980
				479.832	114.789	114.917	69.710	70.470
				-2.552	10.413	10.228	22.920	17.129
N				42510	43831	43831	635010	635010

Notes: The "Lasdt chance sample" estimates refer to those based on the last chance sample (see text and Table 1 for details). The "Complete grade 12" sample refers to the subset of the full sample that are observed in the enrollment data up to and including the final two-week period in grade 12.

Appendix Figure 1: Baseline case (assuming $\gamma_1>0$, $\gamma_3>0$)

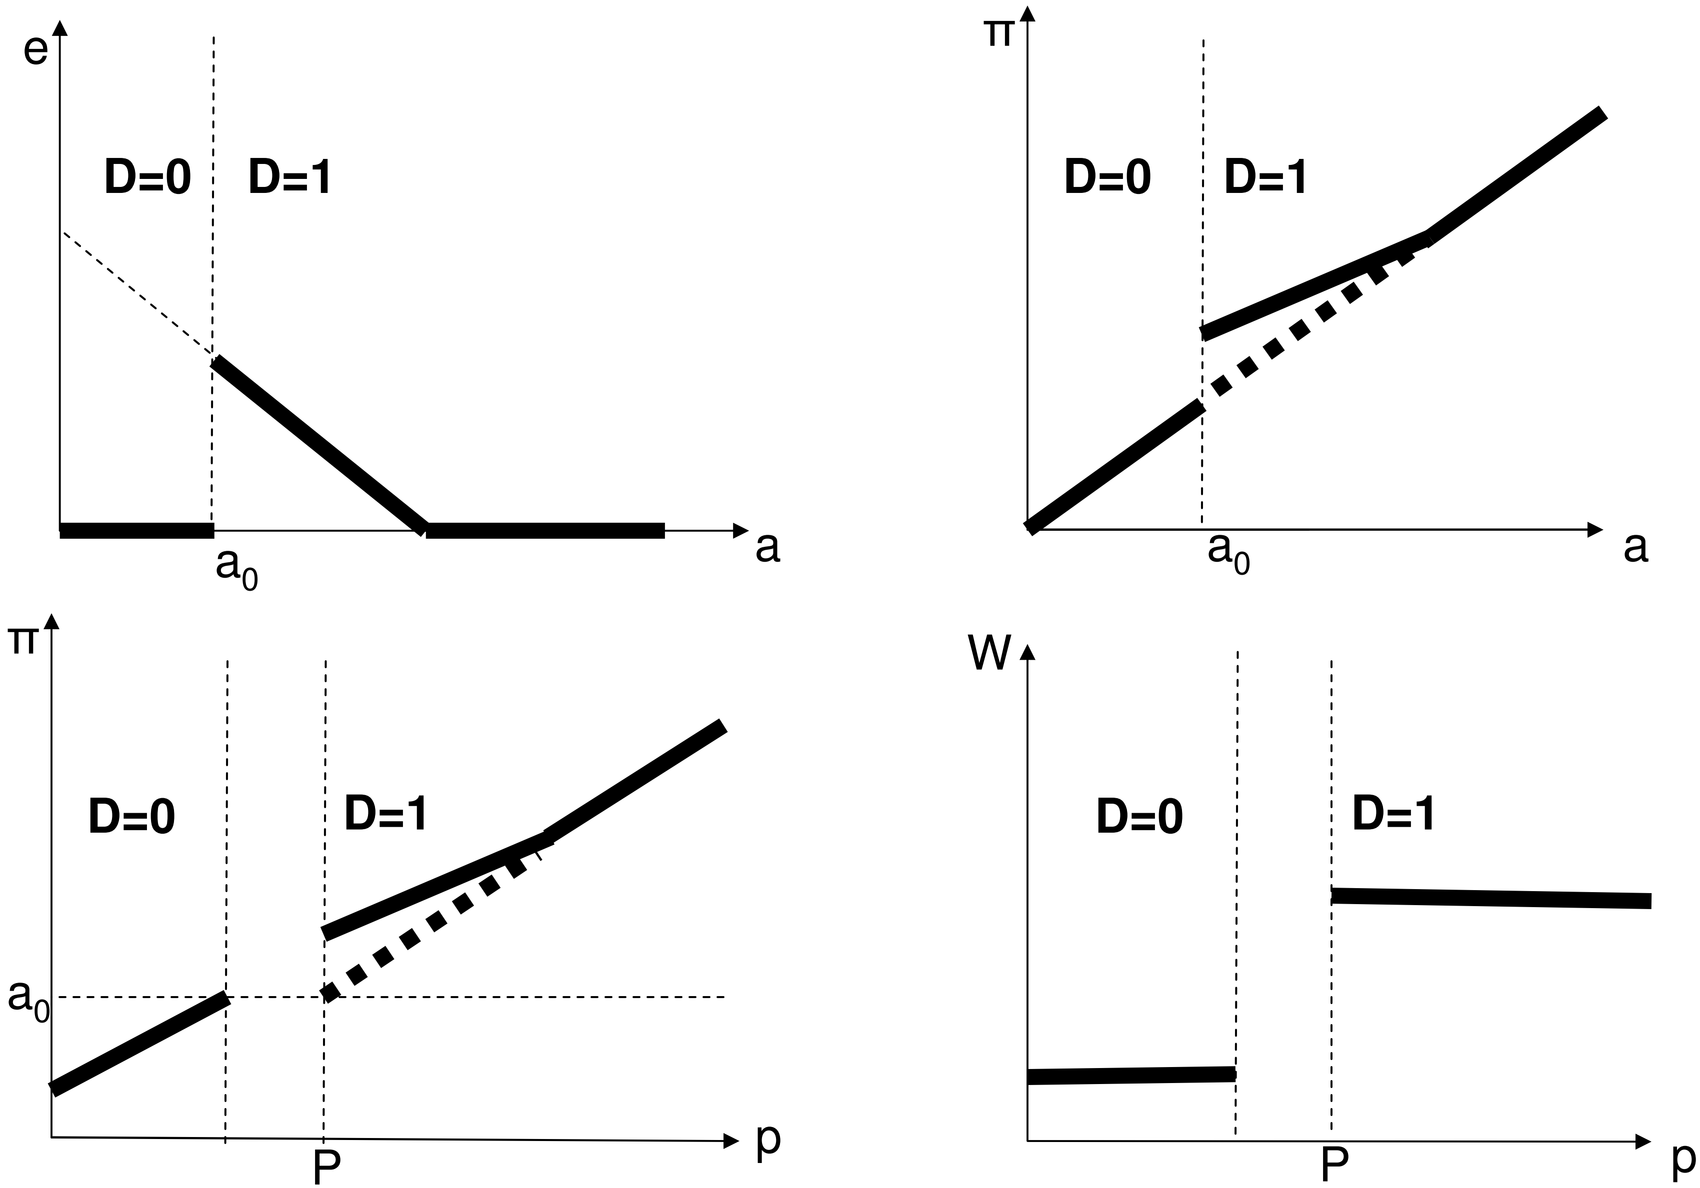

Appendix Figure 2: Firms observe second productivity signal (assuming γ_1>0, γ_3>0)

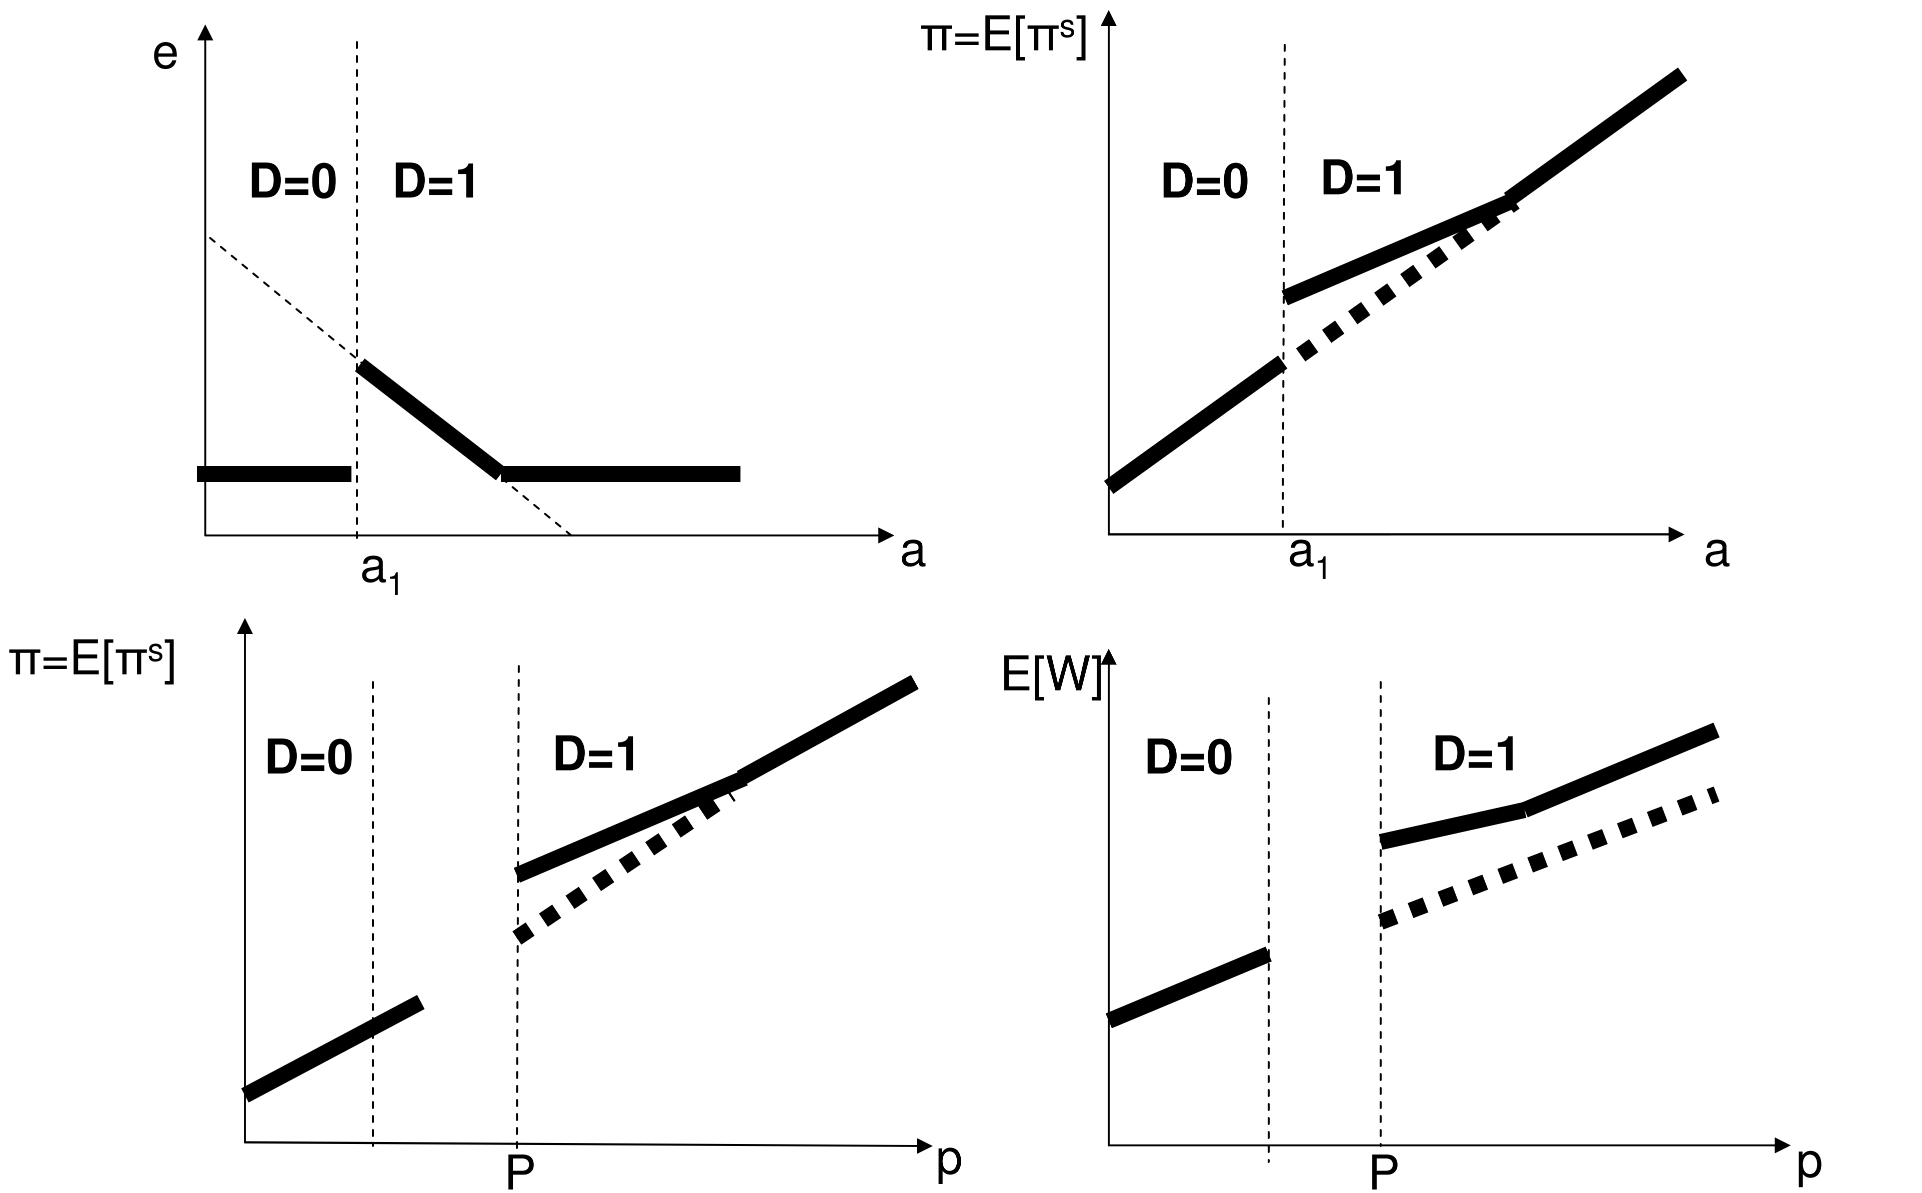

Appendix Figure 3: High school performance uncertain (assuming $\gamma_1>0$, $\gamma_3>0$)