The Signaling Value of a High School Diploma

Paco Martorell*
RAND

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

PRELIMINARY AND INCOMPLETE: PLEASE DO NOT CITE WITHOUT AUTHORS' PERMISSION

7 March 2010

Abstract

Economists accept that educational attainment could act as a signal of underlying ability (Spence, 1973) and that some types of education might send stronger signals than others. In particular, it is often said that education credentials such as a high school diploma could send an especially strong signal since, among other reasons, they are widely understood and should be easy to verify. In this paper we estimate the signaling value of a high school diploma using a regression discontinuity approach that compares the earnings of individuals that barely passed and barely failed high school exit exams, standardized tests that some states require students to pass in order to earn a high school diploma. Since these barely passers and barely failers should have the same distribution of productive characteristics as observed by firms, any earnings differences between these groups can be attributed to the signaling value of the diploma. Using linked administrative data on education and earnings from Florida and Texas, we estimate the signaling value of a high school diploma to be small: we cannot reject a zero effect and we can reject effects larger than 10 percent, much lower than those estimated in the previous literature. We interpret these findings as reflecting the non-trivial cost of verifying whether workers have a diploma and the relative ease of assessing their productivity using other information (e.g. trying them out on the job). For the same reasons, we argue that other types of lower-level education (e.g. time spent in school) are unlikely to act as strong signals of ability. We conclude that education-based signaling may be relatively unimportant at the lower end of the education distribution.

*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. In a labor market context, it implies that firms with incomplete information about worker productivity may base productivity expectations, and thus wages, on unalterable "indices", such as race and sex, and alterable "signals", including various types of education. In principle, several dimensions of education could act as signals of ability. At the lower end of the completed education distribution (e.g., for workers with some college or lower), these could include high school course choices, high school performance (e.g., GPA), length of time spent in high school (e.g., highest grade completed) and receipt of credentials such as a high school diploma or GED. At the higher end of the completed education distribution, these could include college choice, college performance and receipt of a college degree.

Many tests of education-based signaling have focused on the signaling content of educational credentials such as a high school diploma or college degree.[1] This is for two reasons. First, it has long been thought that credential receipt may send a particularly strong signal. As Layard and Pschacharopoulos (1974) observed, "graduation from a course should provide more evidence of ability and staying power than mere attendance for a number of years" (p.989). Second, because credentials are, ultimately, pieces of paper, they cannot have a direct effect on productivity. In principle then, one could estimate the signaling value of a credential by randomly assigning credentials across equally productive workers and then estimating the wage return to holding the credential. Since workers are equally productive, any wage return must reflect the signaling value of the credential.

[1] Tests that do not include analyses of whether the returns to education decline with labor market experience (as they would if firms gradually learned about ability), whether the private returns to education exceed the social returns to education (as they would if signaling dominated any positive externalities associated with education), whether education policy changes impact the educational decisions of students that they do not directly affect (as they would if the students wished to differentiate themselves from the directly affected students), whether the fit of the wage equation is relating to whether the occupation is screened and whether there are returns to education among the self-employed (as there would be if education increased productivity). See Farber and Gibbons (1996), Altonji and Pierret (2001) and Lange (2007) for the first, **?** for the second, Lang and Kropp (1986) and Bedard (2001) for the third, Riley (1979) for the fourth and Wolpin (1977) for the fourth.

In this paper we focus on the signaling value of the high school diploma, the most commonly held credential in the U.S. This may be an especially important signal of ability for workers in the lower part of the completed education distribution. The existing evidence is consistent with this, and suggests that a high school diploma carries an ability signal worth between 10% and 20% of the earnings of those without a diploma. This evidence comes from studies that estimate the signaling value of the diploma (often termed the "sheepskin effect") by regressing wages on diploma receipt conditional on having completed 12 years of schooling (e.g., Hungerford and Solon (1987), Jaeger and Page (1996)). The intuition is that a positive diploma return is evidence that firms view diplomas as signals of underlying ability. The problem is that positive return can also be given a productivity explanation, if workers with diplomas are more productive than workers without them and if firms observe some of these productivity differences. Many of the papers in this literature acknowledge these problems.[2]

More convincing estimates of the signaling value of a high school diploma might be obtained from research designs that better approximate the random assignment of diplomas across equally productive workers. In a seminal paper in the empirical signaling literature, Tyler, Murnane, and Willett (2000) employ a strategy that mimics this experiment to estimate the signaling value of a GED.[3] They note that to earn a GED, workers must pass nationally administered and nationally graded GED exams for which the passing thresholds differ by state. This generates across-state variation in GED receipt among workers with comparable GED scores (hence comparable ability). This suggests basing estimates of the signaling value of a GED on across-state earnings comparisons among workers with comparable scores (hence comparable ability). To disentangle the signaling value of the GED from across-state earnings differences driven by other factors, Tyler, Murnane, and Willett (2000) use a difference-in-difference strategy, comparing across-state

[2] Hungerford and Solon (1987) note that "our regression analyses may be biased by omission of ability variables or other factors correlated with degree completion" (p.177). Jaeger and Page (1996) note that "we are reluctant to interpret our results as purely causal effects of diploma receipt on wages" (p.739).

[3] The GED is a credential pursued by high school dropouts. It is sesigned to be equivalent to a high school diploma.

earnings differences among these workers with across-state earnings differences among workers with scores that always (or never) earn a GED. Their estimates suggest that the signaling value of the GED is large, between 12% and 20%.

If this result generalizes to the signaling value of a high school diploma, it would imply, consistent with the earlier literature, that the signaling value of a high school diploma is large. There are, however, reasons for questioning whether the result does generalize. First, subsequent papers in the GED literature, based on versions of the same strategy, find smaller GED effects. Tyler (2004) uses various strategies to compare the earnings of workers that passed and failed a GED in the same state and finds effects of around 10%. Lofstrum and Tyler (2007) use difference-in-difference strategies based on within-state variation in GED passing scores and cannot reject zero signaling values. There is, therefore, no consensus as to the signaling value of a GED. Second, the signaling value of a GED may differ from the signaling value of a high school diploma. For example, if firms have less information regarding the abilities of high school dropouts (because they do not have high school transcripts, have patchy employment histories and so on), we might expect the signaling value of the GED to be particularly large.

This paper estimates the signaling value of a high school diploma by exploiting variation in the receipt of a high school diploma generated by high school exit exams, standardized tests that some states require students to pass in order to earn a high school diploma. More specifically, we use a regression discontinuity strategy to compare the earnings of students that narrowly pass these exams (and earn a high school diploma) with the earnings of students that narrowly fail these exams (and do not). The key assumption is that workers that narrowly passed and narrowly failed these exams will, on average, be equally productive, such that any earnings discontinuity can be interpreted as the signaling value of the diploma. This argument is intuitive, but we formalize it in the context of a simple theoretical framework. This makes precise the conditions under which the signaling value of the diploma can be identified and the likely determinants of these signaling values.

Using datasets that link detailed education records to high-quality administrative earnings data, we estimate the high school diploma premium in Florida and Texas. The individuals in these data are observed for up to seven years after they leave high school, roughly until age 25. Our estimates suggest that the signaling value of the diploma is small in both states. For example, our main estimates suggest that the diploma premium is less than 5 percent, and we can rule out premiums larger than 10 percent in either direction. Further analysis by sex and race suggests that premiums are small for all of these subgroups; further analysis by labor market experience suggests no consistent patterns.

Although we think that our approach is highly credible, two caveats must be applied to our estimates. First, because students can retake high school exit exams, we base our estimates on the subsample of students that sit for the last administration of the exam in high school, a group for whom the exam outcome is found to exert a strong influence on eventual diploma status. If signaling values are heterogeneous (e.g., because they vary with the other information held by firms), our estimates will be specific to the students on the margin of passing this "last chance exam". This might not be the signaling value for the typical high school student, although it might be the signaling value relevant to similarly marginal students whose diploma status will be affected by policies targeted at high school students at risk of dropping out. 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, our 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 the next part of the paper we discuss how our results ought to be interpreted. There are two main explanations for the small signaling values that we estimate. First, receipt of a high school diploma might be costly to verify. In other words, firms might observe very little credible education information. Second, firms might observe the exit exam scores (or other productivity information that substitutes for them), such that a high school diploma provides no additional information. In other words, firms might observe a lot of credible education and/or productivity information. For two reasons we think it is more plausible to imagine that firms have relatively little information. First, anecdotal evidence suggests that firms may find it costly to verify the existence of high school diplomas and the information on high school transcripts. This is supported by studies that suggest that firms hold little of the information contained on high school transcripts and by evidence suggesting that few hiring firms require high school diplomas or check high school diplomas.

In the final part of the paper we consider whether our results will generalize to other settings and to other types of education. We believe that our results will generalize to settings in which exit exams are not used and to other states in which they are. We also see no reason to suppose that our estimates are any smaller than the signaling values that would have been estimated with data from the 1970s and 1980s (our data begin in the 1990s). In support of this claim, a sheepskin-style analysis applied to our data generates estimates comparable to sheepskin-style analyses applied to earlier data (i.e., we estimate large signaling values when using these methods). Finally, we believe that our estimates will generalize to most types of education at the lower end of the education distribution. That is because information on these types of education (e.g., high school GPA and highest grade completed) is contained on high school transcripts and because we argue that this should be at least as hard to verify as the existence of a high school diploma. These conclusions imply, among other things, that the social returns to lower-level education may be at least as large as the estimated private returns to this education.

2 Theoretical Framework

We begin the paper by laying out a simple theoretical framework. This helps to clarify important issues surrounding the identification and interpretation of credential signaling values. In particular, we define the signaling values of a credential, describe how these can be identified and analyze their likely determinants. This discussion motivates the empirical strategy used in this paper and clarifies how this relates to strategies used in the previous literature. It also clarifies why the same credential may have a different signaling value for different types of workers (i.e., heterogeneous effects) and why different credentials may have different (average) signaling values.

2.1 Basic Setup

We assume that individuals live for two periods. In the first period they attend school, in the second period they work. We characterize individuals in terms of a one-dimensional index of ability as represented by the continuous random variable a. We assume that ability is known to individuals at the start of the first period but is never observed by firms.

In our setup, students earn a credential when a one-dimensional measure of school performance p exceeds some threshold level P. That is, $D = 1$ iff $p \geq P$.[4] We assume that p is determined by $p = \widehat{p} + \eta > P$, where $\widehat{p} = \gamma_0 a + \gamma_1 s$. The variable s is the amount of study effort exerted in the schooling period. We think of study effort as encompassing various behaviors and activities that improve school performance (given ability).[5] We assume that study effort $C(s, a)$ is costly for all individuals ($C_s > 0$) and is, potentially, more costly for lower-ability individuals ($C_{sa} \leq 0$). The variable η is a mean-zero random variable assumed to be independent of $\widehat{p}$ and continuously distributed on the real line.[6] This assumption ensures that all students will obtain a credential with probability between zero and one.

We assume that second-period productivity (denoted π) is given by $\pi = a + \gamma_2 s$, a function of

[4] 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 score on the exit exam.

[5] In the absence of exit exams, studying could correspond to attendance; in the presence of exit exams, it could correspond to exam preparation.

[6] There are two other 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). These will have similar implications to the case considered in the text.

ability and first-period study effort. This captures the idea that productivity is related to underlying ability and also implies that studying might be valued by firms (if $\gamma_2 > 0$). We assume that firms do not observe π. Instead, they observe D, whether a worker has a diploma, and π^s, a noisy signal of productivity. We assume $\pi^s = \pi + \varepsilon$, where ε is a mean-zero random variable assumed to be independent of π. We will describe π^s as a worker's type, as observed by a firm. This could be based, for example, on performance at a job interview. On average, more productive workers will perform better at job interview, but there is some slippage between interview performance and productivity. We assume individuals are perfect substutes in production and assume there are a large number of risk-neutral firms competing for their services. This implies that the second-period wage equals expected productivity. Productivity expectations will be based on a worker's type (π^s) and whether she holds a credential (D).

We note in passing that this framework encompasses two special cases. If $\gamma_0 = \gamma_2, \gamma_1 = 1, C_{sa} < 0$, we have the classic signaling idea that students obtain a diploma by engaging in unproductive studying, the costs of which are decreasing in ability (Spence (1973)). If $\gamma_0 = \gamma_1 = \gamma_2 = 1$ and $C_{sa} = 0$, then productivity equals school performance and credential status is a direct signal of productivity. This captures the idea that education standards can induce students to engage in productivity-enhancing studying (c.f., Weiss (1983), Costrell (1994), Betts (1998)).

2.2 Signaling Equilbria and Signaling Values

These assumptions describe a game of incomplete information in which firms choose wage functions $W(D = 1, \pi^s)$ and $W(D = 0, \pi^s)$ and workers choose study level $s(a)$. A perfect Bayesian equilibrium of this game must satisfy two properties. First, given the wage functions, given knowledge of the distribution of η and ε and assuming rational expectations, the study choices $s(a)$ must maximize the expected utility of workers. Second, given these study choices, given knowledge of the distribution of η and ε and assuming rational expectations, these wage functions must ensure that firms make zero profits.

We can define the signaling value of a credential for a worker of type π^s as:

$$sv(\pi^s) = W(D = 1, \pi^s) - W(D = 0, \pi^s)$$

Notice that this expression defines the signaling value as the wage difference between workers with and without a credential conditional on the other productivity information observed by firms. It is crucial that we define the credential signaling value as conditional on this other information. It would, for example, we incorrect it as the unconditional difference between the wages of those with and without a credential.[7]

Example of a Signaling Equilibrium

It is worth considering the signaling equilibrium in a simple special case of this framework. Suppose that (1) ability is uniformly distributed on $[a_L, a_H]$ where $a_H = \frac{P}{\gamma_0}$ (i.e., all students must exert some study effort to obtain the credential) (2) study costs are $C(s) = cs$, where $c > \gamma_2 + \frac{\gamma_1}{2\gamma_0}$ (i.e., sufficiently high so that not all students will obtain the credential) (3) $\eta = 0$, so that there is no uncertainty in the credential acquisition process and (4) the firm observes only D, so that wages will be conditioned on D only. This implies that a signaling equilibrium must feature an ability cutoff below which workers exert zero effort (and do not get a credential) and above which workers exert just enough study effort to get the credential. It is simple to show that there is only

[7] If the firm had more information than D, there would not be a single wage difference between those with and without the credential. Moreover, the expected difference would depend on the signaling values of the diploma and the difference in the distribution of types among workers with and without the credential. That is, $E[W(D = 1) - W(D = 0)] = \int \{sv(\pi^s) f(\pi^s | D = 0) + W(D = 1, \pi^s) \Delta f(\pi^s)\} d\pi^s$ where $\Delta f(\pi^s) = f(\pi^s | D = 1) - f(\pi^s | D = 0)$.

one such equilibrium, which can be characterized by:

$$
\begin{aligned}
sv^* &= [a_H - a_L][\frac{c}{2(c-\gamma_2)}] \\
s(a) &= \left\{ \begin{array}{c} 0 \text{ if } a < a^* \\ \\ \frac{\gamma_0}{\gamma_1}(a_H - a) \text{ if } a \geq a^* \end{array} \right\} \\
\frac{a_H - a^*}{a_H - a_L} &= \frac{\gamma_1}{2\gamma_0}[\frac{1}{(c-\gamma_2)}] < 1
\end{aligned}
$$

The signaling value is not a function of π^s because we assumed that firms observe only D.[8] We return to this equilibrium below.

2.3 Identification

There are two possible approaches to identifying $sv(\pi^s)$. First, the researcher could assume that he observes the worker's type π^s. Then he could estimate the type-specific signaling value by focusing on these types of workers and comparing the wages of these with and without the credential. The obvious limitation is that some elements of π^s may not observed by the researcher. Since we would expect D to be positively correlated with π hence π^s, these unobserved elements will bias upwards estimates of the signaling value of the credential.

Papers in the sheepskin literature have been criticized for assuming that the researcher can proxy for π^s. Some (e.g., **?**) have also criticized these papers for not explainin what drives variation in credential status among workers who appear similar to firms (i.e., what drives variation in D conditional on π^s). This framework shows that this variation can be driven by either η or ε. Variation driven by η implies that conditional on ability a, some workers are lucky (i.e., get a good draw for η) and get a diploma while other workers are unlucky and do not. Variation driven

[8] An equilibrium is a value of sv such that, given the optimal effort strategies, $E[\pi|D=1] - E[\pi|D=0] = sv$. For a given sv, the cutoff ability level a^* solves $C[\frac{\gamma_0}{\gamma_1}(a_H - a^*)] = sv$. For a given a^*, $E[\pi|D=1] - E[\pi|D=0] = \frac{a_H - a_L}{2} + \frac{\gamma_0\gamma_2}{\gamma_1}(a_H - a^*)$. Equating and using $C(x) = cx$ gives $\frac{a_H - a^*}{a_H - a_L}$, which can be used to derive sv. There is only one solution to the equation $E[\pi|D=1] - E[\pi|D=0] = sv$. Given the assumption on c, this satisfies $0 < \frac{a_H - a^*}{a_H - a_L} < 1$.

by ε implies that some lucky low-ability workers (i.e., low a, $D = 0$ but large ε) will have the same type as some unlucky high-ability workers (i.e., with high a, $D = 1$ but small ε).

Variation in credential status driven by η can be exploited under weaker assumptions. This is the second approach to identifying $sv(\pi^s)$. For example, if the researcher could observe ability, he could estimate signaling values via a comparison of wages among those with and without the credential conditional on some ability level a':

$$\begin{aligned}\delta(a') &= E[W|D=1,a'] - E[W|D=0,a'] \\ &= \int_{\pi^s} sv(\pi^s) f(\pi^s|a=a') d\pi^s\end{aligned}$$

This would identify a weighted average of signaling values, where the weights were the probabilities that the worker was of a specific type conditional on having ability a'. We get from the first to the second line by noting that conditional on ability, credential status is determined by the random component of school performance, the realization of which has no impact on π^s.[9]

Since the researcher cannot, in general, observe ability, a more practical approach is to compare the wages of workers with school performance at the passing threshold (who obtain a diploma) and the wages of workers with school performance just below the passing threshold (who do not):

$$\begin{aligned}\delta(P) &= E[W|p=P] - \lim_{\Delta\to 0^+} E[W|p=P-\Delta] \\ &= \int_{\pi^s} sv(\pi^s) f(\pi^s|p=P) d\pi^s\end{aligned}$$

This will also identify a weighted average of signaling values, where the weights are the probabilities that workers are of different types conditional on school performance being close to the threshold. We get from the first line to the second line provided $f(\pi^s|p=P) = \lim_{\Delta\to 0^+} f(\pi^s|p=P-\Delta)$ which will follow if $f_a(a|p=v)$ is smooth through P.[10] That is, if the distribution of ability is the

[9] In other words, $f(\pi^s|a=a', D=1) = f(\pi^s|a=a', D=0)$.

[10] This is because $f(\pi^s|p=v) = f(\pi^s|a) f_a(a|p=v)$ and because $f(\pi^s|a)$ is not a function of p conditional on a.

same for the workers in a small interval around the passing threshold.

It can be shown that $f_a(a|p = v)$ will be smooth through P provided school performance is uncertain (i.e., the distribution of η is non-degenerate).[11] When the distribution of η is degenerate, this will not be the case. Instead, the distribution of school performance will have zero density just to the left of the threshold because, intuitively, it will not be worth a worker studying almost - but not quite - hard enough to pass. The example equilibrium described above exhibited this property. School performance was certain, workers with ability $a = a^*$ exerted positive study effort and just passed while workers with ability $a < a^*$ exerted zero study effort and failed by a margin of at least $\gamma_1 s(a^*)$.[12] It is easy to show that this "zero density" result carries through to the more general case in which the firm also observes worker's type π^s.[13]

2.4 Interpretation

Because signaling values are conditioned on π^s, they can differ across types. To the extent that specific subgroups are associated with specific types, this implies that they can also vary across subgroups (e.g., sex and race). Average signaling values can also vary across credentials. To see this, return to the simple example considered above and note that the equilibrium signaling value was decreasing in c and increasing in γ_2. The parameter c measures the costs of the studying required to obtain the credential; the parameter γ_2 measures the productivity returns to the studying required to obtain the credential. The dependence on γ_2 implies that the signaling value

[11] Because we have assumed that η is a continuous random variable distributed on the real line, we have that for any equilibrium study strategies and for every ability type $a, 0 < P(fail|a) = G_\eta[P - p(a)] < 1$. This implies that for all a, $G_\eta(v - p(a))$ and $g_\eta(v - p(a))$ are continuous at $v = P$ and $g_\eta(v - p(a)) > 0$ when $v = P$. Since $f_p(v) = \int_a g_\eta(v - p(a))f(a)da$, this implies that $f_p(v)$ is continuous at $v = P$. In turn, this implies that $f_a(a_0|p = v) = \frac{f_p(v|a_0)f_a(a_0)}{f_p(v)} = \frac{g_\eta(v-p(a_0))f_a(a_0)}{f_p(v)}$ is continuous at $v = P$.

[12] Since $\widehat{p}(a^*) = \gamma_0 a^* + \gamma_1 s(a^*) = P$, $\lim\limits_{\Delta \to 0^{+-}} (P - \widehat{p}(a^* - \Delta)) = \gamma_1 s(a^*) > 0$.

[13] We want to show that we cannot have an equilibrium in which $\eta = 0$, $f(\widehat{p})$ is continuous at $\widehat{p} = P$ and some $sv(\pi^s) > 0$. Notice first that if $f(\widehat{p})$ is continuous at $\widehat{p} = P$, then there exist some a for which $\widehat{p}(a) \geq P$ and some a for which $\widehat{p}(a) < P$. Given our assumptions on study costs and $\gamma_0, \gamma_1, \gamma_2$, $\widehat{p}(a)$ must be increasing in a among both groups and $\widehat{p}(a) < \widehat{p}(a^*)$ for all $a < a^*$ where $\widehat{p}(a^*) = P$. Since $f(\widehat{p})$ is continuous at $\widehat{p} = P$, $\lim\limits_{m \to 0^+} \widehat{p}(a^* - m) = \widehat{p}(a^*) = P$ and $\lim\limits_{m \to 0^+} s(a^* - m) = s(a^*)$. This implies that study costs are smooth through a^* which implies that for these strategies to be optimal, $\lim\limits_{m \to 0^+} E[W|a^* - m, s(a^* - m), D = 0] \geq E[W|a^*, s(a^*), D = 1]$ hence $E[W|p = P] - \lim_{\Delta \to 0^+} E[W|p = P - \Delta] \leq 0$. This cannot be the case if there exists a positive signaling value.

will be higher the more strongly correlated are the credential requirements and productivity. 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. More generally, when firms observe credential status and additional productivity information (i.e., type), we would expect signaling values to be decreasing in the informativeness of this information. This implies that credentials may have especially large signaling values when firms observe little else about the workers that pursue them.

Note finally that there are two special cases of this framework in which signaling values are zero. First, if $C_{as} = 0$ and $\gamma_0 = \gamma_2 = 0$, there is no mechanism linking s and π, school performance does not predict productivity and $W(D = 1, \pi^s) - W(D = 0, \pi^s) = 0$. Second, if the firm observes π (i.e., $\varepsilon = 0$) then wages equal observed productivity, $W(D, \pi^s) = W(D, \pi) = W(\pi)$ and $W(D = 1, \pi^s) - W(D = 0, \pi^s) = 0$. There are also two variations on the assumptions made in this framework that would generate zero signaling values. First, if the credential is non-verifiable, then D is, in effect, unobservable, and $W(D = 1, \pi^s) = W(D = 0, \pi^s) = W(\pi^s)$ so that $W(D = 1, \pi^s) - W(D = 0, \pi^s) = 0$. Second, if the firm observes p, then the wage will be conditioned on p and π^s rather than D and π^s. Conditional on p, there is no overlap in D hence the signaling value is no longer defined as it was above. Moreover, the difference in wages among workers either side of the passing threshold will be zero, since $f(\pi^s|p = P) = \lim_{\Delta \to 0^+} f(\pi^s|p = P - \Delta)$ hence $E[W|p = P] - \lim_{\Delta \to 0^+} E[W|p = P - \Delta] = 0$. We return to these possibilities below.

3 High School Exit Exams

High school exit exams were introduced to combat the idea that high school diplomas were awarded for "time-serving" (see Martorell (2010) for more details). They do so by making diploma receipt conditional on passing statewide exams of basic skills. For the purposes of identifying the high school diploma premium, two aspects of these exams are relevant. First, they introduce an element of uncertainty into the diploma acquisition process. In particular, it is reasonable to suppose that students cannot forecast their performance on these exams, hence reasonable to suppose that, by chance, there exist workers with and without these diplomas for whom productivity overlaps. Second, performance on these exam is, potentially, observed by the econometrician. This implies that the diploma premium can be identified via a threshold comparison of the workers that narrowly pass and narrowly fail these exams. Intuitively, these two sets of workers will differ with respect to diploma status but not in expected productivity. If true, this implies that this comparison can identify the diploma premium *irrespective of whether the econometrician observes as much productivity information as the firm.*

The intuition underlying this approach is most clear under the hypothetical scenario in which students remain in school until the end of twelfth grade, then take a single test that determines whether or not they graduate. In practice, there are several ways in which high school exit exams deviate from this stylized description. While the details of their operation differ across states, the important details are common to all of the states that employ them. In what follows we focus on Texas and Florida, the states that we focus on here.

Multiple test subjects

Typically, students must pass exams in several subjects. 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. In Florida, since the exit exam became the tenth grade FCAT, it has also consisted of reading, math and writing tests. 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, opportunities have ranged from six to eight, depending on the cohort. In Florida, students have five opportunities to pass. 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 $i \leq L-1$).

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.[14] With regard to interpretation, multiple test-taking and 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. 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.[15] 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.

Imperfect compliance

[14] 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.

[15] 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.

In practice, there is some slippage between passing the exam and obtaining a diploma, even among students in the last-chance sample. Imperfect compliance could occur because students fail the last-chance exam but are exempted from the passing requirement. It could also occur because students fail the last-chance exam but re-take and pass it later, typically by returning to high school for a "thirteenth" year. It could also occur because students pass the last-chance exam but are denied a diploma because they have not fulfilled other high school graduation requirements. Texas and Florida all three sources of non-compliance. The main difference between them is that in Florida, students that are exempted from the testing requirements receive an "alternative" diploma. Students exempted in Texas receive a standard diploma. Martorell (2010) provides more details on the process that students must follow in order to be exempted in Texas.

With regard to identification, imperfect compliance can be handled via a modified version of 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)). While we might expect exemptions to affect the equilibrium diploma premium, only a small number of exemptions are granted.[16] The same can be said for the practice of denying diplomas to students that pass

[16] To see this, recall the equilibrium in the model in which high school performance is certain and firms observe

the exit exam but do not complete other requirements.

High school dropout and further education enrollment

The stylized description assumed that students complete high school, take the exam then work. It is easy to modify this so as to allow for student dropout. If firms can observe years of schooling, hence 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. 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.[17] 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

only diploma status 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.

[17] 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.

could also generate biased estimates of the diploma premium. This is because enrollment in further education is an endogenous downstream outcome that could be affected by whether or not students receive a diploma.[18]

4 Empirical Framework

In this section we provide a formal description of our empirical strategy. We do so in the context of a general model of wage determination that draws on the main insights of our theoretical framework and our discussion of the practical operation of exit exams.

Without loss of generality, we assume that firms observe whether workers hold a high school diploma (D) and a vector of other productivity signals π^s. We also assume that firms observe years of completed schooling, such that the following equation can be viewed as a linear approximation to the expected productivity (hence 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 parameter β_1 measures the diploma premium: the extent to which the diploma predicts productivity conditional on the other productivity information observed by firms. For ease of exposition, we assume this parameter is the same for all workers. As noted above, a less restrictive assumption is that it is heterogenous, a function of π^s. In either case, β_1 will depend on the accuracy of the productivity signal π^s. As this becomes more accurate, the diploma premium will weaken. In the full information case in which firms can observe productivity, $\beta_1 = 0$. The same

[18] 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 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

result will be obtained if firms observe exit exam scores. That is because, conditional on those, dipoma status will not predict productivity hence wages.

In our theoretical framework, we argued that there are two approaches to identifying β_1. One assumes that the productivity information available to the researcher (variables X) can proxy for productivity information available to firms' (π^s). In this case, β_1 can be identified via regression of earnings on diploma status (D) and X. This is the method used in the sheepskin literature. The concern is that X will not capture all of the productivity information observed by firms, some of which loads onto D, generating upward-biased estimates of β_1. The second approach exploits variation in diploma status among equally productive workers. If this type of variation can be found, the signaling value of the diploma can be identified without any assumptions regrading π^s. For this type of variation to exist, the key requirement is that there exists some uncertainty in the diploma acquisition process.

In this paper, we focus on settings in which receipt of a high school diploma requires that students pass high school exit exams. The uncertainty that generates productivity overlap across workers with and without the diploma is that surrounding the outcome of the high school exit exam. While students may be able to predict this with some degree of accuracy, we assume they do not have perfect control over this outcome. This implies that the expectation of π (hence π^s) conditional on the exit exam score is smooth through the passing threshold. In order to use regression discontinuity methods, we make the slightly stronger assumption that this conditional expectation is smooth at all values of p and that it can be represented by some polynomial $g(p)$.

With perfect compliance, this assumption would imply that in the projection of π^s onto D and $g(p)$, the projection coefficient on D should be zero. This would allow us to rewrite (1) as:

$$W_i = \beta_0 + \beta_1 D_i + g(p_i) + \varepsilon_i \tag{2}$$

where ε is the projection error. We could then identify the signaling value of the diploma by regressing wages on diploma status and a flexible function of the exit exam score. Under imperfect complicance, the projection coefficient on D need not be zero and this no longer follows. However, in the projection of π^s onto $PASS$ (a dummy variable for obtaining a passing score) and $g(p)$, the projection coefficient on $PASS$ should be zero. This implies that we can use $PASS$ to instrument for D in the "outcome" equation (2). The associated "first-stage" equation is:

$$D_i = \alpha_0 + \alpha_1 PASS_i + g(p_i) + \omega_i \tag{3}$$

4.1 Estimation Issues

This approach allows us to identify the signaling value of the diploma under weak assumptions. In implementing this approach however, various issues must be confronted. First, one must deal with the fact that these exams are administered multiple times. In other words, should one estimate equations (3) and (2) using scores obtained on the first attempt, the second attempt or the last attempt? We focus on the last attempt, since this generates by far the largest correlation between passing status and diploma status (i.e., the strongest instruments). It is tempting to increase power by combining the scores on various attempts. Preliminary investigations suggest that this is possible, but that it would have only a marginal impact on the precision of the instrumental variables estimates. We therefore focus on the last-chance exam, which also has the virtue of being transparent.

A potential drawback with focusing on the last chance sample is that our estimates will be specific to this sample. To think more clearly about effect heterogeneity, suppose again that workers belong to one of a set of possible types (as perceived by firms). In this case, a more precise statement is that our estimates will be a weighted average of the type-specific treatment effects, where the weights are the probabilities that workers of specific type appear in the last chance sample. This differs from the average treatment effect, which is also a weighted average of type-specific treatment effects, but where the weights are the probabilities that workers of a specific type appear in the population of students that finish twelfth grade. The difference between the two sets of weights will depend how persistent are exit exam scores. If they are highly persistent, the workers in the last chance sample will be those that scored poorly in each administration of the test. If they are less persistent, the last chance sample will also include workers that are of relatively high ability.

A second issue is how to determine what portion of the high school diploma premium can be attributed to the downstream effects of passing the exam. We deal with this in two steps. First, we assess high school diploma effects on downstream outcomes. Second, we consider whether effects of these sizes could account for a large portion of the diploma premium that we estimate. Inevitably, this requires that we make assumptions about the earnings effects of these downstream outcomes.

A third issue concerns specification choices such as the form of g(.), the range of data over which the models are estimated and whether the models adjust for covariates.We estimated two types of models. First, we estimated models in which all the data were used and $g(.)$ was chosen from among a wide set of polynomials, all of which were fully interacted with the "pass" variable (i.e., allowed to take different slopes on either side of the threshold). The preferred $g(.)$ was the one that minimized the AIC statistic. As discussed by **?**, this statistic helps choose a functional form that balances the trade-off between generating a good fit and generating precise estimates. Using the preferred $g(.)$, we also estimated the reduced-form earnings effects using the **?** efficient estimator. This exploits the availability of two regression discontinuity estimators in the case in which the assignment variable is discrete (as it is in this case): one based on estimation of the limits on either side of the threshold; one in which data at the threshold are used to estimate the limit on the right of the threshold. The optimal weights depend on the goodness of fit of the regression specification (i.e., how well the fitted regression corresponds to the cell means). Second, we estimated models in which $g(.)$ was linear and a smaller range of data were used. As discussed by Imbens and Lemieux (2008) and **?**, this approximates a nonparametric regression approach.

The data range was chosen via a cross-validation procedure of the type discussed by Imbens and Lemieux (2008).

As we discuss in greater detail below, in both Texas and Florida, a linear function in the test score approximates the conditional expectation of earnings quite well. Not surprisingly then, the AIC statistic was minimized for a linear function. In contrast, the "first stage" relationship between high school graduation and the exit exam score is nonlinear over the whole range of scores. Thus, first stage estimates would be biased unless a flexible polynomial were chosen on the full range of data or a linear polynomial was applied to a narrow range of data. For Florida, the cross-validation procedure suggested that a bandwidth of 25 scale score points would be close to optimal for the first stage estimates. At this bandwidth, earnings estimates obtained with a linear $g(.)$ are comparable to but more precise than earnings estimates obtained using flexible polynomials applied to the full range of data. For Texas, the cross-validation procedure suggested that the optimal bandwidth for the first-stage estimates is very small (roughly five scale score points). At these bandwidths, the earnings effects are estimated imprecisely.

With all of this in mind, our preferred specification for the Florida analysis is one that uses a linear $g(.)$ on a narrow range of data - 25 points on either side of the passing threshold. For the Texas analysis, our preferred was the optimal polynomial (which turned out to be second-order) estimated over the full range of data. For both states, we discuss the results of additional specifications. The results turn out to be quite robust across a wide range of models. Note finally that all models reported here were estimated with and without covariates. As expected in a valid regression discontinuity design, the addition of covariates had minimal effect on the estimates.

5 Data

In this section we describe the data that we have acquired from Texas and Florida.

5.1 Texas

The Texas data come from the Texas Schools Microdata Panel (TSMP), a database housed at the Green Center for the Study of Science and Society on the University of Texas–Dallas campus. The TSMP is a collection of administrative records from various Texas state agencies that permits a longitudinal analysis of individuals as they proceed through high school and later as they enter college or the workforce.

The datasets used in this analysis are part of the the Texas Education Agency's (TEA) records. These files have information on enrollment, attendance, dropout status, GED acquisition, and performance on the exit-level TAAS. We examine five cohorts of test takers, students in 10th grade in the spring of 1991 through 1995. Linking multiple test records on the basis of a unique student identifier, we construct the test-taking history for each student. The resulting panel includes the test scores as well as the student's ultimate passing or exemption status. The main outcomes used in this paper are graduation status, further education acquisition and earnings.

Graduation

Information on a student's graduation status comes from the TEA's roster of students who receive a high school diploma from a Texas public school during an academic year. In this paper, a student is classified as a graduate if a matching record appears in these files indicating that the student earned a diploma within two years of the year in which they would graduate if they were not held back after initially taking the test. We refer to this year as the expected date of high school completion.

GED Test Taking and Certification

The TSMP's GED data includes a record for each time an individual in Texas who earned a GED certificate took the GED qualifying exam. Beginning in 1995, it includes records of all GED test takers regardless of eventual passing status. We construct two outcomes using these data.

The first is whether or not the student earned a GED within five years of the expected high school completion year. The second is whether the student took the GED exam within five years of this date.

Post-Secondary Schooling

The TSMP's data on post-secondary education comes from the administrative records of the Texas Higher Education Coordinating Board (THECB). Included in the THECB's files is a list of all students enrolled in public two and four-year colleges in Texas in a given academic term as well as the number of semester credit hours in which they are enrolled. In addition to looking at attendance, we also examine the amount of post-secondary schooling acquired. Using the THECB's data on degrees conferred by Texas post-secondary institutions, we construct an outcome that denotes a student earning a bachelor degree, associate degree, or certificate (such as a nursing or teaching certificate). Since the available data stops in 2001, only degrees awarded four years after the expected end of high school are included in this measure. Lastly, we look at the number of semester credit hours a student enrolled in up to five years after the expected high school completion date.

Earnings

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. The use of administrative earnings data presents advantages and drawbacks relative to more commonly used survey data such as the CPS or NLSY. A strength of administrative data is their accuracy. 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 estimates of the GED premium were based on similar types of data, it is straightforward to compare our estimates with those. Although previous estimates of the high school diploma premium were based mainly on survey data, we can use our data to compare discontinuity-based estimates of the diploma premium with estimates based on the approach taken in this literature. This can help reveal whether differences between our estimates and the estimates in that literature are driven by differences in approach or differences in data used. 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.

The Last Chance Sample

The main estimation sample covers all five cohorts and includes students who had not passed the exit exam by the final time it was offered in high school. We refer to this administration as the "last chance test".[19] Figure 1-T illustrates how this "last chance test" sample was constructed. 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.[20]

The first graph is based on data for all students that had valid scores on the initial test (nearly 100% of the sample). 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

[19] For the 1992-1995 cohorts, last chance test was a special test given in May or April to seniors who had yet to pass. No such test was available for the 1991 cohort so their "last chance" test is the exam administered in the spring of their senior year. Despite our labelling these as "last chance" tests, students could have re-taken the tests at a later date, typically by returning to high school for a thirteenth year of study. We label the tests "last chance" since they were students' last chance to pass in order to graduate on time and because many of the students that failed these tests did not retake them later.

[20] Figure 1-T is based on a subset of cohorts (1993-1995) since earlier cohorts initially took the test in 11th grade and had fewer retest opportunities. However, similar patterns (i.e., the emergence of a large discontinuity in the probability of graduating in the final retest offered in 12th grade) are observed for the first two cohorts.

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 also 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 students who took the exam in the summer after 10th grade. All of these students failed the initial attempt (that is, have a score less than zero on the first attempt) or did not take the test when it was first offered (e.g., because they were absent on the testing date). 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-T, 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-T 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-T 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 (**?**).

The characteristics that we observe are listed in the rows of Table 1-T. The columns present the means of these variables among the last chance sample, the number of observations associated with the last chance sample (around 42,000) 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. This confirms that these various observable characteristics are smooth through the last chance passing threshold.

5.2 Florida

The Florida data were obtained from the Florida Education and Training Placement Information Program (FETPIP), a data collection system that follows students from school into post-secondary education and into the labor market. As in Texas, these various datafiles were linked at the student level using identifying information such as the student's name, date of birth and social security number. In contrast to the Texas data, which we linked ourselves, the Florida data were linked by FETPIP staff.

We use information on students in the 2001-2004 tenth grade cohorts. We have earnings data out to 2009. Although this implies that the Florida samples cannot be observed as far into their careers as the Texas samples, it is interesting to compare the results across samples from different states and from different peiods of time. Since we obtain similar results for the samples, this suggests that our findings might generalize to other states and time periods. The Florida analysis focuses on a similar set of outcomes and the estimation sample is defined in the same way as it is for Texas: the set of students that sit for the "last chance test", the final test offered before the expected date of the end of high school.

Figure 1-F illustrates how this "last chance test" was constructed. The graphs are similar to those in Figure 1-T, the only difference being that there are slightly larger discontinuities in the probability of obtaining a high school diploma for the penultimate test. This may be because Florida students have slightly fewer chances to retake the test. Figure 1b-F shows that Florida last chance students also score towards the bottom of the initial test score distribution. Figure 1c-F shows that last-chance students' last-chance scores are smoothly distributed - evidence in support of our main identifying assumptions. Futher support for those assumptions comes from Table 1-F. This shows that Florida last chance students are, like Texas last-chance students, relatively disadvantaged. It also shows that at the last-chance passing threshold, there are no discontinuities in the observable characteristics listed in Table 1-F.

6 Results

We present results based on the Texas data and results based on the Florida data. We then assess the robustness of these two sets of results.

6.1 Texas

We present two sets of estimates for Texas: estimates of the impact of passing the last chance exam on the probability of obtaining a high school diploma and estimates of the high school diploma premium.

Diploma effects of passing the last-chance exam

We have already seen that passing the last chance exam is associated with a sharp increase in the probability of obtaining a high school diploma (Figure 1-T). We now investigate this relationship in more detail, using data on all cohorts used in the analysis of earnings. To that end, Table 2-T reports estimates of the diploma effects of passing the last-chance exam at various points in time relative to the last chance test. The first row shows the effect of passing the test on the likelihood of graduating within the summer after 12th grade. The estimates range from 0.43 to 0.51 depending on the regression specification. To determine the preferred specification, we examined AIC statistics, goodness of fit statistics and implemented a cross-validation procedure proposed by Imbens and Lemieux (2008). All of these, combined with visual examination of the quality of the regression fit, suggest that a nonlinear polynomial (fully interacted with passing status indicators) using the full range of data produces more reliable estimates than a linear function (unless a very narrow bandwidth is chosen). Therefore, the preferred estimates suggest a discontinuity of about 0.44. One year after the test, the discontinuity is about 0.38 (again the linear specification overstates the discontinuity), and remains stable at longer intervals (two and three years). 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. Graphs of the relationship between last-chance test scores and the probability of obtaining a diploma via an exemption and of the relationship between last-chance test scores and the probability of obtaining a diploma via a retake show that the exemption route is more common among those with low scores, while the retake route is more common among those with higher scores. 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, these exemptions should have little effect on the size of the credential wage premium. We assume that firms cannot observe the number of test administrations taken to pass the exam. This is consistent with the absence of a wage discontinuity associated with passing at the first attempt (not reported).

Earnings effects of receiving a high school diploma

With these "first stage" estimates in hand, we can use the last-chance sample to estimate the earnings effects of obtaining a diploma. Table 3-T presents the estimated earnings effects of obtaining the diploma between one and seven years after the last chance exam. As seen in the first row, mean earnings among this sample is increasing over this period. This will be reflect both an increase in average hourly wages and an increase in labor supply. The increase in labor supply will be driven in part by the addition of some college enrollees to the sample; average hourly wages would be expected to increase as all workers acquire more experience.

Columns (1)-(4) show reduced-form estimates using different polynomial specifications and column (6) shows estimates from a local linear regression where the bandwidth of 15 test score items is chosen by examining the cross validation function. These estimates are small in magnitude (ranging between -$200 and $200), statistically insignificant, and largely insensitive to the regression specification. Additionally, they show no obvious pattern across the seven year window. Column (5) shows estimates produced by an efficient estimator proposed by **?**. The estimates are a little larger and statistically significant in some years, but the estimates are small and statistically insignificant in the early years (years1 and 2) and the end of the observation period (years 6 and 7). Further, when we examine total earnings received in years 1-7 (final row), the estimates are all small (ranging from $57 to $800) and statistically insignificant across all specifications. The instrumental variables estimates of the effect of a high school diploma on earnings in columns (7) and (8) are a little more than twice as large (in magnitude) as the reduced-form estiamtes, and the effects remain statistically insignificant and display no obvious time profile.[21] To put these effect sizes into perspective, the point estimates are about 3 percent of mean earnings on either side of zero. In addition, the upper bound of the 95 percent confidence interval is never more than 14 percent of mean earnings and for earnings received over the full 7 year period is only 6.5 percent of mean earnings.

6.2 Florida

Estimates from Florida reveal similar patterns. The first stage estimates (Table 2-F) are of a similar magnitude (around 0.6), albeit slightly larger than the Texas estimates. The reduced-form earnings estimates (Table 3-F) are tiny, typically within 1 percent of zero and never statistically significant. These results appear robust to specification. When these estimates are scaled up by the first stage estimates (i.e., the instrumental variables estimates in the right-hand part of Table 3-F), they also suggest effects within five percentage points of zero. The similarity of the Florida

[21] For the IV estimates, we chose a quadratic polynomial in the test score but the IV estimates are similar using higher order polynomials in the test score (not shown). The AIC statistic for the reduced-form is actually minimized with a linear function in the test score. But, as noted above, a linear function overstates the magnitude of the first stage discontinuity.

and Texas results suggests that these findings likely generalize to other states and other types of exit exam regime. The similarity of findings across the two states is especially notable given that the Florida students left high school roughly ten years after the Texas students.

6.3 Robustness

Our estimates suggest that the high school diploma premium is small. Before discussing how these estimates should be interpreted, we assess how robust they are.

Downstream impacts

If diploma receipt affects downstream outcomes such as further education choices, then diploma premia cannot be interpreted through the lens of statistical discrimination and signaling.[22] We cannot assess diploma effects on every possible outcome, but we can assess diploma effects on further education outcomes, perhaps the most obvious and important downstream outcomes that could be affected. We examine diploma effects on the probability of enrolling in high school beyond the last chance test, the probability of pursuing and earning a GED, college enrollment and credit accumulation and college diploma receipt. Since we are especially interested in whether GED receipt can explain the zero sgnaling values that we estimate, and since we only have GED data for Texas, in most of what follows we discuss only those estimates based on the Texas data, reported in Table 4-T.[23] We do however produce graphs for Texas and Florida (Figures 4-F and 4-T). These are consistent in showing that passing the last chance test has little impact on college enrollment in both states.

In Table 4-T, the columns to the left of the vertical lines refer to outcomes for which we expect to find negative effects of passing the last chance test. The first outcome is being enrolled

[22] If passing or failing the exam has a direct effect (e.g., on pyschological well-being), then we cannot estimate the impact of the diploma even without downstream effects. That is because these psychological effects result in a violation of the exclusion restriction.

[23] These are again based on models that include a second-order polynomial in the last chance test scores. The estimates are quite similar when using other regression specifications.

in high school after the last chance exam. This should be more common among students that fail the last chance exam because retaking the exam is an additional reason why students may return to school.[24] This is exactly what we find (i.e., the effect is negative), although the effect is small and short-lived. This is consistent with the first-stage estimates, which did not change more than two years after the last-chance test. For the second outcome, receiving a GED, the effects are in the range of five percentage points - about 10 percentage points when scaled up by the first stage estimates. The fact that the effect of passing the test on GED acquisition is so much smaller than the effect of earning a high school diploma suggests that GED acquisition is not nearly enough to explain why we find little evidence of a signaling premium. This conjecture was further supported by an analysis we did using our data to generate observational estimates of the GED earnings premium.[25] We then combine these estimates with estimates of the GED effects of obtaining a high school diploma. The conclusion is that downstream GED effects might bias down the signaling value of the diploma premium by around one percentage point - a very small part of any story for why we find such small effects.

The remainder of Table 5-T shows results for college outcomes. We find clear evidence that college enrollment in Year 1 is affected by passing the test, but this effect disappears by Year 2. Furthermore, the estimated effects on college attainment, measured either using college academic credits or completion of a college degree, are very small and statistically insignificant. Notice that in Figures 4-T and 4-F, among the group close to the passing threshold, enrollment rates seem relatively high while credit completion is low. This suggests that many of these students attend part-time or dropout early.[26] This is a likely explanation for why we find no effects on college attainment despite finding evidence of a temporary effect on college enrollment.

[24] Students may also return if they were away from school because of illness or suspension.

[25] These estimates are based on samples of workers that failed the last-chance test and did not obtain a diploma. Specifically, they are obtained from models that regress earnings on two dummy variables: one for whether a GED was attempted unsuccessfully and one for whether a GED was attempted successfully. The omitted category are workers that did not attempt a GED. Our estimates suggest that workers that are successful in a GED attempt earn around 10 percent more than workers that are unsuccessful in a GED attempt.

[26] It is worth remembering that the last chance sample is one with low academic ability, at least as measured by the initial exit exam score.

Alternative outcomes

We have estimated the impact of dipoma status on earnings. To do this, we used the full sample of individuals and assigned zero earnings to workers not observed in the earnings data. It is possible that the small effects that we find on this outcome mask larger effects on one of those two underlying outcomes: the probability of having positive earnings and earnings conditional on having positive earnings.[27]

It is straightforward to test whether diploma status affects the probability of having positive earnings. It is harder to test whether diploma status affects conditional-on-positive earnings. One approach would be to repeat the analysis performed above but restrict to the sample observed to have positive earnings. Yet these estimates will be biased if diploma status impacts the probability of reporting positive earnings. Without auxilliary assumptions on the selection process into positive earnings (e.g., passing the test affects selection into this sample in a monotonic fashion for everyone, Lee (2009)), we cannot determine whether it does. Moreover, it is not clear in which direction any biases would go.[28] In Tables 4-F and 4-T we report estimates of the impacts of a high school diploma on the probability of having any earnings and on the impacts on earnings among those with positive earnings. For Florida, across both outcomes, the estimates are small, positive and statistically indistinguishable from zero. For Texas, there are some small negative effects on the probability of having positive earnings between years 3-5 and some small positive effects on conditional-on-positive earnings in years 3-5. These two phenomena could well be related, if there was negative selection out of the labor force among workers with a diploma. It is notable that by year 7, diploma status has no significant impact on either outcome.

[27] They could mask large effects on both of those outcomes, although those effects would have to be working in opposing directions. For example, a diploma would have to increase the probability of having positive earnings but decrease conditional-on-positive earnings. It would be hard to rationalize that combination of effects.

[28] In the classic selection scenario, we might expect estimates on conditional-on-positive earnings to be biased downwards. That is because positive diploma effects on the probability of having positive earnings will attract lower-quality diploma holders into the labor force. On the other hand, if diploma status increases the probability of attending college, diploma status may have negative effects on labor force status and may take lower-productivity diploma holders out of the labor force.

Zero earnings

To this point we have 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. For example, some individuals will be out of the earnings data because they work for the federal government (e.g., individuals in the military). Others 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 versions of the paper we will 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 individuals cannot join the military unless they have a high school diploma. There is no way of knowing how many of our sample joined the military. However, the military also screens on AFQT and typically selects very few individuals below the 31st percentile in the national distribution of AFQT (Angrist, 1998). Given that individuals in the last chance sample who score near the passing cutoff were at about the 15th percentile of the initial test score distribution, it is unlikely that they would meet the military's eligibility standards unless they did much better on the AFQT than they did on the exit exam. Furthermore, we find small effects for women as well (not reported), which pushes against a military-based explanation since women are also much less likely to enlist in the military than men.

7 Interpretation and Discussion

Our findings suggest that the signaling value of a high school diploma is small. Across both states and across various subgroups, we could not reject the hypothesis that the signaling value was zero; our preferred estimates allowed us to reject signaling values in excess of 10%. In this section we discuss what might explain these findings and consider how generalizeable they are likely to be.

7.1 What might explain these findings?

Recall that in the theoretical framework outlined above, there were three main reasons why the signaling value of a credential might be small.[29] First, receipt of a high school diploma might be costly to verify. Second, firms might observe exit exam scores. Third, firms might observe other productivity signals such that school performance information is redundant. In what follows we discuss the evidence relating to each of these explanations.

Is a high school diploma verifiable?

Under the Federal Education Records and Privacy Act (FERPA), a school can release "directory information" on names, dates attended and degrees earned without the written consent of the applicant or employee.[30] They are, however, under no legal obligation to do so unless the applicant signs a form requesting that this information be released.[31] Even then, **?** provides evidence suggesting that in many cases this information is either not provided or provided too late to be of use to a hiring firm.[32]

It is difficult to quantify the costs of verifying a diploma. It should be easier to verify the extent

[29] Another possible explanation is that there is no correlation between school performance and productivity. This is simple to dismiss, since all graphs of the relationship between earnings and last-chance exam scores sloped upwards. Provided wages reflect productivity, this implies that performance and productivity are positively correlated.

[30] www.ed.gov/offices/OM/fpco/ferpa/index.html

[31] In principle, firms could ask applicants to include the diploma as part of the application packet. Presumably though, a worker could claim to have lost his diploma. Since a worker could also forge a diploma, the firm would still need to verify this information.

[32] According to Bishop (1989), Nationwide Insurance, "one of Columbus, Ohio's most respected employers", obtained permission to get all high school records for its applicants. It sent over 1,200 such records to high schools in 1982 and received only 93 responses.

to which diplomas are requested and verified, but even this is problematic. First, whether or not diplomas are checked is an equilibrium outcome that depends also on the benefits of verifying this information. Second, data on the extent to which diplomas are requested and checked is difficult to come by.

One source of information is surveys of firms that have recently filled a position. These often ask whether a worker employed in that position is required to have a high school diploma. Firms often respond that they are. But this does not imply that an applicant without a high school diploma would not have been considered. In the surveys analyzed by Holzer (1998), while 75% of firms say that a high school diploma is "necessary or strongly preferred", 80% claim they would hire welfare recipients or applicants with a GED. In the Bay Area Longitudinal Surveys (Maxwell (2004), which target firms that have hired into entry-level jobs, we calculate that 55% say that a high school diploma is required, but 70% say that experience can substitute for this requirement. On this question, a more reliable source of information might be the information contained in the advertisements posted by firms. Bertrand and Mullainathan (2004) collected information on a large sample of advertisements in Boston and Chicago newspapers. Only 1% of these specified that a high school diploma was required.[33,34,35]

Can firms observe exit exam scores?

To verify exit exam scores, a firm would require access to high school transcripts. Under FERPA, this cannot be obtained without the written consent of the applicant. As noted above, obtaining this information could be difficult and time-consuming. Since exit exam scores are

[33] According to Holzer (1998), newspaper advertisements are the second most common method of recruitment (after informal contacts and referrals).

[34] Only 11 percent of these adverts required that workers have any credential. Only 8.8 percent of these specified a high school diploma. We are currently assessing analyze online job postings to see whether this finding holds in other states and for adverts targeted towards low-skilled workers.

[35] One piece of our analysis may seem inconsistent with the claim that a high school diploma is costly to verify: our finding that workers that just failed the last chance test were more likely to pursue a GED. One explanation is that these workers wish to enroll in college. This would be rational if the costs of college enrollment were lower for students with a high school diploma or GED (e.g., because they do not then need to enroll in remediation classes) and if colleges require that students include the high school transcript as part of the application packet (as many do).

scale scores (as opposed to letter grades, percentages or GPA measures), it may also be hard to interpret this information.[36] Again, firms' decisions as to whether to check transcripts for these scores are an equilibrium outcome. Again, the evidence suggests that they are not checked. In particular, according to the BALS, less than 4% of hiring firms request high school transcripts. Moreover, if firms observed scores, they might request that applicants have scores in excess of certain thresholds. We were, however, unable to find any online job postings or resume postings that mentioned them.[37]

Do firms observe a lot of productivity information?

Anecdotal evidence suggests that firms might observe a lot of productivity information. This could come from a job interview, tests given by the employer and the worker's employment history. Information could also be provided in the process of referring a worker to a job. According to Holzer (1998) and an analysis of the BALS data, this is the most common means by which positions are filled. Analysis of the BALS data also reveals that references are required for more than 80% of the positions filled and that employment tests are given in around 22% of cases.

Summary

These three explanations need not be mutually exclusive. In particular, a combination of the first and third explanations may go a long way towards explaining our findings. If firms face non-trivial costs of verifying receipt of a high school diploma and if they have already observe reasonably precise signals of productivity, the benefits to be gleaned from verifying receipt of a high school diploma may be less than costs of verifying this information. This could explain why few firms appear to require a diploma and why few firms appear to request high school transctipts.

Our preferred explanation is that firms observe neither diploma status nor last chance test scores. The leading alternative is that they observe both. Under some assumptions, the evolution

[36] In Florida, the scale ranges from 100 to 500 in both reading and math.

[37] This is based on an analysis of a small number of job and resume postings on monster.com. For future drafts of this paper we plan to review a larger sample of postings.

of the correlation between wages and last chance test scores can distinguish between these two cases. If firms observe these scores at labor market entry, the correlation between scores and wages will not change as workers acquire more experience. If firms do not observe these scores at labor market entry and instead learn about worker's ability over time, then as workers acquire more experience, the wage-score correlation will strengthen. This test is in the spirit of Altonji and Pierret (2001) (see the Appendix for a more formal exposition).

The data strongly suggest that the correlation becomes much stronger as workers acquire more experience. Although estimates are not reported, the strengthening wage-score correlation can be seen in Figures 3-T and 3-F. Subject to the same caveats discussed by Altonji and Pierret (2001), this implies that firms do not observe exit exam scores upon labor market entry.[38]

7.2 External validity

We assess generalizabilty, or external validity, in several dimensions.

Other types of students

As noted in the context of our theoretical and empirical frameworks, the signaling values that we identify are a weighted average of type-specific signaling values. The weights are equal to the probabilities of being of a certain type (as observed by firms) conditional on being a complier with last chance test score close to the passing threshold (Imbens and Angrist (1994)). This implies that the signaling value that we estimate is local in two senses: to marginal students with high school performance close to the passing threshold and to compliers.

The signaling value for marginal students might not be the same as the signaling value for "good" types (i.e., with a lot of favorable productivity signals) or "bad" types (i.e., with less

[38] These include the assumption that information is public (i.e., available to the market) and the assumption that any productivity growth is "neutral" (i.e., the experience profile does not depend on exit exam scores). This would be violated if, for example, workers with better exit exam scores were more likely to receive training. We do not have any information on training, although it is somewhat reassuring that the Altonji and Pierret (2001) estimates are robust to the inclusion of training measures (see columns (2) and (4) of Table V).

favorable productivity signals). In general, it will also differ from the population-weighted average signaling value (i.e., where the weights are the fraction of workers that belong to each type). There are however, two reasons for thinking that marginal students are an interesting group. First, there is no obvious relationship between signaling value and type quality. This will depend on the joint distribution of productivity (π) and productivity signals (π^s). For example, if these variables are distributed biva riate normal, the signaling value will be largest for extremely good and extremely bad types.[39] Second, the signaling value for marginal students may be comparable to the signaling value for students whose high school graduation status is affected by interventions designed to increase high school graduation rates. In that sense, marginal students might be policy-relevant.[40]

The signaling value among compliers might differ from the signaling value among non-compliers. In this context, since nearly all of those that pass the last chance exam get a diploma, a complier is someone that would not have obtained a diploma had they failed the last chance exam - someone that is less likely to retake the exam after failing the last chance exam. A particular concern is that this group is unlikely to retake because they have least to gain from the high school diploma signal. These might be workers for whom a lot of productivity information already exists.[41]

Perhaps the most plausible story in support of this argument is that compliers are likely those that already have relatively stable jobs or relatively high earnings by the time they are eligible to

---

[39] For a variable x distributed normal with mean μ and variance σ^2, $D(a) = E[x|x > a] - E[x|x < a] = \frac{\sigma\phi(\alpha)}{\Phi(\alpha)[1-\Phi(\alpha)]}$ where $\alpha = \frac{a-\mu}{\sigma}$. The derivative $D'(a)$ is positive (negative) if $\alpha > 0$ ($\alpha < 0$). Hence $D(\alpha)$ becomes larger as the mean moves away from the truncation point. For variables (x,y) distributed bivariate normal, where y is the signal and x is the variable of interest, $f(x|y)$ is normal with conditional mean $\mu_x + r(y - \mu_x)$ where $0 < r < 1$ measures the signal to noise ratio. Suppose $\mu_x < a$. If $\mu_x < a < y^* < y$, where y^* is the value at which the conditional mean equals the truncation point, an increase in y will increase $D(a, y)$ by moving the conditional mean away from the truncation point. If $y < \mu_x$, a decrease in y will have the same effect. Hence signaling values will be largest at extreme values of the signal. The same is true when $\mu_x > a$.

[40] Suppose for example that a high school intervention increases school performance and increases high school graduation rates. Earnings effects could operate via productivity improvements or via the signaling value of obtaining a high school diploma. Since only marginal students will have their high school graduation status affected by this intervention, the signaling value of the high school diploma for these students should be comparable to the signaling values that we identify.

[41] In the theoretical framework, firms had imperfect information about the productivity of all workers. In an extended framework, one could imagine that firms would have more productivity information for some workers than others.

retake the last chance test. To test this hypothesis, we explored whether the impact of passing the last-chance exam was related to high school earnings and high school employment. The estimates suggest that it is not. For example, among those without any earnings in high school (around 35% of the last-chance sample), the impact of passing on the probability of obtaining a high school diploma is 0.57. Among those with positive earnings below the 75th percentile of the high school earnings distribution, the impact is only slightly higher, at 0.59. Among those in the quartile of the high school earnings distribution, the impact is 0.6. These estimates are consistent with the hypothesis that compliers are more likely to be those that already have some contact with the labor market, but these differences are extremely small. This suggests that the decision to retake may be based on other considerations. These could include the desire to attend college or permanent characteristics such as persistence. These permanent characteristics could be correlated with the productivity signals observed by firms, but since there is no obvious relationship between the quality of these signals and the signaling value of the diploma, we need not worry that, for example, the compliers might be "worse" than the non-compliers.

High school diplomas acquired in other settings

Roughly one half of all states now operate exit exams. Hence, even if our results were specific to these states, we think they would still be of interest.[42] Since we interpret our results as reflecting the costs of verifying high school diploma receipt and the availability of alternative productivity information, we believe that they will generalize to settings in which high school diplomas are not employed. We also believe they will generalize to states other than the two considered in this paper. That is because our results are remarkably consistent across both of the states considered and across all of the major subgroups within them.

It is interesting to consider whether we would obtain similar results if we could perform a similar

[42] Assuming that exit exam scores are more strongly correlated with productivity than the high school graduation requirements in states that do not operate these exams, we woulf expect the signaling value of a diploma to be *higher* in states that operate exit exams.

exercise on workers that graduated high school in the period before that covered by our data. This depends how the labor market for these types of workers has changed over this period. Goldin and Katz (2008) estimate that the elasticity of substitution between high school dropouts and high school graduates has been stable since 1950. To the extent that this is roughly proportional to the productivity difference between those with and without a diploma conditional on having twelve years of schooling, it suggests that the signaling value of the diploma might have been relatively stable over this period. Since Goldin and Katz (2008) estimate the elasticity of substitution to have been much lower before 1950, it is interesting to speculate as to whether the signaling value of a high school diploma would have been larger during this period.[43]

It is difficult to assess whether the signaling values that we estimate would also apply in ealier periods. The best we can do is compare "sheepskin" estimates of the signaling value of a high school diploma obtained in the earlier literature (using data from the 1970s and 1980s) with the estimates obtained when the same methods are applied to our data. These estimates are presented in Tables 5-F and 5-T. The column (1) estimates do not control for any worker characteristics, the column (2) estimates control for the types of variables controlled for in the previous literature, the column (3) estimates also include the score obtained on the initial test, a variable we expect to be strongly correlated with productivity. In the left-hand columns, we report estimates from regressions applied to the full sample of workers that completed twelve years of school but did not attend college (the sample used in the previous literature). In the right-hand columns we report estimates from regressions applied to the last-chance sample.

Two patterns are apparent in these tables. First, the estimates that correspond most closely to the previous literature (full sample, demographic controls) suggest large effects - over 25% in year 6 - that are comparable to those estimated in the earlier sheepskin literature. This exercise therefore provides some support for the claim that the signaling value of a high school diploma

[43] Interestingly, Goldin and Katz (2008) note that in this period "hiring employees described certain jobs as requiring a high school diploma or particular high school courses" (p.307).

may have relatively stable over the post-war period. It also confirms our conjecture that the estimates obtained in the earlier literature were biased upwards. Since we find effects of around 7% when using the last-chance sample and controlling for initial test scores, it is seems that one half of the sheepskin effect is eliminated by focusing on the last-chance sample with the remainder eliminated by focusing on the earnings discontinuity at the passing threshold.

Other types of education

If our explanation for our findings holds, the signaling value of other types of education will depend on the costs and benefits of verification. There is little reason to believe that the costs of verifying other types of education in the lower part of the education distribution (e.g., highest grade completed) will be smaller than the costs of verifying high school diploma receipt.[44]

One could also argue that the benefits of verifying these types of education will be no larger than the benefits of verying high school diploma receipt. That is because workers with these other types of education should be in roughly the same labor market as workers on the margin of obtaining a high school diploma (i.e., the market for non-college labor). It is possible that firms have less information about workers with lower levels of education (e.g., high school dropouts), such that there is a greater return to verifying types of education that distinguish among this group (i.e., a GED). This could explain some of the positive estimates found in the GED literature, although this literature has produced mixed results, as noted above. It is difficult to know whether these results generalize to types of education in the top part of the education distribution (e.g., college attended). This type of education might be easier to verify (e.g., colleges might devote more resources to providing this information). There might also be a stronger correlation between college performance and productivity.

[44] Some types of education may be even costlier to interpret (e.g., course offerings and GPA), since they might be school-specific hence hard to interpret.

8 Conclusion

It is difficult to determine whether various types of education act as signals of underlying ability. In this paper we estimate the signaling value of one type of education - a high school diploma. Using a strategy that identifies the signaling value of a high school diploma under much weaker assumptions than those made in the previous literature, we show that the signaling value of the diploma is small. We believe these findings reflect the cost of verifying receipt of a high school diploma and the existence of other productivity information that reduces the benefits of verification.

Although these estimates are local to the types of students on the margin of graduating high school, these are the students whose graduation status is likely affected by high school interventions designed to increase high school graduation rates. Moreover, we see no reason why these estimates should be larger than those found in non exit exam settings or in other states. It is difficult to know whether the signaling value of a high school diploma is smaller now than it was in the past, although it is notable that the earnings difference that we see between high school graduates with and without diplomas is similar to that seen in the previous literature. Since the costs and benefits of verifying high school diploma receipt should be comparable to the costs and benefits of verifying other types of lower-level education (e.g., high school GPA, highest grade completed), we suspect that our results are representative of the signaling value of other types of education at the lower end of the education distribution. While the GED may be an exception, since firms might have less information about high school dropouts, findings in the GED literature are mixed, and a relatively small fraction of workers hold a GED.

Economists care about education-based signaling because it implies that the private returns to education could exceed the social return to education, the parameter of interest relevant to education policy. Estimates of the private returns to education at the bottom end of the education distribution (e.g., those reviewed by Card (1999) are relatively high. The analysis in this paper suggests that the social returns to this education might be at least as high.

A Appendix

In this Appendix we provide details of the test discussed in the Discussion and Interpretation section of the text. This is a simplified version of the test proposed by Altonji and Pierret (2001) and is valid under similar assumptions. These are discussed below.

Suppose we express productivity π as a function of variables observed by the firm but not the econometrician (q_f), variables observed by the both the firm and the econometrician (q_b), variables observed by the econometrician but not the firm (q_e) and variables observed by neither (q_n):

$$\pi = \beta_0 + \beta_1 q_f + \beta_2 q_b + \beta_3 q_e + q_n$$

Now suppose we estimate the following equation:

$$W_t = c_{0t} + c_{1t} p + e_t$$

where p is the last-chance exam score, W_t is earnings in year t and the coefficients in this equation can vary over time. We can show that the coefficient on p in the above equation allows us to distinguish two cases: (1) p observed by the econometrician and the firm (hence belongs to q_b) (2) p observed by the econometrician but not the firm (hence belongs to q_e).

Case 1: Variable p observed by the econometrician and the firm

Suppose that p belongs to q_b. Suppose without loss of generality we can assume that this is the only variable in q_b and that q_e is empty. Then:

$$\begin{aligned}\pi &= \beta_0 + \beta_1 q_f + \beta_2 p + q_n \\ &= \widetilde{\beta_0} + \widetilde{\beta_1} q_f + \widetilde{\beta_2} p + v\end{aligned}$$

where the second line is obtained by substituting the linear projection of q_n on q_f and p in the first line (v is the projection error).[45]

When the worker enters the labor market, the market's best guess of the worker's productivity is $\widetilde{\beta_0} + \widetilde{\beta_1} q_f + \widetilde{\beta_2} p$. In every subsequent period, the market observes productivity realizations $\widehat{\pi}_t = \pi + \varepsilon_t$, where ε_t is a mean-zero random variable that captures the noise in these realizations. This gives the firm additional information with which to predict v. Let $\mu_t = v - E[v|t]$ be a measure of how well the firms predict v in period t. Substituting into the above equation, taking expectations and noting that $E[\pi|t] = W_t$ gives:

$$W_t = \widetilde{\beta_0} + \widetilde{\beta_1} q_f + \widetilde{\beta_2} p + E[v|t]$$

where the final term captures how wages change as firms learn about v. We can write this equation as a function of p using the projections of q_f on p and $E[v|t]$ on p (denoting the projection coefficients Φ_{fp} and Φ_{vp}). This implies that in the wage equation:

$$c_{1t} = \widetilde{\beta_1}\Phi_{fp} + \Phi_{vp} = \widetilde{\beta_1}\frac{Cov(q_f, p)}{V(p)} + \beta_2 \frac{Cov(v, p)}{V(p)}\theta_t$$

where $\theta_t = \frac{Cov(E[v|t], p)}{Cov(v,p)}$ summarizes firms' knowledge of v at time t. Since $Cov(v, p) = 0$, c_{1t} is constant over time.

Case 2: Variable p observed by the econometrician but not the firm

[45] To derive this equation, we substitute $q_n = \alpha_0 + \alpha_1 q_f + \alpha_2 p + v$ into the earlier equation.

Suppose p belongs to q_e. Suppose without loss of generality that this is the only variable in q_e and that q_b is empty. Then:

$$\begin{aligned} \pi &= \beta_0 + \beta_1 q_f + \beta_3 p + q_n \\ &= \widetilde{\beta_0} + \widetilde{\beta_1} q_f + (\beta_3 u + v) \end{aligned}$$

where the second line is obtained by substituting the linear projections of p and q_n on q_f into the first line (u and v are the projection errors). As above, we can derive:

$$W_t = \widetilde{\beta_0} + \widetilde{\beta_1} q_f + E[\beta_3 u + v|t]$$

and an expression for c_{1t}:

$$c_{1t} = \widetilde{\beta_1}\Phi_{fp} + \frac{Cov(\beta_3 u + v, p)}{V(p)}\theta'_t = \widetilde{\beta_1}\Phi_{fp} + \beta_3 \frac{V(u)}{V(p)}\delta_t$$

where $\delta_t = \frac{Cov(E[\beta_3 u+v|t],p)}{Cov(\beta_3 u+v,p)}$ summarizes firms' knowledge of $(\beta_3 u + v)$ at time t. If firms cannot predict p at labor market entry (i.e., $V(u) > 0$) and if they learn over time (i.e., $\delta'_t > 0$), then c_{1t} will increase with time in the labor market.

Summary

We have shown that under some assumptions, a regression of wages on last-chance test scores can distinguish between the case in which the firm observes the last-chance score and the case in which it does not. These assumptions are first, that firms learn about productivity as workers accumulate labor market experience (and learning is public) and second, that any experience profile in earnings is "neutral" with respect to last chance test scores. In that case, it is straightforward to show that the preceding analysis is unaffected by the addition of a neutral experience profile to the productivity equation. This will be captured by the (time-varying) constant in the wage equation. See Altonji and Pierret (2001) for more discussion of these assumptions.

References

Altonji, J., and C. Pierret (2001): “Employer learning and statistical discrimination,” *Quarterly Journal of Economics*, 116, 313–350.

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

Bertrand, M., and S. Mullainathan (2004): “Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination,” *American Economic Review*, 94(4), 991–1013.

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.

Card, D. (1999): “The Causal Effect of Education on Earnings,” in *Handbook of Labour Economics Volume 3A*, ed. by O. Ashenfelter, and D. Card. Elsevier, Amsterdam.

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.

Farber, H., and R. Gibbons (1996): “Learning and wage dynamics.,” *Quarterly Journal of Economics*, 111, 1007–47.

Goldin, C., and L. Katz (2008): *The Race Between Education and Technology*. Harvard University Press, Cambridge, MA.

Holzer, H. (1998): *What Employers Want: Job Prospects for Less-Educated Workers*. Russell Sage Foundation, New York.

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.

Imbens, G. W., and T. Lemieux (2008): “Regression discontinuity designs: A guide to practice,” *Journal of Econometrics*, 142(2), 615–635.

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. (2009): "Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects," *Review of Economic Studies*, 76, 1071–1102.

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.

Martorell, P. (2010): "The Impact of High School Exit Exams on High School Graduation and Persistence: Evidence from Texas," RAND, mimeo.

Maxwell, N. L. (2004): "The Bay Area Longitudinal Surveys of Firms," HIRE Center Discussion Paper D04-06-04.

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

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.

Weiss, A. (1983): "A Sorting-cum-]Learning Model of Education.," *Journal of Political Economy*, 91, 420–42.

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

Figure 1a-F: Fraction awarded HSD, by score on exit exam attempt

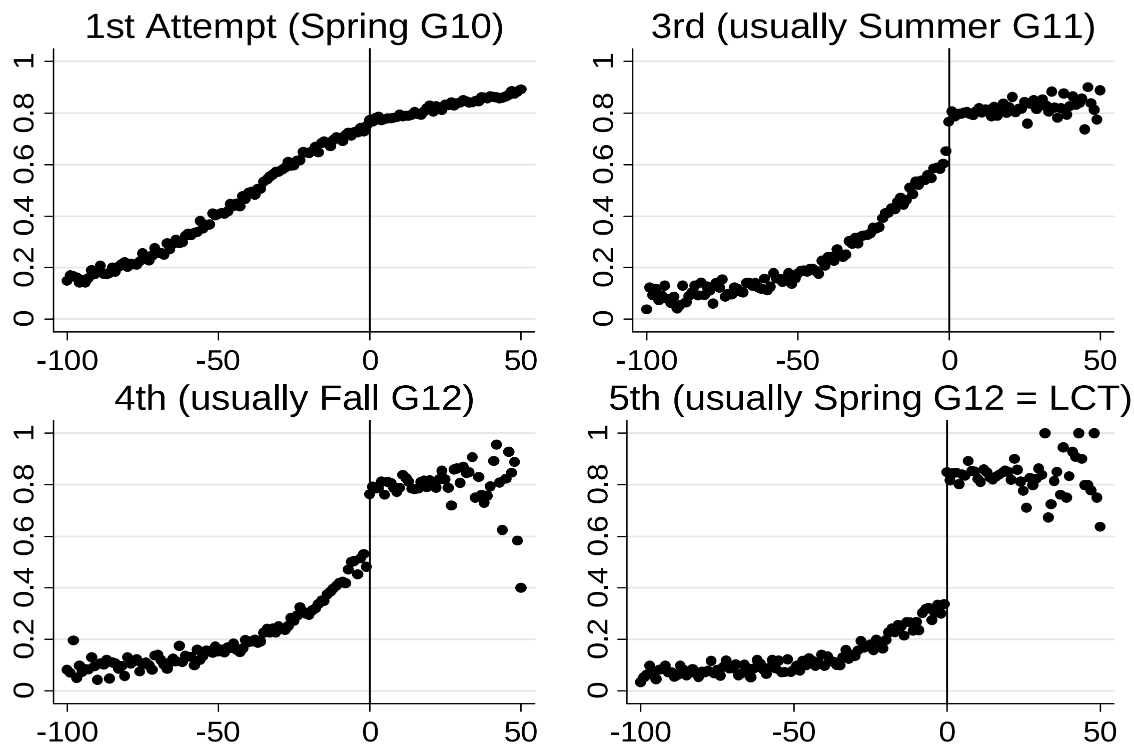

Notes: The sample underlying these figures includes all of those in grade 10 in 2001-02 (170,597) who took

Figure 1b-F: Distribution of first-attempt scores for full sample and last-chance sample

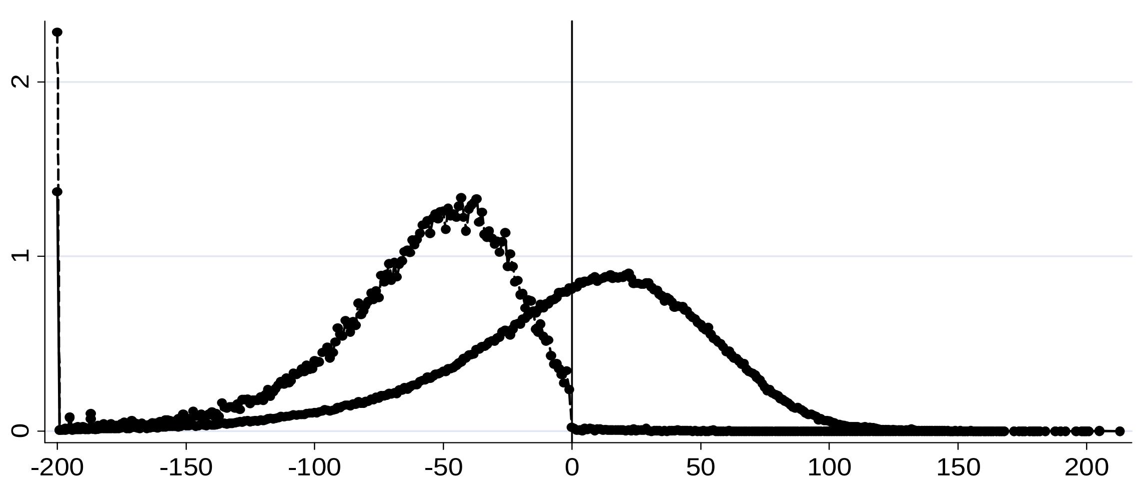

Notes: The samples underlying the two densities are (a) all of those observed to take at least one part of the test once (b) the subset of those that are in the last-chance sample (see Table 1).

Figure 1a-T: Fraction awarded HSD, by score on exit exam attempt

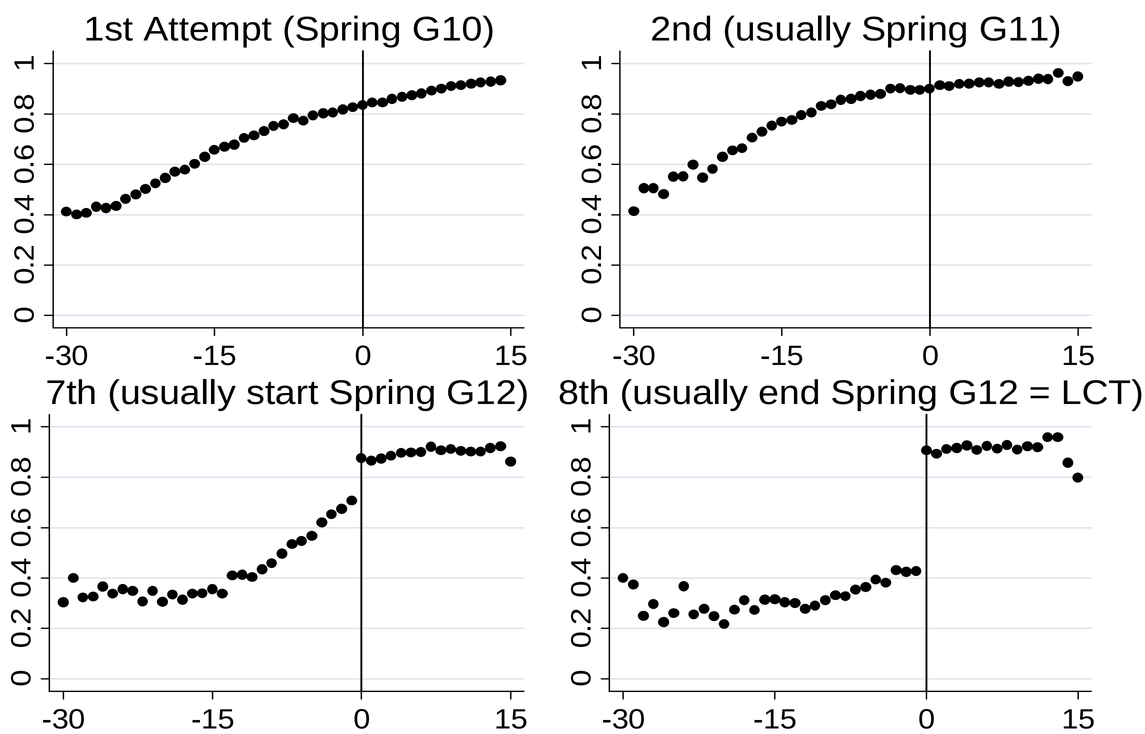

Notes: see notes to Figure 1a-F

Figure 1b-T: Distribution of first-attempt scores for full sample and last-chance sample

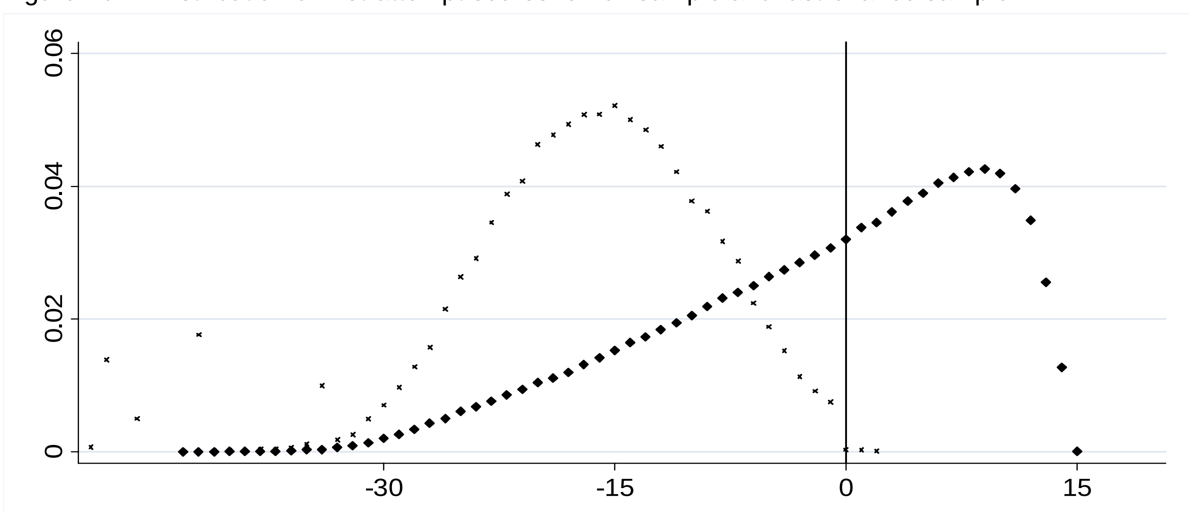

Notes: see notes to Figure 1b-F

Figure 1c-F: Distribution of last-chance scores (among last-chance sample)

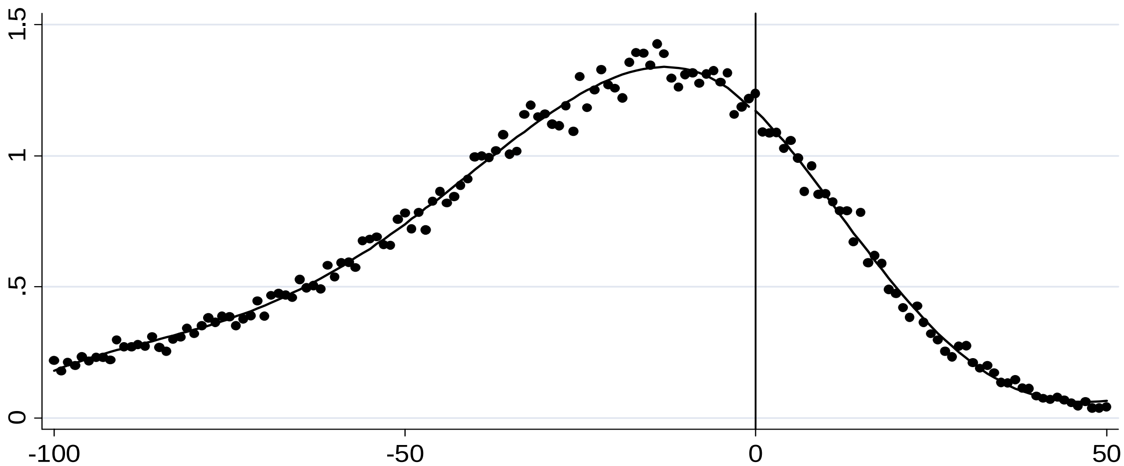

Notes: the figure is based on the last-chance sample (see Table 1). The lines are fitted values from a regression of cell counts on a fourth-order polynomial in the last-chance score, interacted with pass dummy variable.

Figure 1c-T: Distribution of last-chance scores (among last-chance sample)

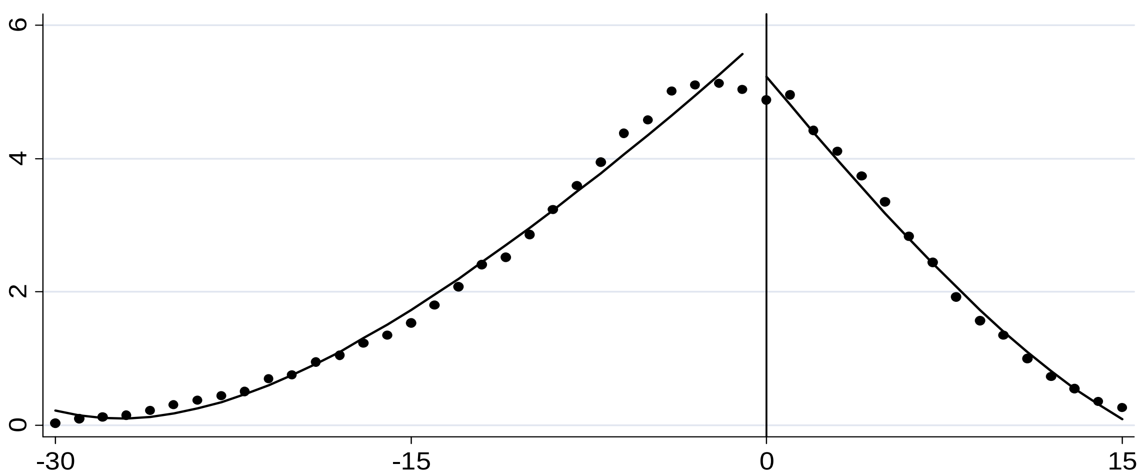

Notes: see notes to Figure 1c-F

Figure 2-F: Certificates awarded, by last-chance scores

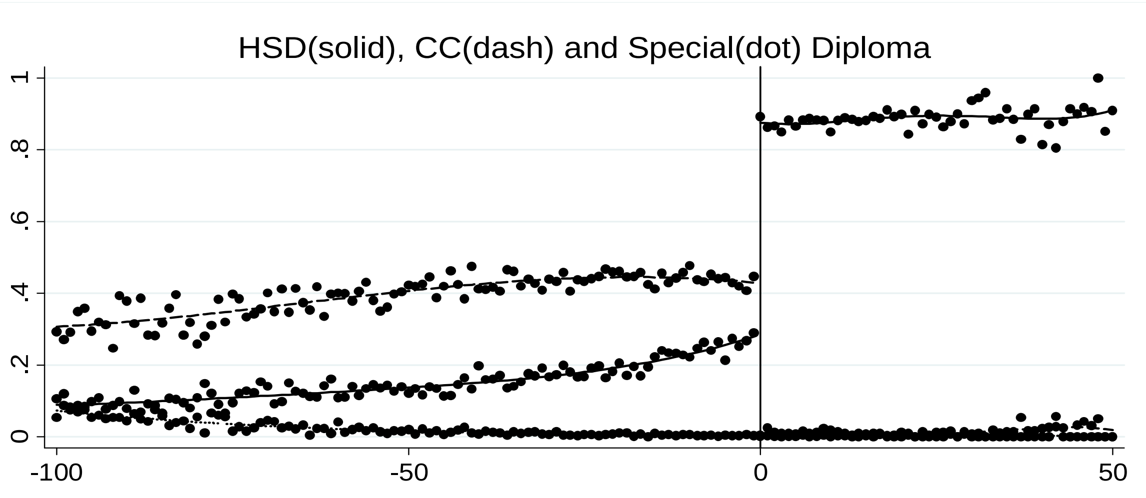

Notes: lines are fitted values from a regression of these outcomes on a fourth-order polynomial (interacted) in the last-chance test score. "HSD", "CC" and "Special" refer to high school diplomas, certificates of completion and special high school diplomas. Sample is last-chance sample (see Table 1).

Figure 2-T: Certificates awarded, by last-chance scores

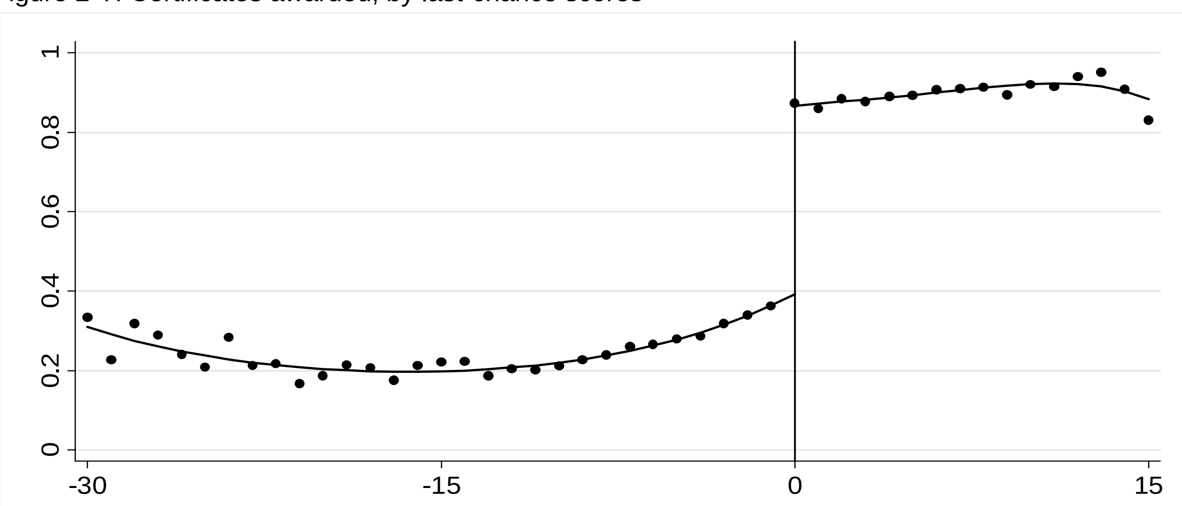

Notes: see notes to Figure 2-F

Figure 3a-F: Earnings in years after last-chance exam, by last-chance scores

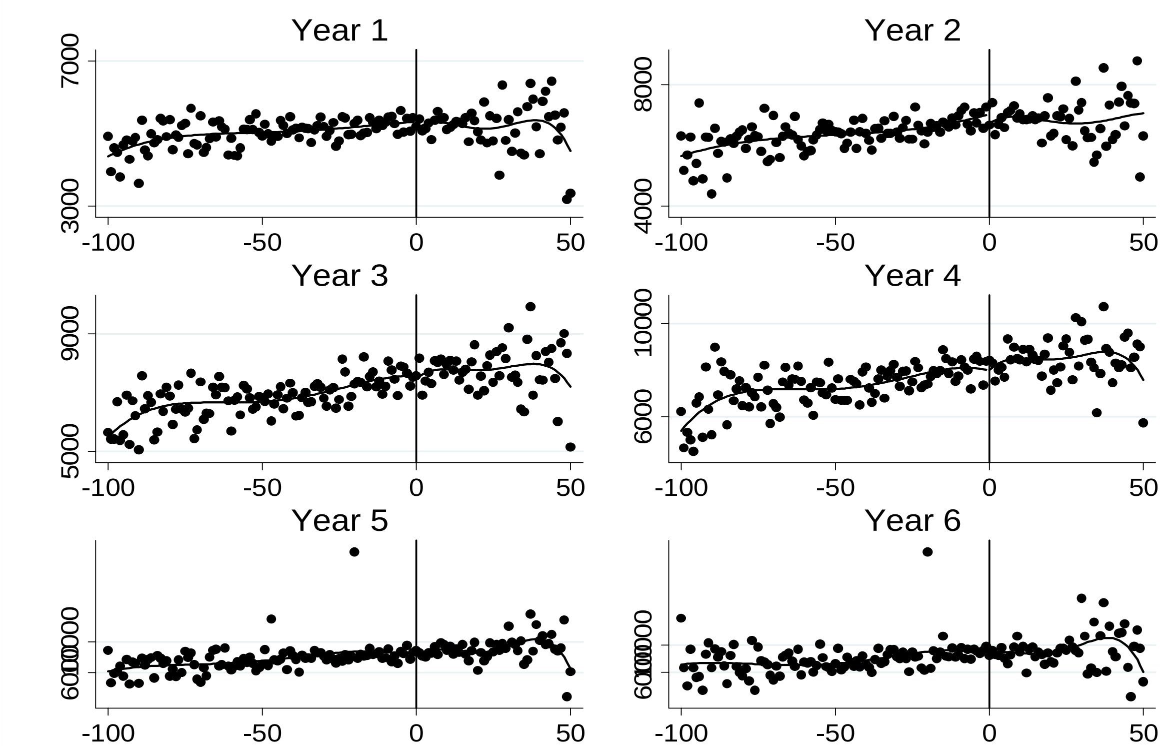

Notes: lines are fitted values from a regression of these outcomes on a fourth-order polynomial (interacted) in the last-chance test score. Sample is last-chance sample (see Table 1).

Figure 3b-F: Cumulative earnings up to six years after last-chance exam, by last-chance scores

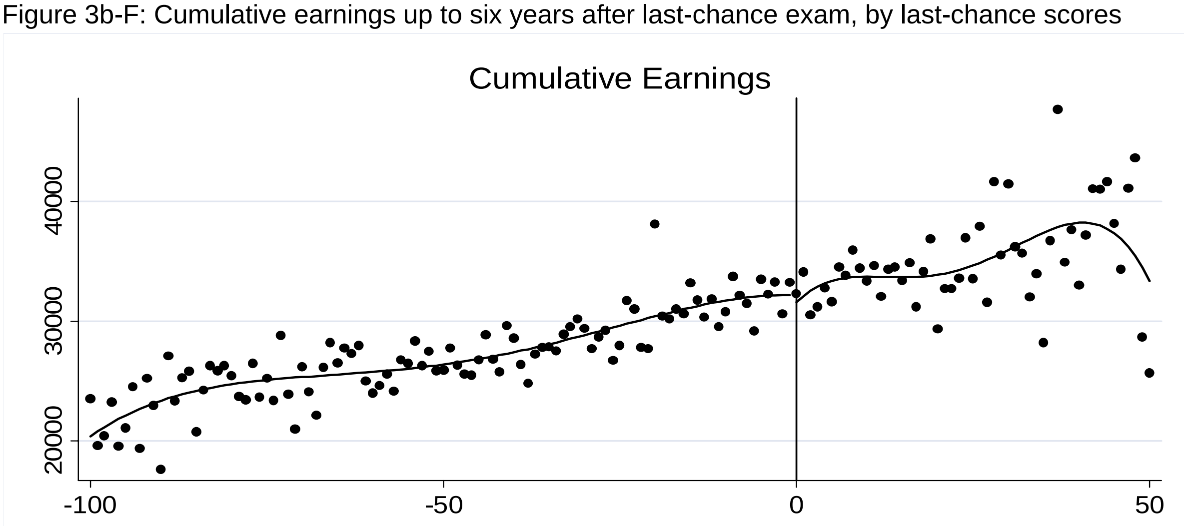

Notes: lines are fitted values from a regression of this outcome on a fourth-order polynomial (interacted) in the last-chance test score. Sample is last-chance sample (see Table 1).

Figure 3a-T: Earnings in years after last-chance exam, by last-chance scores

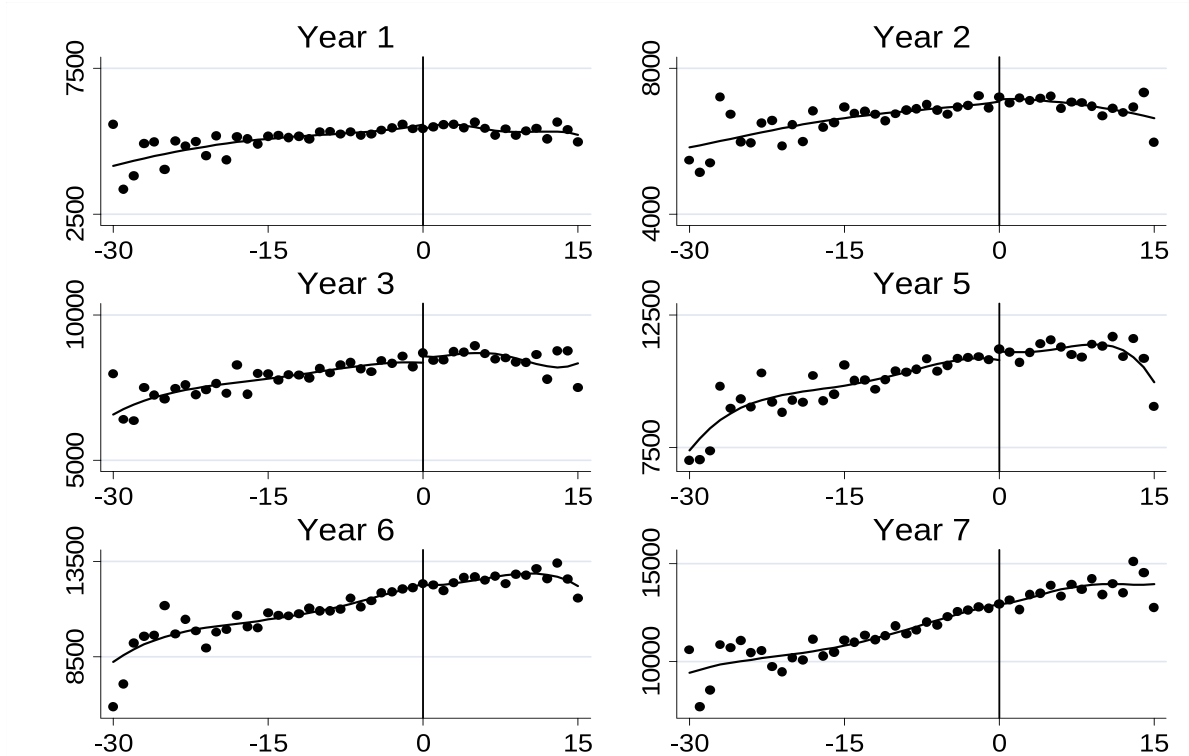

Notes: see notes to Figure 3a-F

Figure 3b-T: Cumulative earnings up to six years after last-chance exam, by last-chance scores

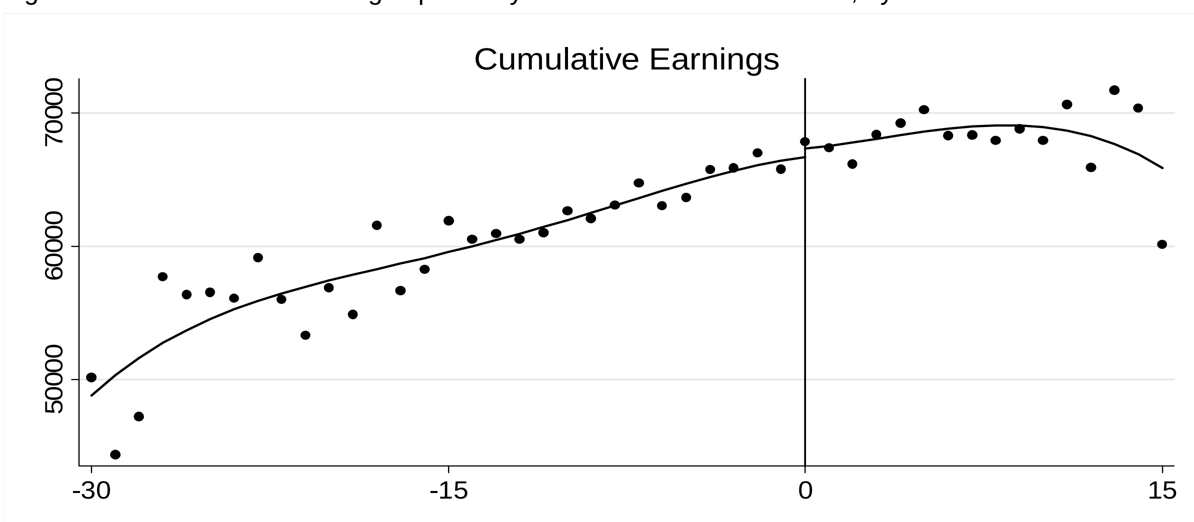

Notes: see notes to Figure 3b-F

Figure 4-F: Post-secondary enrollment, by last-chance scores

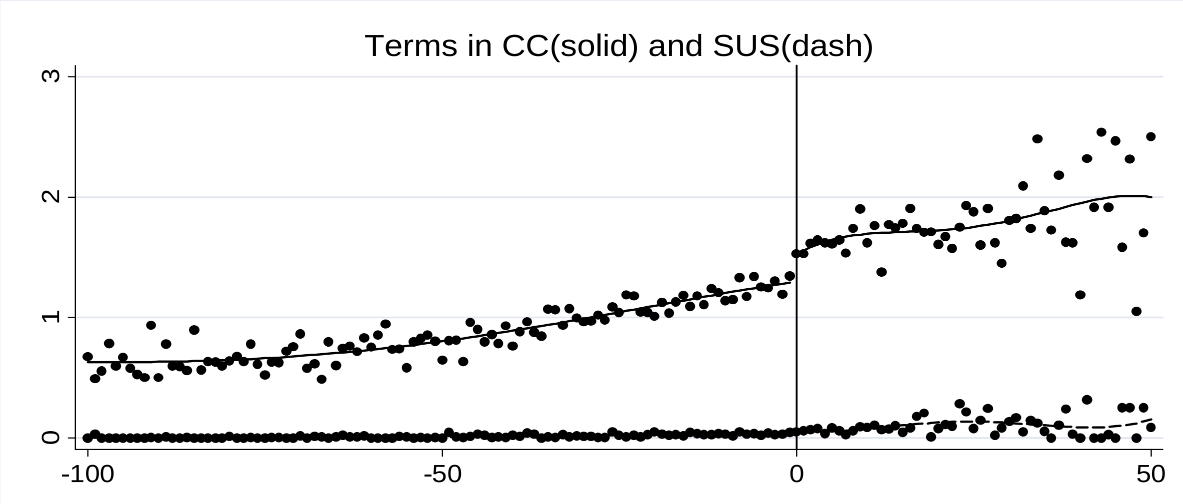

Notes: lines are fitted values from a regression of this outcome on a fourth-order polynomial (interacted) in the last-chance test score. "CC" and "SUS" refer to community colleges and state univeristy system colleges. Sample is last-chance sample (see Table 1).

Figure 4-T: GED Certification and college credits, by last-chance scores

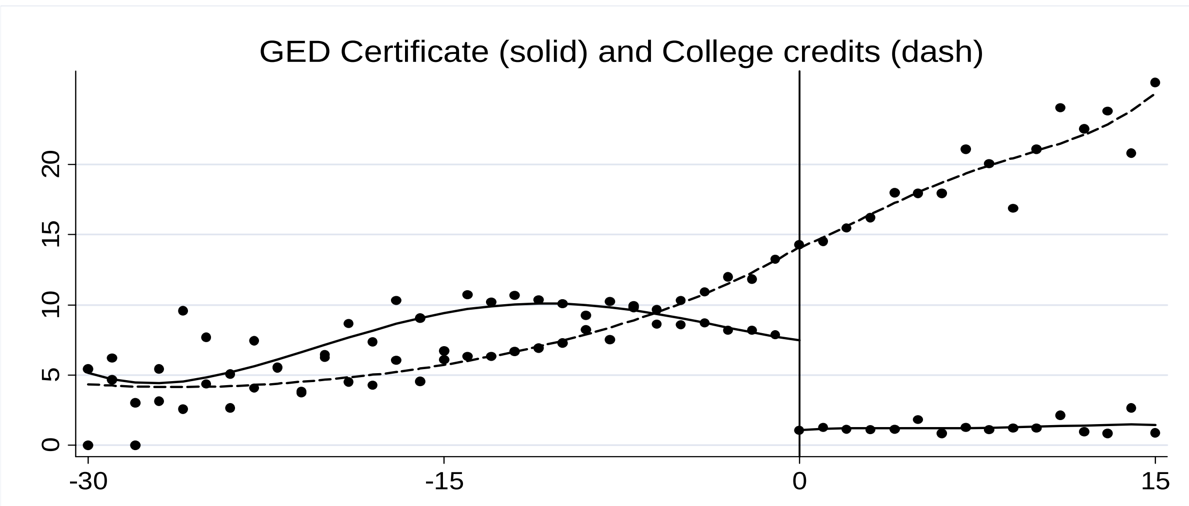

Notes: see notes to Figure 4-F

Figure 5a-F: Probability of positive earnings, by last-chance scores

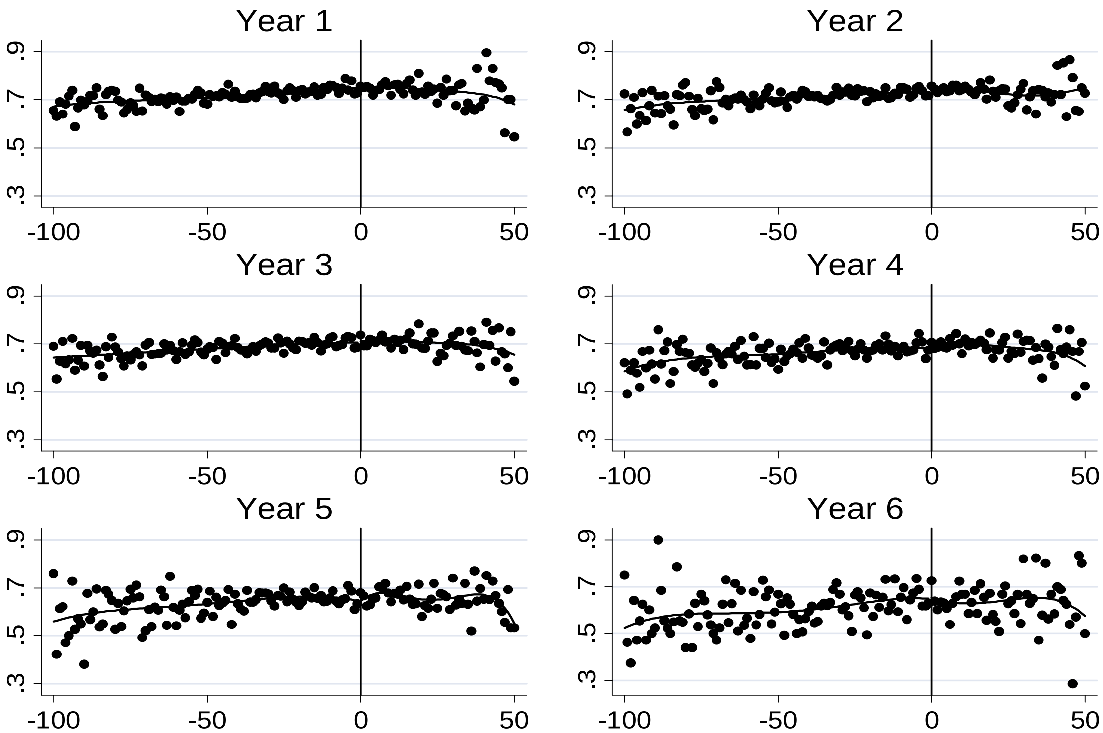

Notes: lines are fitted values from a regression of this outcome on a fourth-order polynomial (interacted) in the last-chance test score. Sample is last-chance sample (see Table 1).

Figure 5b-F: Conditional-on-positive earnings, by last-chance scores

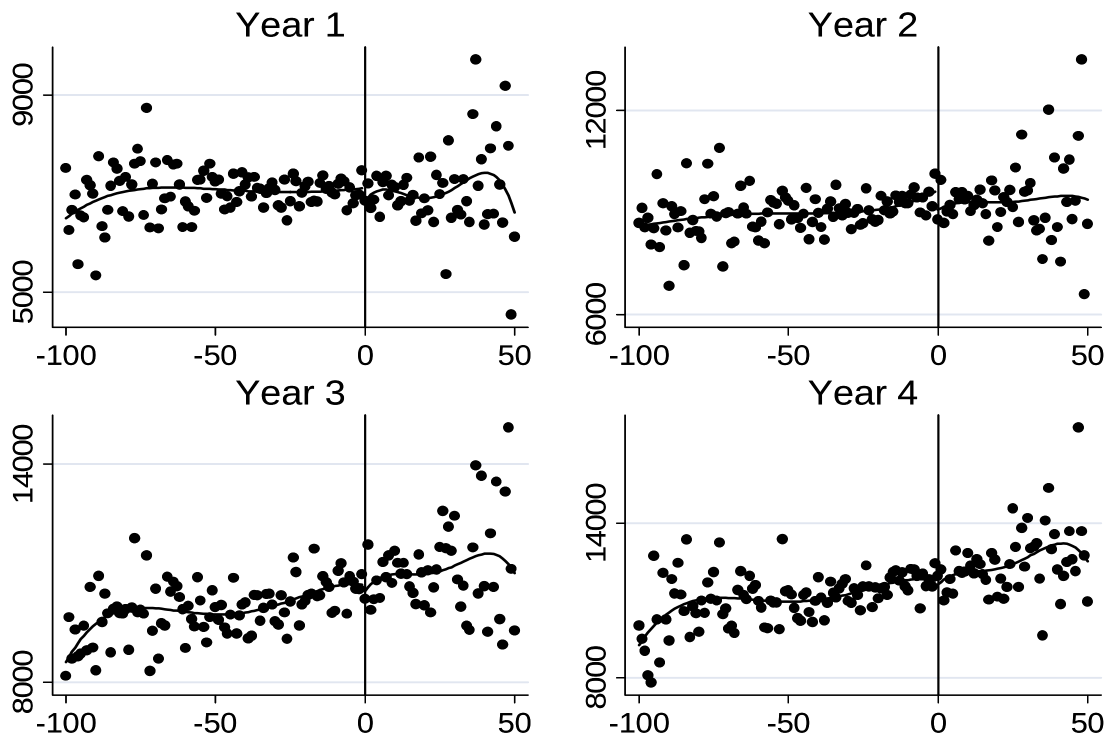

Notes: lines are fitted values from a regression of this outcome on a fourth-order polynomial (interacted) in the last-chance test score. Sample is last-chance sample (see Table 1).

Figure 5a-T: Probability of positive earnings, by last-chance scores

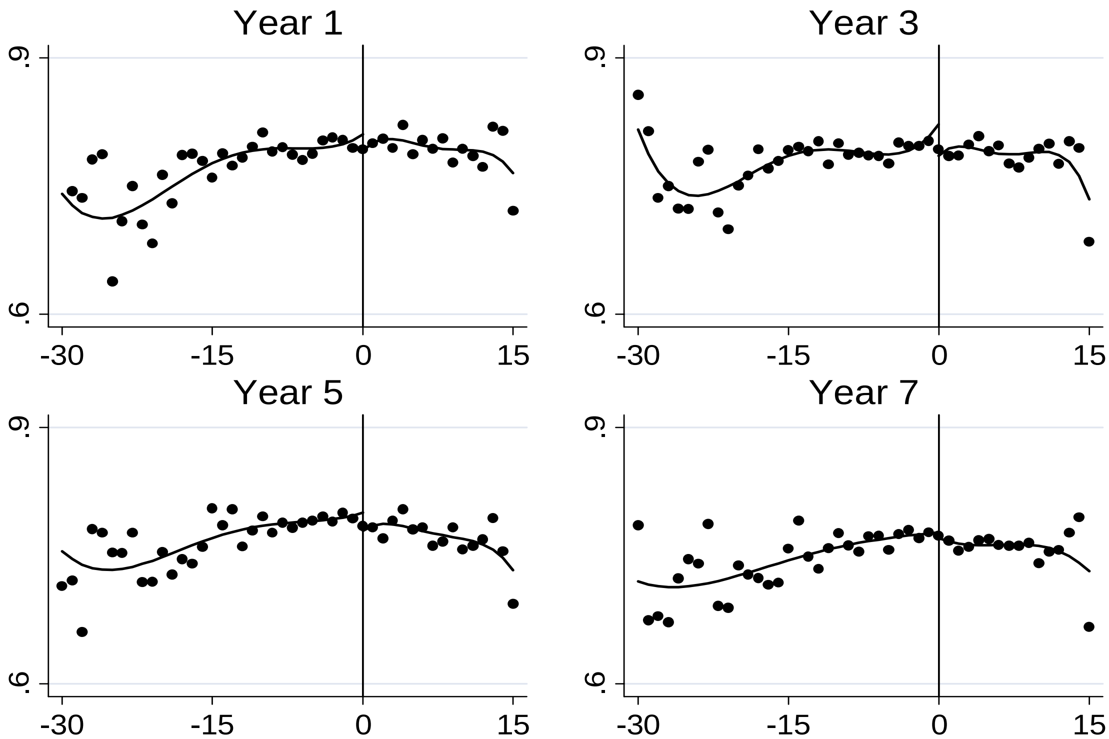

Notes: see notes to Figure 5a-F

Figure 5b-T: Conditional-on-positive earnings, by last-chance scores

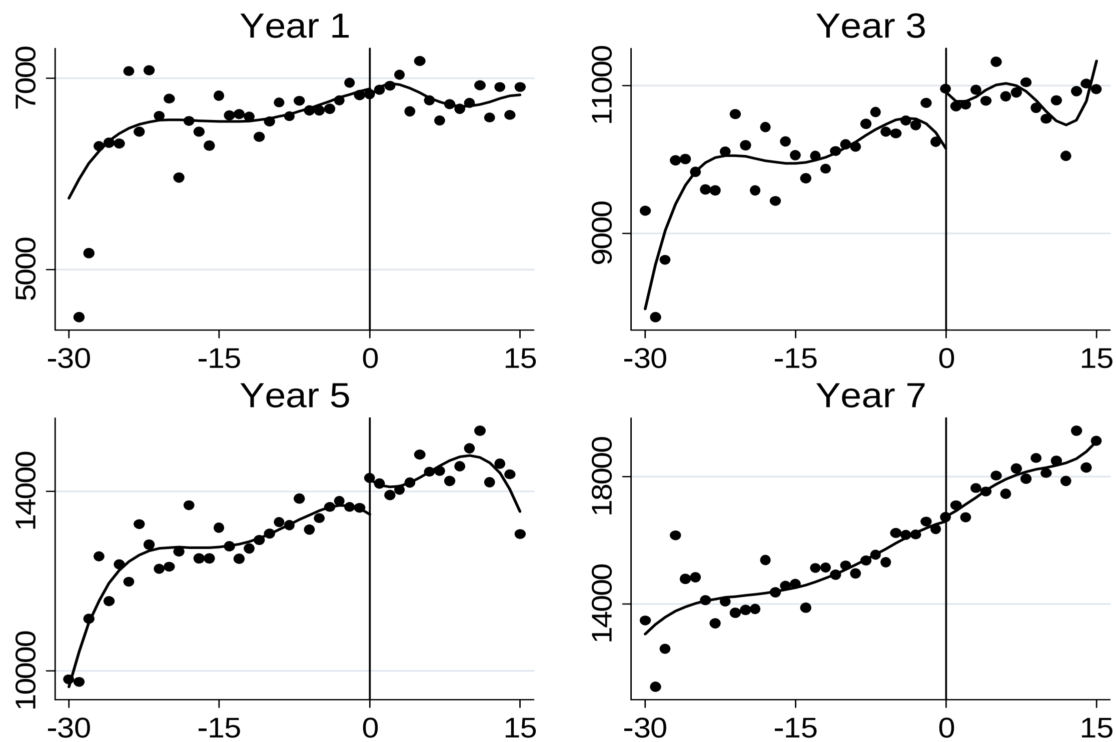

Notes: see notes to Figure 5b-F

Table 1-F: Descriptive Statistics

	All		Fail	Pass	Discontinuity	
	Mean	N	Mean	Mean	Coeff	SE
10th grade in 1999	0.18	51638	0.17	0.22	-0.009	-0.011
10th grade in 2000	0.27	51638	0.25	0.34	0.015	-0.013
10th grade in 2001	0.26	51638	0.26	0.25	0.002	-0.017
10th grade in 2002	0.29	51638	0.32	0.18	-0.008	-0.010
First attempt - math score	-24.59	50934	-28.23	-12.41	0.461	-1.243
First attempt - reading score	-55.52	51077	-60.82	-37.79	-1.947	-1.668
First attempt - math missing	0.01	51638	0.01	0.01	-0.007	-0.009
First attemt - reading missing	0.01	51638	0.01	0.01	0.007	-0.008
Male	0.44	51638	0.44	0.45	-0.003	-0.056
Black	0.44	51638	0.46	0.39	0.000	-0.032
Hispanic	0.27	51638	0.27	0.27	0.015	-0.022
Free or reduced price lunch	0.48	51638	0.50	0.42	0.005	-0.051
Limited english proficient	0.34	51638	0.35	0.31	-0.001	-0.033
Special education	0.20	51638	0.23	0.13	-0.031	-0.025

Notes: statistics are for the last-chance sample, defined as those that took the test for the first time in 10th grade, were enrolled in 12th grade two years later, took the test again in the spring of their 10th grade and scored between -100 and 50 (normalized relative to the passing cutoff) on this test. Discontinuity column refers to coefficient and standard error on pass dummy in a regression of the characteristic on as last-chance pass dummy and a fourth-order polynomial (interacted) in the last chance score.

Table 1-T: Descriptive statistics

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

Notes: statistics are for the last-chance sample, defined as those that took the test for the first time in 10th grade, were enrolled in 12th grade two years later, took the test again in the spring of their 10th grade and scored between -30 and 15 (normalized relative to the passing cutoff) on this test. Discontinuity column refers to coefficient and standard error on pass dummy in a regression of the characteristic on as last-chance pass dummy and a fourth-order polynomial (interacted) in the last chance score.

Table 2-F: Estimated impact of passing last change test on probability of earning high school diploma

	Poly-1	Poly-2	Poly-3	Poly-4	Poly-1 bw=25
	Receive diploma before summer of 12th grade (N=51,638)				
Pass	0.632	0.601	0.590	0.591	0.592
	(0.007)	(0.010)	(0.012)	(0.013)	(0.010)
	Receive diploma within 1 year of LCT (N=51,638)				
Pass	0.627	0.593	0.580	0.580	0.584
	(0.008)	(0.009)	(0.011)	(0.012)	(0.009)
	Receive diploma within 2 years of LCT (N=51,638)				
Pass	0.626	0.592	0.581	0.581	0.584
	(0.007)	(0.009)	(0.011)	(0.012)	(0.009)

Notes: Column heading Poly-X implies that estimates in this column are based on a regression of the outcome on a X'th order polynomial (interacted) in the last-chance score. "Bw=25" implies that the estimates are based on a subset of those in the last-chance sample that scored within 25 points of zero. Clustered (on the last-chance score) standard errors in parentheses.

Table 2-T: Estimated impact of passing last change test on probability of earning high school diploma

	Poly-1	Poly-2	Poly-3	Poly-4	Poly-1 bw=15
	Receive diploma before summer of 12th grade (N=51,638)				
Pass	0.507	0.440	0.439	0.427	0.482
	(0.015)	(0.012)	(0.007)	(0.011)	(0.013)
	Receive diploma within 1 year of LCT (N=51,638)				
Pass	0.448	0.380	0.389	0.379	0.422
	(0.013)	(0.009)	(0.008)	(0.012)	(0.010)
	Receive diploma within 2 years of LCT (N=51,638)				
Pass	0.440	0.376	0.383	0.373	0.415
	(0.013)	(0.009)	(0.008)	(0.013)	(0.011)
	Receive diploma within 3 years of LCT (N=51,638)				
Pass	0.437	0.373	0.381	0.370	0.413
	(0.013)	(0.010)	(0.008)	(0.013)	(0.011)

Notes: see notes to Table 2-F

Table 3-F: Estimated earnings impact of passing last chance test, earing high school diploma

	Reduced-Form Estimates						IV Estimates	
	Poly-1	Poly-2	Poly-3	Poly-4	Poly-1 Lee-Card	Poly-1 bw=25	Poly-1 bw=25	Poly-1 bw=25, X
	Year 1 (Mean at score=0: 5,196.64; N=51,638)							
Pass/HSD	-66.3	-47.5	-146.6	-218.7	-120.1	-98.6	-168.8	-222.3
	(99.2)	(136.1)	(172.9)	(210.1)	(46.8)	(120.91)	(207.0)	(223.9)
Upper bound on CI							4.7	4.3
	Year 2 (Mean at score=0: 6,661.51; N=51,638)							
Pass/HSD	22.9	-20.9	-254.5	-296.3	-189.7	-110.5	-189.3	-243.7
	(124.3)	(170.1)	(214.3)	(256.6)	(71.2)	(185.62)	(317.8)	(335.7)
Upper bound on CI							6.7	6.4
	Year 3 (Mean at score=0: 7,573.54; N=51,638)							
Pass/HSD	108.6	-48.7	-202.6	-72.1	18.3	63.5	108.9	54.6
	(143.5)	(193.7)	(241.3)	(291.9)	(192.9)	(227.52)	(389.8)	(402.7)
Upper bound on CI							11.7	11.4
	Year 4 (Mean at score=0: 8,428.39; N=36,611)							
Pass/HSD	106.9	-25.0	-155.3	16.7	241.6	76.5	128.7	111.3
	(172.7)	(236.9)	(297.9)	(360.1)	(238.2)	(273.96)	(460.2)	(480.1)
Upper bound on CI							12.4	12.7
	Year 5 (Mean at score=0: 8,897.68; N=23,274)							
Pass/HSD	-328.4	-49.3	552.3	315.4	-39.7	700.6	1208.9	1234.8
	(472.9)	(424.2)	(433.1)	(809.8)	(654.5)	(551.23)	(948.3)	(981.8)
Upper bound on CI							34.9	35.9
	Year 6 (Mean at score=0: 8,428.36; N=9,338)							
Pass/HSD	-860.3	-1,103.8	-95.4	-272.4	-915.2	-120.3	-201.1	-350.4
	(620.8)	(642.8)	(691.0)	(1086.1)	(954.5)	(736.60)	(1070.7)	(1085.5)
Upper bound on CI							23.0	21.6
	Years 1-6 (Mean at score=0: 32,253.60; N=51,638)							
Pass/HSD	-56.7	-523.7	-667.4	-554.8	-123.4	142.5	244.1	258.6
	(636.4)	(805.6)	(967.4)	(1231.9)	(1094.3)	(867.13)	(1485.4)	(1571.0)
Upper bound on CI							10.0	10.5

Notes: Column heading Poly-X implies that estimates in this column are based on a regression of the outcome on a X'th order polynomial (interacted) in the last-chance score. "Bw=30" implies that the estimates are based on a subset of those in the last-chance sample that scored within 30 points of zero. Clustered (on the last-chance score) standard errors in parentheses. "Lee-Card" refers to estimates based on the Lee and Card (2008) "efficient" estimator. IV estimates use passing the last-chance test as an instrument for whether a high school diploma was received within two years of taking the last-chance test. "Upper bound on CI" refers to the upper bound on the 95% confidence interval around the estimate, expressed as a fraction of the mean earnings observed among those who scored zero (i.e., barley passed) on the last-chance test. "X" (last column) indicates that estimates based on regression models that include as covariates the variables listed in Table 1.

Table 3-T: Estimated earnings impact of passing last chance test, earing high school diploma

	Reduced-Form Estimates						IV Estimates	
	Poly-1	Poly-2	Poly-3	Poly-4	Poly-1 Lee-Card	Poly-1 bw=15	Poly-2	Poly-2 X
Year 1 (Mean at score=0: 5,418.19; N=42,510)								
Pass/HSD	12.4	19.6	-117.1	-182.5	-65.7	49.3	52.1	42.2
	(67.4)	(83.5)	(99.0)	(130.4)	(64.19)	(71.7)	(222.6)	(190.4)
Upper bound on CI							9.2	7.8
Year 2 (Mean at score=0: 7,217.57; N=42,510)								
Pass/HSD	104.0	121.4	89.5	66.1	173.8	163.3	323.3	348.1
	(81.8)	(119.3)	(162.5)	(207.6)	(104.92)	(96.2)	(312.9)	(285.2)
Upper bound on CI							13.2	12.7
Year 3(Mean at score=0: 8,688.07; N=42,510)								
Pass/HSD	161.1	114.8	100.5	221.6	257.4	187.8	305.7	371.8
	(116.6)	(159.4)	(216.3)	(229.3)	(114.86)	(128.8)	(419.6)	(421.7)
Upper bound on CI							13.2	14.0
Year 4 (Mean at score=0: 10,049.82; N=42,510)								
Pass/HSD	94.5	133.0	195.6	263.4	283.9	188.1	354.2	481.2
	(127.2)	(148.1)	(157.5)	(176.0)	(86.55)	(129.1)	(389.8)	(400.7)
Upper bound on CI							11.3	12.8
Year 5 (Mean at score=0: 11,217.34; N=42,510)								
Pass/HSD	62.2	8.5	131.5	320.5	221.0	124.6	22.6	214.4
	(133.8)	(180.0)	(153.1)	(171.3)	(90.68)	(142.1)	(479.0)	(502.3)
Upper bound on CI							8.7	10.9
Year 6(Mean at score=0: 12,331.90; N=42,510)								
Pass/HSD	99.7	-108.6	-145.0	-37.1	59.4	58.3	-289.2	-57.8
	(124.8)	(144.0)	(128.6)	(160.8)	(71.19)	(120.3)	(386.4)	(412.7)
Upper bound on CI							3.9	6.2
Year 7(Mean at score=0: 12,932.83 N=42,510)								
Pass/HSD	69.7	-232.2	-143.0	-8.2	-132.5	3.3	-618.2	-339.9
	(147.7)	(172.0)	(171.8)	(169.0)	(102.40)	(145.0)	(460.3)	(470.5)
Upper bound on CI							2.3	4.6
Year 1-7 (Mean at score=0: 67,855.71; N=42,510)								
Pass/HSD	603.5	56.6	111.9	643.8	797.5	774.8	150.6	1059.8
	(622.2)	(800.0)	(956.4)	(1091.0)	(545.65)	(666.5)	(2128.3)	(2183.0)
Upper bound on CI							6.5	8.0

Notes: see notes to Table 3-F

Table 4-F: Estimated impacts of passing last chance test, earning HSD on other earnings outcomes

	P(positive earnings)			Conditional-on-positive earnings		
	Reduced-form		IV	Reduced-form		IV
	Poly-4	Poly-1 bw=25	Poly-1 bw=25	Poly-4	Poly-1 bw=25	Poly-1 bw=25
Year 1 (Mean %Any Earnings = 75.74; Mean COP Earnings = 6861)						
Pass/HSD	-0.195	-0.350	-0.599	-272.14	-98.57	-167.39
	(1.44)	(1.08)	(1.85)	(192.69)	(144.65)	(246.24)
Upper bound on CI						4.74
Year 2 (Mean %Any Earnings = 75.74; Mean COP Earnings = 8795)						
Pass/HSD	1.413	1.597	2.735	-578.69	-353.20	-609.22
	(1.28)	(0.86)	(1.46)	(386.56)	(233.91)	(403.87)
Upper bound on CI						2.26
Year 3 (Mean %Any Earnings = 73.55; Mean COP Earnings = 10,297)						
Pass/HSD	0.505	0.120	0.206	-166.61	69.35	119.03
	(1.82)	(1.14)	(1.95)	(461.09)	(300.24)	(515.24)
Upper bound on CI						11.16
Year 4 (Mean %Any Earnings = 70.54; Mean COP Earnings = 11,948)						
Pass/HSD	1.136	1.120	1.883	-143.09	69.35	-144.11
	(2.02)	(1.43)	(2.40)	(411.16)	(300.24)	(493.07)
Upper bound on CI						7.05
Year 5 (Mean %Any Earnings = 68.0; Mean COP Earnings = 13,085)						
Pass/HSD	0.106	1.796	3.099	486.12	706.81	1216.66
	(2.72)	(1.69)	(2.91)	(1137.12)	(799.79)	(1375.05)
Upper bound on CI						30.32
Year 6 (Mean %Any Earnings = 72.57; Mean COP Earnings = 11,615)						
Pass/HSD	0.583	-1.020	-1.704	-474.52	37.43	62.43
	(4.28)	(2.76)	(4.61)	(1539.36)	(1045.48)	(1744.47)
Upper bound on CI						30.58

Notes: Column heading Poly-X implies that estimates in this column are based on a regression of the outcome on a X'th order polynomial (interacted) in the last-chance score. "Bw=30" implies that the estimates are based on a subset of those in the last-chance sample that scored within 30 points of zero. Clustered (on the last-chance score) standard errors in parentheses. IV estimates use passing the last-chance test as an instrument for whether a high school diploma was received within two years of taking the last-chance test. "Upper bound on CI" refers to the upper bound on the 95% confidence interval around the estimate, expressed as a fraction of the mean earnings observed among those who scored zero (i.e., barley passed) on the last-chance test.

Table 4-T: Estimated impacts of passing last chance test, earning HSD on other earnings outcomes

	P(positive earnings)			Conditional-on-positive earnings		
	Reduced-form		IV	Reduced-form		IV
	Poly-2	Poly-1 bw=15	Poly-2	Poly-2	Poly-1 bw=15	Poly-2
Year 1 (Mean %Any Earnings = 79.17; Mean COP Earnings = 6834)						
Pass/HSD	-0.125	-0.055	-0.146	30.88	71.11	80.70
	(0.64)	(0.59)	(1.57)	(81.74)	(73.69)	(213.28)
Upper bound on CI						7.42
Year 2 (Mean %Any Earnings = 79.31; Mean COP Earnings = 8918)						
Pass/HSD	0.707	0.479	1.275	100.76	125.75	260.86
	(0.79)	(0.95)	(2.50)	(102.78)	(88.32)	(265.63)
Upper bound on CI						8.88
Year 3 (Mean %Any Earnings = 78.9; Mean COP Earnings = 10,958)						
Pass/HSD	-0.194	-0.765	-2.036	250.01	263.09	642.35
	(0.57)	(0.60)	(1.62)	(190.49)	(149.77)	(482.28)
Upper bound on CI						14.66
Year 4 (Mean %Any Earnings = 78.82; Mean COP Earnings = 12,644)						
Pass/HSD	-0.309	-0.389	-1.036	233.94	284.27	601.34
	(0.65)	(0.72)	(1.93)	(111.56)	(101.13)	(281.75)
Upper bound on CI						9.21
Year 5 (Mean %Any Earnings = 78.22; Mean COP Earnings = 14,295)						
Pass/HSD	-0.658	-1.369	-3.646	256.04	269.95	673.89
	(0.59)	(0.56)	(1.53)	(176.13)	(139.11)	(456.18)
Upper bound on CI						11.10
Year 6 (Mean %Any Earnings = 77.37; Mean COP Earnings = 15,960)						
Pass/HSD	-0.972	-1.737	-4.626	207.76	253.54	542.89
	(0.40)	(0.38)	(1.05)	(164.70)	(157.84)	(428.96)
Upper bound on CI						8.78
Year 7 (Mean %Any Earnings = 76.22; Mean COP Earnings = 16,728)						
Pass/HSD	-0.876	-1.164	-3.100	-51.07	180.99	-134.62
	(0.55)	(0.63)	(1.72)	(190.53)	(174.54)	(500.91)
Upper bound on CI						5.18

Notes: see notes to Table 4-F

Table 5-T: 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	-5.380			8.118		
	(0.785)			(0.860)		
Year 2	-0.271			0.371		
	(0.205)			(0.760)		
Year 3	-0.203			0.092		
	(0.140)			(0.634)		
Year 4				-1.214		
				(0.503)		
Year 5		-7.200	-5.908	0.167	0.277	-0.378
		(0.512)	(0.274)	(0.492)	(0.370)	(0.206)
	42510	42510	42510	42510	42510	42510

Notes: estimates based on second-order polynomial. All estimates except "college credits" expressed in terms of percentage points.

Table 6-F: "Sheepskin estimates" of the High School Diploma premium

	Full Sample			Last-chance sample		
	No controls	Demographic controls	Full controls	No controls	Demographic controls	Full controls
	Year 1 (Mean=4,904; N=257,091)			Year 1 (Mean=5,054; N=51,638)		
Diploma	345.94	336.75	710.51	258.74	239.59	65.46
	(28.26)	(29.53)	(32.10)	(56.87)	(56.98)	(58.23)
% Mean earnings	7.05	6.87	14.49	5.12	4.74	1.30
	Year 2 (Mean=6,030; N=257,091)			Year 2 (Mean=6,425; N=51,638)		
Diploma	462.12	450.77	951.66	482.49	345.08	164.42
	(33.60)	(34.06)	(37.43)	(71.61)	(71.75)	(73.19)
% Mean earnings	7.66	7.47	15.78	7.51	5.37	2.56
	Year 3 (Mean=6,372; N=257,091)			Year 3 (Mean=6,935; N=51,638)		
Diploma	699.65	699.60	1,240.50	829.21	560.88	358.93
	(31.41)	(66.35)	(68.43)	(81.15)	(83.51)	(85.39)
% Mean earnings	10.98	10.98	19.47	11.96	8.09	5.18
	Year 4 (Mean=6,852; N=164,649)			Year 4 (Mean=7,518; N=36,611)		
Diploma	910.95	836.83	1,390.00	915.36	736.74	500.42
	(54.56)	(55.98)	(65.96)	(101.12)	(101.65)	(103.38)
% Mean earnings	13.29	12.21	20.29	12.17	9.80	6.66
	Year 5 (Mean=7,190; N=100,777)			Year 5 (Mean=8,164.2; N=23,274)		
Diploma	1,352.24	1,210.47	1,539.72	687.11	686.16	429.09
	(74.77)	(75.18)	(79.54)	(310.32)	(238.00)	(223.51)
% Mean earnings	18.81	16.84	21.42	8.42	8.40	5.26
	Year 6 (Mean=7,054; N=47,255)			Year 6 (Mean=8,217.5; N=9,338)		
Diploma	1,979.97	1,799.71	1,877.14	840.71	786.16	553.99
	(117.98)	(114.92)	(121.97)	(373.08)	(319.75)	(294.36)
% Mean earnings	28.07	25.51	26.61	10.23	9.57	6.74
	Years 1-6 (Mean=25,867; N=257,091)			Years 1-6 (Mean=28,274; N=51,638)		
Diploma	2,908.33	2,827.92	4,745.19	4,992.30	2,233.19	1,364.26
	(142.25)	(139.45)	(154.83)	(359.58)	(341.64)	(343.60)
% Mean earnings	11.24	10.93	18.34	17.66	7.90	4.83

Notes: estimates from regressions of outcome on high school diploma and three types of controls: no controls, demographic controls (sex and race) and full controls (the variables listed in Table 1). Mean earnings refer to those without a high school diploma in the relevant sample.