Returns to Education: The Causal Effects of Education on Earnings, Health, and Smoking

James J. Heckman,
University of Chicago and the American Bar Foundation

John Eric Humphries, and
University of Chicago

Gregory Veramendi
Arizona State University

Abstract

This paper estimates returns to education using a dynamic model of educational choice that synthesizes approaches in the structural dynamic discrete choice literature with approaches used in the reduced form treatment effect literature. It is an empirically robust middle ground between the two approaches which estimates economically interpretable and policy-relevant dynamic treatment effects that account for heterogeneity in cognitive and non-cognitive skills and the continuation values of educational choices. Graduating college is not a wise choice for all. Ability bias is a major component of observed educational differentials. For some, there are substantial causal effects of education at all stages of schooling.

James J. Heckman: Department of Economics, University of Chicago, 1126 East 59th Street, Chicago, IL 60637; phone: 773-702-0634; fax: 773-702-8490; jjh@uchicago.edu. John Eric Humphries: Department of Economics, University of Chicago, 1126 East 59th Street, Chicago, IL 60637; phone: 773-980-6575; johneric@uchicago.edu. Gregory Veramendi: Arizona State University, 501 East Orange Street, CPCOM 412A, Tempe, AZ 85287-9801; phone: 480-965-0894; gregory.veramendi@asu.edu.

James J. Heckman, Department of Economics, University of Chicago, 1126 East 59th Street, Chicago, IL 60637, Phone: 773-702-0634, jjh@uchicago.edu

John Eric Humphries, Department of Economics, University of Chicago, 1126 East 59th Street, Chicago, IL 60637, Phone: 773-980-6575, johneric@uchicago.edu

Gregory Veramendi, Department of Economics, Arizona State University, 501 East Orange Street, CPCOM, 412A, Tempe, AZ 85287-9801, Phone: 480-965-0894, gregory.veramendi@asu.edu

[84]Variation in the expected treatment effect comes from the variation in observed variables (X) and the unobserved endowments ($\boldsymbol{\theta}$).

1 Introduction

In his pioneering analysis of human capital, Gary Becker (1962; 1964) emphasized the importance of the rate of return for evaluating the effectiveness of human capital investments. He launched an active industry estimating returns to schooling.[1]

At the time Becker crafted his analysis, modern economic dynamics was in its infancy, as was research on the economics of uncertainty in dynamic sequential models. In an early contribution, Burton Weisbrod (1962) noted that each year of schooling attained opened up options for additional schooling and training and provided opportunities for learning about personal abilities and life opportunities.[2]

A parallel development in empirical economics was the growing awareness of heterogeneity and diversity among individual cognitive and non-cognitive abilities.[3] Agents differ in their returns to schooling. Failure to account for this heterogeneity leads to confusion in interpreting estimated effects of schooling.

Becker's early work focused on internal rates of return that equated *ex post* discounted values of earnings streams net of monetary and psychic costs at different levels of education. He noted that the full return to schooling includes non-market benefits and non-pecuniary costs. In modern parlance, individuals should continue their schooling as long as their *ex ante* marginal return exceeds their *ex ante* marginal opportunity cost of funds.

Formidable empirical challenges arise in estimating *ex post* internal rates of return: lifetime earnings profiles are required; observed earnings profiles are subject to the selection bias that arises from the fact that earnings are observed only at schooling levels selected by agents; and quantifying non-market benefits and non-pecuniary costs is a difficult task. For estimating *ex ante* returns, information on how agents forecast future events is also required.

In a neglected paper, Becker and Chiswick (1966) developed a tractable framework for measuring *ex post* returns to schooling that utilizes cross-section synthetic cohort data on earnings to approximate life cycle earnings data.[4] Mincer (1974) improved on this model by adding work experience. The "Mincer Equation" has become the workhorse of the empirical literature on estimating *ex post* rates of return:

---

[1]Becker (1964) also estimated rates of return. For surveys of this literature, see, e.g., Card (1999, 2001); Heckman et al. (2006a); Oreopoulos and Salvanes (2011); McMahon (2009); Oreopoulos and Petronijevic (2013).

[2]Weisbrod's paper stimulated research on the option value of schooling. See, e.g., Comay et al. (1973); Dothan and Williams (1981); Bamberger (1987); Altonji (1993); Cameron and Heckman (1993); Keane and Wolpin (1997); Arcidiacono and Miller (2011); Heckman et al. (2008); Stange (2012); Eisenhauer et al. (2015a).

[3]See Heckman (2001).

[4]It is based on the assumption that the earnings of a person age a in a given cross section when that person turns $a'(> a)$ is well-approximated by the earnings of agents a' in that same cross section. This synthetic cohort assumption is standard in the literature.

$$\ln Y(S_i, X_i) = \gamma_i + \rho_i \underbrace{S_i}_{\substack{\text{years of} \\ \text{schooling}}} + \phi(\underbrace{X_i}_{\substack{\text{other} \\ \text{determinants}}}) \quad (1)$$

where $Y(S_i, X_i)$ is the earnings of individual i with S_i years of schooling and a vector of other determinants X_i.

This equation is interpreted as a causal relationship generated by hypothetical variations of each of γ_i, ρ_i, and $\phi(X_i)$, holding other components on the right-hand side of (1) fixed.[5] γ_i is what person i would earn independent of any influence of schooling X_i. Correlation between γ_i and S_i is the source of "ability bias" (Griliches, 1977). Strictly speaking, γ_i may or may not be related to ability. It is a determinant of earnings that may also be correlated with S_i. ρ_i is the "return to a unit of schooling" for person i and is allowed to vary among individuals. It is a causal parameter realized by acquiring one more unit of schooling. There are both *ex ante* and *ex post* definitions of γ_i and ρ_i. The early literature and most of the empirical literature today focuses on estimating *ex post* returns.

This paper examines the economic foundations of Equation (1) and its generalizations accounting for the dynamics of educational decision-making and multidimensional heterogeneity in abilities among agents. We develop and estimate an empirically robust dynamic discrete choice model that allows for agent fallibility arising from imperfect information and learning, as well as time inconsistency. We allow agents to make schooling decisions based on expected future values. We test and reject strong forms of forward-looking behavior, but nonetheless find that agents sort on *ex post* gains.

We develop and estimate a variety of economically motivated and policy-relevant treatment effects. For most of the outcomes studied in this paper, we find strong evidence of ability bias at all levels of education, where ability includes both cognitive and non-cognitive skills, but only find sorting on gains (a relationship between ρ_i and S_i) at higher levels of schooling.

1.1 Interpreting Returns to Education

The Becker-Chiswick-Mincer Equation (1), and variants of it, have become the standard framework for estimating *ex post* returns to schooling for a variety of outcomes.[6] While ρ_i is not, in general, an internal rate of return for individual i, it is the *ex post* causal effect of increasing final schooling by exactly one year from any base state of schooling, holding γ_i and X_i fixed.[7] It is the slope of an hedonic wage function—the derivative of the aggregate production function evaluated at $S = s$ for a fixed γ_i and X_i.

ρ_i ignores the continuation values arising from the dynamic sequential nature of the schooling decision where information is updated and schooling at one stage opens up

---

[5]See Heckman (2008) and Heckman and Pinto (2015) for a discussion of causality and the role of fixing.

[6]See, e.g., Cutler and Lleras-Muney (2010) who apply model (1) to estimate the causal effect of education on health.

[7]The stringent conditions under which ρ_i is an internal rate of return, and evidence that they are not satisfied in many commonly used samples, are presented in Heckman et al. (2006a).

options for schooling at later stages. More generally, for a person at $s-1$, the perceived *ex ante* gain in log earnings of moving to schooling level s is the anticipated direct effect ρ_i and the (undiscounted) perceived continuation value of schooling for person i:

$$R_{s,i} = \rho_i + \underbrace{\rho_i \sum_{l=s+1}^{\bar{s}} P_{s,l,i}}_{\text{Continuation Value}} \cdot \quad (2)$$

Under an *ex ante* interpretation, $P_{s,l,i}$ is the agent's perceived probability of attaining (at least) schooling level $S = l$ for a person starting at schooling level s, including any relevant discounting of future benefits; $\bar{s}$ is the highest attainable value of S. $R_{s,i}$ captures Weisbrod's notion of valuing the future options that attaining schooling level s opens up.[8] It is the individual causal effect of an extra year of schooling inclusive of continuation values. As long as $\rho_i \neq 0$, it is distinct from $R_{s,i}$. One can define different versions of $R_{s,i}$ depending on how $P_{s,l,i}$ and ρ_i are specified.

Determining (2) poses major empirical challenges. There are multiple sources of heterogeneity in $R_{s,i}$. Individuals may differ in their values of ρ_i. Even if all people have the same ρ_i, they may differ in their expected anticipated probabilities of attaining schooling level s' ($P_{s,s',i}, s' > s$).[9]

The causal effects ρ_i and $R_{s,i}$ are formulated at the individual level. The modern treatment effect literature defines versions of these parameters for different groups and typically estimates *ex post* effects.[10] Thus, one can define the mean causal effect for the whole population $E(\rho)$.[11] Another possible causal effect is $E(R_s)$ defined for schooling level s for the entire population. One could also define the direct return to schooling for those who *choose* to be at a given level of schooling $E(\rho|S = s)$. This is the causal effect of one more unit of schooling for those who *stop* at $S = s$. One can define causal parameters for samples defined by other choices (e.g., for those indifferent between s and s'; for those who would stop at $s-1$; etc.), and for different notions of returns, e.g., $E(R|S = s)$.

$E(\gamma|S = s)$ is the population mean γ arising solely from statistical dependence between γ and S. It has no causal basis and is the source of ability bias. Since dependence between γ and S may arise from multiple sources, we refer to "ability bias" as selection bias throughout much of this paper.

The early literature adopted a simple approach to identifying returns. It assumed that ρ_i is identical for persons with the same observed characteristics. In this case, the only source of bias in estimating (1) is the statistical dependence between γ_i and S_i (selection bias). The recent literature recognizes heterogeneity in both γ_i and ρ_i. Both may be statistically dependent on S_i, giving rise to both selection bias and sorting on gains. The latter arises

---

[8]Note, however, that the continuation value is different from the option value. See, e.g., Stange (2012) and Eisenhauer et al. (2015a).
[9]Rational expectations models assume that objectively measured probabilities are subjective probabilities. We do not impose this assumption in our analysis. For a survey of the expectation elicitation literature, see, e.g., Manski (2004).
[10]See Heckman (2008).
[11]See, e.g., Card (1999; 2001).

because the causal effect of S may be moderated by other variables. Whether or not sorting gains are a source of bias depends on the question being addressed.

To illustrate the importance of accounting for continuation values, consider a compulsory schooling policy that forces all persons to take a minimum level of schooling ($S \geq \underline{s}$). What causal effect is identified by this "natural experiment?" Abstracting from general equilibrium effects, any estimated treatment effect is defined conditional on the set of people who change their schooling from below $\underline{s}$ to at or above $\underline{s}$. However, there is no presumption that such agents will stop at $\underline{s}$ if they are forced to attain it. They may learn things about themselves and their possibilities, so they continue beyond $\underline{s}$ and thereby generate continuation values.[12] Thus, an experiment that evaluates the effects of this policy does not, in general, estimate $E(\rho)$ or even $E(\rho | S = \underline{s})$. It does not, in general, estimate the marginal effect of a change in S on the log marginal price of schooling.

The analysis just presented can be generalized to incorporate non-linear structural (causal) returns to schooling by allowing the ρ_i to depend on the origin and destination schooling states ($\rho_{s,s',i}$) for $s' > s$. Non-linearities associated with sheepskin effects associated with graduation are a potentially important source of continuation values.

1.2 Approaches to Identifying Causal Effects and Causal Rates of Return

Two general approaches have been developed to estimate returns to schooling in the general case. They are: (i) structural models that jointly analyze outcomes and schooling choices; and (ii) treatment effect models that use instrumental variables methods (including randomization and regression discontinuities as instruments) as well as matching on observed variables to identify causal parameters.[13]

The structural approach explicitly models agent decision rules that generate $P_{s,1,i}$ and the dependence between ρ_i and S_i. The modern version explicitly models agent expectations and distinguishes *ex ante* from *ex post* returns.[14] It uses a variety of sources of identification, including exclusion restrictions (instrumental variables), conditional independence assumptions about unobservables, and functional form assumptions (see, e.g., Blevins, 2014). Among other features, the structural approach identifies causal effects at well-defined margins of choice and can evaluate the impacts of different policies never previously implemented.[15]

The treatment effect approach is typically agnostic about agent decision rules and relies on exclusion restrictions to identify its estimands. It rarely distinguishes *ex ante* from *ex post* returns.[16] This approach is more transparent in securing identification than the structural approach.[17] However, the economic interpretation of its estimated parameters is often quite

---

[12]This is recognized in the LATE literature. See Angrist and Imbens (1995). What is not recognized in that literature is that LATE estimates the returns expected by agents only under a rational expectations assumption.

[13]See, e.g., Angrist and Imbens (1995) and Angrist and Pischke (2009) for IV, and Heckman et al. (1998) for matching.

[14]See, e.g., Keane and Wolpin (1997); Eisenhauer et al. (2015a).

[15]See Heckman (2010) and Heckman and Urzúa (2010).

[16]Eisenhauer et al. (2015b) distinguish and estimate *ex ante* and *ex post* returns in an instrumental variable model.

[17]The modern instrumental variables case requires assumptions about the validity of the instruments. If there are heterogeneous treatment effects, additional assumptions such as "monotonicity" (better termed uniformity) are required to interpret IV estimates. See Imbens and Angrist (1994); Heckman and Vytlacil (2005); Angrist and Pischke (2009) for details.

obscure. In a model with multiple levels of schooling, LATE typically does not identify returns at the various margins of choice that generate outcomes or the sub-populations (defined in terms of observables and unobservables) affected by the instruments used.[18] Its estimands do not identify a variety of well-posed policy questions except when the variation induced by the instruments corresponds closely to the variations induced by the policies of interest.[19]

We build on the analyses of Heckman and Vytlacil (1999, 2005, 2007a,b), Carneiro et al. (2010, 2011), and Eisenhauer et al. (2015b), who introduce choice theory into the modern analysis of instrumental variables. They focus on binary choice models but also analyze ordered and unordered choice models with multiple outcomes to estimate economically interpretable treatment effects. Expanding on that body of research, we consider multiple sources of identification besides instrumental variables. We do not rely on continuous instruments. In addition, we link our analysis to the dynamic discrete choice literature.

1.3 Our Approach

This paper develops a methodological middle ground between the reduced form treatment approach and the fully structural dynamic discrete choice approach. As in the structural literature, we estimate causal effects at clearly identified margins of choice. Our methodology identifies which agents are affected by instruments as well as which persons would be affected by alternative policies not previously implemented. As in the treatment effect literature, we are agnostic about the precise rules used by agents to make decisions. Unlike that literature, we recognize the possibility that people make decisions and account for the consequences of their choices. We approximate agent decision rules and do not impose the cross-equation restrictions that are the hallmark of the structural approach, nor do we explicitly model agent expectations about costs and returns.[20]

Using a generalized Roy framework, we estimate a multistage sequential model of educational choices and their consequences. An important feature of our model is that educational choices at one stage open up educational options at later stages. Each educational decision is characterized using a flexible discrete choice model. The anticipated consequences of future choices and their costs can be assessed in a variety of ways by individuals in deciding whether or not to continue their schooling. Our model approximates a dynamic discrete choice model without taking a stance on exactly what agents are maximizing or their information sets.

Like structural models, our model is identified though multiple sources of variation. Drawing from the matching literature, we identify the causal effects of schooling at different stages of the life cycle by using a rich set of observed variables and by proxying unobserved endowments. Unlike previous work on matching, we correct the match variables for measurement error and the bias introduced into the measurements by family background. We also use exclusion restrictions to identify our model as in the IV and control function

---

[18]See Heckman et al. (2006c) and Heckman et al. (2016) for a discussion.
[19]See Heckman (2010).
[20]Such approximations are discussed in Heckman (1981), Eckstein and Wolpin (1989), Cameron and Heckman (2001), and Geweke and Keane (2001).

literatures. Unlike many structural papers, we provide explicit proofs of model identification. [21]

Our framework allows agents to make *ex ante* valuations as in dynamic discrete choice models but does not explicitly identify them.[22] However, we estimate a variety of *ex post* returns to schooling, and model how they depend on both observed and unobserved variables. We decompose *ex post* treatment effects into (i) the direct benefits of going from one level of schooling to the next;[23] and (ii) continuation values arising from access to additional education beyond the next step.

Estimating our model on NLSY79 data, we investigate foundational issues in human capital theory. We report the following findings.

1. There are substantial returns/causal effects of education on wages, the present value of wages, health, and smoking.[24]

2. The continuation values arising from sequential choices are empirically important components of returns to education. Low-ability individuals gain mostly from graduating high school and stopping there. High-ability individuals have substantial post-high school continuation values.

3. Estimated returns (causal effects) differ by schooling level and depend on observed and unobserved characteristics of individuals. Graduating high school benefits all—and especially low-ability persons. Only high-ability individuals receive substantial benefits from college graduation. There is positive sorting on gains only at higher educational levels.

4. People sort on *ex post* gains, especially more able people at higher schooling levels, confirming a core tenet of human capital theory. Yet, at the same time, people do not know or act on publicly available information when making decisions about high school graduation.

5. This paper contributes to an emerging literature on the importance of both cognitive and non-cognitive abilities in shaping life outcomes.[25] Consistent with the recent literature, we find that both types of abilities are important predictors of educational attainment. Within schooling levels, cognitive and non-cognitive abilities have impacts on most outcomes.[26]

6. Selection bias arising from both observed and unobserved variables accounts for a substantial portion (typically over one half) of the observed differences in wage outcomes classified by education. This finding runs counter to a common

---

[21]Heckman and Navarro (2007) and Blevins (2014) also proof identifiability of structural models.
[22]See, e.g., Eisenhauer et al. (2015a).
[23]The human capital literature traditionally focused on the direct causal benefits of one final schooling level compared to another, but makes sequential comparisons from the lowest levels of schooling to the highest (Becker, 1964).
[24]There is a small, but growing literature on the effects of education on health and healthy behaviors. See Grossman (2000); McMahon (2000); Lochner (2011); Oreopoulos and Salvanes (2011); Cutler and Lleras-Muney (2010). For a review of this literature see Web Appendix A.1.
[25]See, e.g., Borghans et al. (2008); Heckman et al. (2006b); Almlund et al. (2011).
[26]Our estimates of the causal effects of education do *not* require that we separately isolate the effects of individual cognitive and non-cognitive endowments on outcomes, just that we control for them as a set.

interpretation in the literature based on comparing IV and OLS estimates of Equation (1).[27]

Using our estimated model, we conduct two policy experiments. In the first, we examine the impact of a tuition subsidy on college enrollment. We identify who is affected by the policy, how their decisions change, and how much they benefit. Those induced to enroll benefit from the policy, and many go on to graduate from college. In a second experiment, we analyze a policy that improves the ability endowments of those at the bottom of the distribution to see how this impacts educational choices and outcomes. Such improvements are produced by early intervention programs.[28] Increasing cognitive endowments positively impacts all outcomes, while increasing non-cognitive endowments mostly impacts smoking and health outcomes.

Our paper proceeds in the following way. Section 2 presents our model. Section 3 presents economically interpretable treatment effects (rates of return) that can be derived from it. Section 4 discusses identification. Section 5 discusses the data analyzed and presents unadjusted associations and regression-adjusted associations between different levels of education and the outcomes analyzed in this paper. Section 6 reports our estimated treatment effects and interprets them. Section 7 uses the estimated model to address two policy-relevant questions. Section 8 tests a key identifying assumption. Section 9 compares our estimates to those derived from alternative methodological approaches such as OLS and matching. Section 10 concludes.

2 Model

This paper estimates a multistage sequential model of educational choices with transitions and decision nodes shown in Figure 1. Let $\mathscr{J}$ denote a set of possible terminal states. At each node there are only two possible choices: remain at j or transit to the next node ($j + 1$ if $j \in \{1, \ldots, \bar{s} - 1\}$). $D_j = 0$ if a person at j does not stop there and goes on to the next node. $D_j = 1$ if the person stops at j for $j \quad 0$. $D_0 = 1$ opens an additional branch of the decision tree. A person may remain a dropout or get the GED.[29] For $D_0 = 1$, we define the attainable set as $\{0, G\}$. Thus, in the lower branch ($D_0 = 1$), agents can terminate as a dropout ($D_0 = 1$, $D_G = 1$) or as a dropout who gets a GED certificate ($D_0 = 1$, $D_G = 0$). $D_j \in \mathscr{D}$ is the set of possible transition decisions that can be taken by the individual over the decision horizon. Let $\mathscr{S} = \{G, 0, \ldots, \bar{s}\}$ denote the set of stopping states with $S = s$ if the agent stops at $s \in \mathscr{S}$ ($D_s = 1$ for $s \in \mathscr{S} \backslash \{0, G\}$). Define $\bar{s}$ as the highest attainable element in $\mathscr{S}$ in the ordered subset $\{0, \ldots, \bar{s}\}$. We assume that the environment is time-stationary and decisions are irreversible.[30]

$Q_j = 1$ indicates that an agent *gets to* decision node j and acquires at least the education associated with j. $Q_j = 0$ if the person never gets there. $Q_G = 1$ if the agent drops out of high school and faces the GED option. The history of nodes visited by an agent can be described by the collection of the Q_j such that $Q_j = 1$. Observe that $D_s = 1$ is equivalent to $S = s$ for $s \in$

---

[27]See, e.g., Griliches (1977) and Card (1999, 2001).
[28]Heckman et al. (2013a).
[29]The GED is a test high school dropouts can take to earn state-issued high school equivalency credentials. For strong evidence on the nonequivalence of GEDs to high school dropouts, see Heckman et al. (2014).
[30]Versions of this model are also analyzed Cunha et al. (2007), Heckman and Navarro (2007), and Heckman et al. (2016).

$\{1, \ldots, \bar{s}\}$ and $D_{\bar{s}} = 1$ if $D_j = 0$, $\forall j \in \mathscr{S} \setminus \{\bar{s}\}$.[31] Finally, $D_0 = 1$ and $D_G = 0$ is equivalent to $S = G$.

2.1 A Sequential Decision Model

The decision process at each node is assumed to be characterized by an index threshold-crossing property:

$$D_j = \begin{cases} 0 & \text{if } I_j \quad 0 \\ 1 & \text{otherwise} \end{cases} \text{ for } Q_j = 1, \quad j \quad = \{G, 0, \ldots, \bar{s} - 1\} \quad (3)$$

where I_j is the agent's *perceived* value at node j of going on to the next node. The requirement $Q_j = 1$ ensures that agents are able to make the transition at j by conditioning on the population eligible to make the transition.

Associated with each final state $s \in \mathscr{S}$ is a set of K_s potential outcomes for each agent with indices $k \in \mathscr{K}_s$. We define the Y_s^k as latent variables that map into potential outcomes Y_s^k:

$$Y_s^k = \begin{cases} Y_s^k & \text{if } Y_s^k \text{ is continuous} \\ \mathbf{1}(Y_s^k \quad 0) & \text{if } Y_s^k \text{ is a binary outcome} \end{cases} \text{ for } k \qquad s' \quad s \qquad . \quad (4)$$

The outcome variables may be in levels, logs, or other transformations. Using the switching regression framework of Quandt (1958, 1972), the observed outcome Y^k for a k common across all decision nodes is

$$Y^k = \left(\sum_{\setminus \{0, G\}} D_s Y_s^k \right)(1 - D_0) + (Y_0^k D_G + Y_G^k (1 - D_G)) D_0. \quad (5)$$

2.2 Parameterizations of the Decision Rules and Potential Outcomes for Final States

Following a well-established tradition in the treatment effect and structural literatures, we approximate I_j using a separable model:

$$I_j = \phi_j \underbrace{(\mathbf{Z})}_{\substack{\text{Observed} \\ \text{by analyst}}} - \underbrace{\eta_j}_{\substack{\text{Unobserved} \\ \text{by analyst}}}, j \qquad , \quad (6)$$

where $\mathbf{Z}$ is a vector of variables observed by the analyst, components of which determine the transition decisions of the agent at different stages, and η_j is unobserved by the analyst. A

---

[31] For notational convenience, we assign $D_j = 0$ for all $j > s$.

separable representation of the choice rule is an essential feature of LATE (Vytlacil, 2002) and is often invoked in dynamic discrete choice models (Blevins, 2014).

This specification of agent decision-making is quite agnostic. It does not impose forward-looking behavior. Agents may be myopic or time-inconsistent and may be confronted by surprises. Because we do not impose particular expectation formation assumptions, we are not tied to a particular set of assumptions about agent rationality. A drawback of this approach is that we cannot identify *ex ante* versions of the economic parameters we estimate.

Outcomes are also assumed to be separable:

$$Y_s^k = \tau_s^k \underbrace{(X)}_{\substack{\text{Observed} \\ \text{by analyst}}} + \underbrace{U_s^k}_{\substack{\text{Unobserved} \\ \text{by analyst}}}, \quad k \qquad s' \quad s \qquad , \quad (7)$$

where X is a vector of observed determinants of outcomes and U_s^k is unobserved by the analyst.[32] Separability of the unobserved variables in the outcome equations is often invoked in the structural literature but is not strictly required in the structural or discrete choice literatures.[33]

2.3 Assumptions about the Unobservables

Central to our main empirical strategy is the existence of a finite dimensional vector $\boldsymbol{\theta}$ of unobserved (by the economist) endowments that generate all of the dependence across the η_j and the U_s^k. We assume that

$$\eta_j = -(\boldsymbol{\theta}\,\lambda_j - \nu_j), \quad j \qquad (8)$$

and

$$U_s^k = \boldsymbol{\theta}\,\boldsymbol{\alpha}_s^k + \omega_{s'}^k, \quad k \qquad s' \, s \qquad , \quad (9)$$

where ν_j is an idiosyncratic error term for transition j. ω_s^k represents an idiosyncratic error term for outcome k in state s.

---

[32]In our model, X and Z can vary by decision or outcome depending on the specification of functions $\tau_s^k(X)$ and $\phi_j(Z)$. See Table 1 for details.

[33]Moreover, we can condition on observable covariates X.

Conditional on $\boldsymbol{\theta}, \boldsymbol{X}, \boldsymbol{Z}$, choices and outcomes are statistically independent. Controlling for this set of variables eliminates selection effects. If the analyst knew $\boldsymbol{\theta}, \boldsymbol{X}, \boldsymbol{Z}$, he/she could use matching to identify the model.[34]

The standard "random effects" approach in the structural literature treats $\boldsymbol{\theta}$ as a nuisance variable and does not interpret it.[35] Our approach is to proxy $\boldsymbol{\theta}$ using multiple interpretable measurements of it. We correct for errors in the proxy variables. The measurements facilitate the interpretation of $\boldsymbol{\theta}$. We develop this intuition further in Section 4, after presenting the rest of our model.

We array the $\nu_j, j \in \mathcal{J}$ into a vector $\boldsymbol{\nu} = (\nu_G, \nu_0, \nu_1, \ldots, \nu_{\bar{s}-1})$, and the η_j into $\boldsymbol{\eta} = (\eta_G, \eta_0, \ldots, \eta_{\bar{s}-1})$. Array the ω_s^k into a vector $\boldsymbol{\omega}_s = (\omega_{s}^1, \ldots, \omega_s^{Ks})$. Array the U_s^k into vector $\boldsymbol{U}_s = (U_{s}^1, \ldots, U_s^{Ks})$, and array the $\boldsymbol{U}_s$ into $\boldsymbol{U} = (\boldsymbol{U}_G, \boldsymbol{U}_0, \ldots, \boldsymbol{U}_{\bar{s}})$.

Letting "⊥⊥" denote statistical independence, we assume that, conditional on $\boldsymbol{X}$

$$\nu_j \quad \nu_{l'} \qquad l \quad j \quad l, j \qquad \text{(A-1a)}$$

$$\omega_s^k \quad \omega_{s'}^k \qquad s \quad s \qquad k \quad \text{(A-1b)}$$

$$\boldsymbol{\omega}_s \quad \boldsymbol{\nu}_{,} \qquad s \qquad \text{(A-1c)}$$

$$\boldsymbol{\theta} \quad \boldsymbol{Z} \quad \text{(A-1d)}$$

$$(\boldsymbol{\omega}_{s'} \boldsymbol{\nu}) \quad (\boldsymbol{\theta}, \boldsymbol{Z}), \qquad s \qquad . \quad \text{(A-1e)}$$

Assumption (A-1a) maintains independence of the shocks affecting transitions; (A-1b) assumes independence of shocks across all states; (A-1c) assumes independence of the shocks to transitions and the outcomes; (A-1d) assumes independence of $\boldsymbol{\theta}$ with respect to the observables; and (A-1e) assumes independence of the shocks with the factors $\boldsymbol{\theta}$ and $\boldsymbol{Z}$. Versions of assumptions (A-1d) and (A-1e) play fundamental roles in the structural dynamic discrete choice literature.[36] Any dependence postulated across the $\boldsymbol{\omega}$ and $\boldsymbol{\nu}$ can be captured by introducing factors in $\boldsymbol{\theta}$.

---

[34] See Carneiro et al. (2003).
[35] See, e.g., Keane and Wolpin (1997); Rust (1994); Adda and Cooper (2003); Blevins (2014).

2.4 Measurement System for Unobserved Factors θ

We allow for the possibility that θ cannot be measured precisely, but that it can be proxied with multiple measurements. We correct for the effects of measurement error in the proxy. We link θ to measurements, and adjoin measurement equations to choice and outcome equations, making θ interpretable.

Let M be a vector of N_M measurements on θ. They may consist of lagged or future values of the outcome variables or additional measurements.[37] The system of equations determining M is

$$w M = (X, \theta, e), \quad (10)$$

where X are observed variables, θ are the factors, and

$$M = \begin{pmatrix} M_1 \\ M_{N_M} \end{pmatrix} = \begin{pmatrix} {}_1(X, \theta, e_1) \\ {}_{N_M}(X, \theta, e_{N_M}) \end{pmatrix},$$

where we array the e_j into $e = (e_1, \ldots, e_{N_M})$. We assume, in addition to the previous assumptions that, conditional on X,

$$e_j \quad e_{l'} \quad j \quad l, \quad j, l \quad \{1, \ldots, N_M\} \quad \text{(A-1f)}$$

$$\text{and} \quad e \quad (X, Z, \theta, \nu, \omega). \quad \text{(A-1g)}$$

For the purpose of identifying treatment effects, we do not need to identify each equation of system (10). We just need to identify the span of θ that preserves the information on θ in (10). That is sufficient to produce conditional independence between choices and outcomes.[38] However, in this paper we estimate equation system (10) to enhance interpretability.

3 Defining Returns/Causal Effects of Education

A variety of *ex post* counterfactual outcomes and associated treatment effects can be generated from our model. There is no single "causal effect" of education. The causal effects we analyze can be used to predict the effects of changing education levels through different

---

[36]For example, the widely-used "types" assumption of Keane and Wolpin (1997) postulates conditional independence between choices and outcomes conditional on types (θ) that operate through the initial conditions of their model.
[37]See, e.g., Abbring and Heckman (2007); Schennach et al. (2012).
[38]See, e.g., Heckman et al. (2013b).

policies for people of different backgrounds and abilities. They allow us to improve on the "effects" reported in the literature on instrumental variables to understand the effectiveness of policies for different identifiable segments of the population, and the benefits to people at different margins of choice. These effects are defined for different conditioning sets and thought experiments. Our dynamic model suggests a new range of treatment parameters that do not arise in models with binary treatments. This section makes precise the notions of returns to education discussed in Section 1.

In principle, we could define and estimate a variety of causal effects, many of which are not plausible. For example, many empirical economists would not find estimates of the effect of fixing (manipulating) $D_j = 0$ if $Q_j = 0$ to be credible (i.e., the person for whom we fix $D_j = 0$ is not at the decision node to take the transition).[39] In the spirit of credible econometrics, we define such treatment effects conditional on $Q_j = 1$. This approach blends structural and treatment effect approaches. Our causal parameters recognize agent heterogeneity and are allowed to differ across different subsets of the population.

The person-specific treatment effect T_j^k for outcome k for an individual selected from the population $Q_j = 1$ with characteristics $X = x, Z = z, \theta = \bar{\theta}$, making a decision at node j between going on to the next node or stopping at j, is the difference between the individual's outcomes under the two actions:

$$T_j^k[Y^k \quad X = x, Z = z, \theta = \bar{\theta}] := (Y^k \quad X = x, Z = z, \theta = \bar{\theta}, Q_j = 1, Fix\, D_j = 0) - (Y^k \quad X (11) = x, Z = z, \theta = \bar{\theta}, Q_j = 1, Fix\, D_j = 1).$$

The random variable ($Y^k | X = x, Z = z, \theta = \bar{\theta}, Q_j = 1, Fix\, D_j = 0$) is the outcome at node j for a person with characteristics $X = x, Z = z, \theta = \bar{\theta}$ from the population that attains node j (or higher), $Q_j = 1$, and for whom we fix $D_j = 0$ so they go on to the next node. They may choose to go even further. Random variable ($Y^k | X = x, Z = z, \theta = \bar{\theta}, Q_j = 1, Fix\, D_j = 1$) is defined for the same population but forces persons with those characteristics not to transit to the next node.

We present population-level treatment effects based on (11). We focus our discussion on means, but we also discuss distributional counterparts for all of the treatment effects considered in this paper.

3.1 Direct Effects and Continuation Values

A principal contribution of this paper is the definition and estimation of treatment effects that take into account the direct effect of moving to the next node of a decision tree, plus the benefits associated with the further schooling that such movement opens up. The associated mean treatment effect is the difference in expected outcomes arising from changing a single

---

[39]The distinction between *fixing* and *conditioning* traces back to Haavelmo (1943). White and Chalak (2009) use the terminology "setting" for the same notion. For a recent analysis of this crucial distinction, see Heckman and Pinto (2015).

educational decision in a sequential schooling model and tracing through its consequences, accounting for the dynamic sequential nature of schooling.

Person-specific treatment effects at node j can be decomposed into two components. The first component is the *direct effect* of going from j to $j + 1$: $DE_j^k = Y_{j+1}^k - Y_j^k$, the effect often featured in the literature on the returns to schooling when comparing schooling levels $j + 1$ and j (Becker, 1964). The second component is the *continuation value* of going beyond $j + 1$ for persons with $D_0 = 0$ (the upper branch of Figure 1), which is

$$C_{j+1}^k := \sum_{r=1}^{\bar{s}-(j+1)} \left[\prod_{l=1}^{r} (1 - D_{j+l}) \right] (Y_{j+r+1}^k - Y_{j+r}^k) . [40]$$

The continuation value for the lower branch of Figure 1 ($D_0 = 1$) is defined for the attainable set $\{0, G\}$. G is the only option available to a high school dropout in that branch. In the following, we analyze the upper branch of Figure 1. The analysis for the lower branch is similar.

At the individual level, the *total effect* of fixing $D_j = 0$ on Y^k is decomposed into

$$T_j^k = DE_j^k + C_{j+1}^k . \quad (12)$$

The associated population level average treatment effect at node j inclusive of continuation values, conditional on $Q_j = 1$, is

$$ATE_j^k := \int \ldots \int E(T_j^k[Y^k \mid X = x, Z = z, \theta = \bar{\theta}]) \, dF_{X,Z,\theta}(x, z, \bar{\theta} \mid Q_j = 1), \quad (13)$$

which can be decomposed into direct and continuation value components.

Integrating over the X, Z, θ, conditioning on $Q_j = 1$, the component of (13) due to the population continuation value at $j + 1$ is

---

[40]The relationship between this notion of continuation values and the definition used in the dynamic discrete choice literature is explored in Web Appendix A.3.

$$E_{X,Z,\theta}(C_{j+1}^k) = E_{X,Z,\theta}\left[\sum_{l=j+1}^{\bar{s}-1}\left\{E(Y_{l+1}^k - Y_l^k \mid X = x, Z = z, \theta = \bar{\theta}, Q_{l+1} = 1, Fix\, Q_{j+1} = 1)\right.\right.$$
$$\left.\left. \cdot Pr(Q_{l+1} = 1 \mid X = x, Z = z, \theta = \bar{\theta}, Q_j = 1, Fix\, Q_{j+1} = 1)\right\} \mid Q_j = 1\right],$$

(14)

where $Q_{\bar{s}} = 1$ if $S = \bar{s}$.

We can also define conditional (on X, Z, θ) population distributions of total effects as in Heckman et al. (1997):[41]

$$Pr(T_j^k < t_j^k \mid X = x, Z = z, \theta = \bar{\theta}, Q_j = 1) \quad (15)$$

and the population counterpart, integrating over X, Z, θ, which can be further decomposed into the distributions of direct effects and of continuation values.[42]

Because we do not specify or attempt to identify choice-node-specific agent information sets, we can only identify *ex post* treatment effects. Hence, we can identify continuation values associated with choices, but cannot identify option values. A benefit of this more agnostic approach is that it does not impose specific decision rules or assumptions about agent expectations. Our model allows for irrationality, regret, and mistakes in agent decision-making associated with maturation and information acquisition and allows us to test the validity of certain assumptions commonly made about agent expectations.

3.2 Average Marginal Treatment Effects

In order to understand the economic returns to an additional unit of schooling for persons at the margin of indifference at each node of the decision tree of Figure 1, we estimate the Average Marginal Treatment Effect (AMTE).[43] It is the average effect of transiting to the next node for individuals at or near the margin of indifference between the two nodes:

$$AMTE_j^k := \int E\left[T_j^k\left(Y^k \mid X = x, Z = z, \theta = \bar{\theta}\right)\right] dF_{X,Z,\theta}(x, z, \bar{\theta} \mid Q_j = 1, \quad I_j \quad \varepsilon),$$

(16)

---

[41] See Abbring and Heckman (2007) for a review of the literature.
[42] The modifications for the unordered case require that we define these terms over the admissible options available for $D_0 = 1$ or $D_0 = 0$.
[43] See Carneiro et al. (2010, 2011).

where e is an arbitrarily small neighborhood around the margin of indifference.[44] These effects are inclusive of all consequences of taking the transition at j, including the possibility of attaining final schooling levels well beyond j.[45] AMTE defines causal effects at well-defined and empirically identified margins of choice. It is the proper measure of the *ex post* marginal gross benefit for evaluating the gains from moving from one stage of the decision tree to the next for those at that margin of choice. In general, it is distinct from LATE, which is not defined for any specific margin of choice, and generally does not estimate $E(\rho)$ or $E(\rho \mid S = s)$, and includes the effects on outcomes for transitions induced by instruments beyond any schooling level at which the instrument operates.[46] Since we identify the distribution of I_j, we can identify the characteristics of agents in the indifference set, something not possible using LATE.[47]

The population distribution counterpart of AMTE is defined over the set of agents for whom $|I_j| \quad e$, which can be generated from our model: $Pr(T_j^k < t_j^k \quad Q_j = 1, \quad I_j \quad \varepsilon)$. Distributional versions can be defined for all of the treatment effects considered in this section.

3.3 Policy-Relevant Treatment Effects

The policy-relevant treatment effect (PRTE) is the average treatment effect for those induced to change their choices in response to a particular policy intervention. Let $Y^k(p)$ be the aggregate outcome under policy p for outcome k. Let $S(p)$ be the final state selected by an agent under policy p. The policy-relevant treatment effect from implementing policy p compared to policy p' for outcome k is:

$$PRTE_{p,p}^k := \quad E(Y^k(p) - Y^k(p) \quad X = x, Z = z, \theta = \bar{\theta})dF_{X,Z,\theta}(x,z,\bar{\theta} \quad S(p) \quad S(p))),$$

(17)

where $S(p) \quad S(p')$ denotes the set of the characteristics of people for whom attained states differ under the two policies. In general, it is different from AMTE because the agents affected by a policy can be at multiple margins of choice. PRTE is often confused with LATE. In general, they are different unless the proposed policy change coincides with the instrument used to define LATE.[48]

---

[44]Note that the limit of (16) as $e \to 0$ is not well-defined without further assumptions. This is the so-called "Borel paradox" discussed in this context in Carneiro et al. (2010). We avoid this problem by assuming a functional form for the distribution of e.
[45]One might also define a version of this treatment effect for two adjacent states ignoring continuation values.
[46]aSee Heckman and Vytlacil (2007a) and Carneiro et al. (2010). The LATE can correspond to people at multiple margins. See Angrist and Imbens (1995) and Heckman et al. (2016).
[47]Note that the indifference set may contain multiple margins, as in Heckman and Vytlacil (2007b) and Heckman and Urzúa (2010).
[48]See Carneiro et al. (2011) for an empirical example. The differences between the two parameters can be substantial as we show in Heckman et al. (2016).

3.4 Differences Across Final Schooling Levels

Becker's original approach to estimating returns to schooling (1964) focused on the upper branch of Figure 1 and reported estimates from pairwise comparisons of returns at final schooling levels. He defines returns to education as the gains from choosing between a terminal base state and a terminal final schooling level, implicitly assuming that the probabilities of all intervening transitions in Equation (2) are 1. Following Becker, but controlling for $\boldsymbol{\theta}$, $\boldsymbol{Z}$, and $\boldsymbol{X}$, the mean gain for the subset of the population that completes one of the two adjacent schooling level $S \in \{s, s'\}$ is:

$$ATE_{s,s}^k : = \qquad E(Y_s^k - Y_s^k \mid \boldsymbol{X} = \boldsymbol{x}, \boldsymbol{Z} = \boldsymbol{z}, \boldsymbol{\theta} = \bar{\boldsymbol{\theta}}) \, dF_{\boldsymbol{X},\boldsymbol{Z},\boldsymbol{\theta}}(\boldsymbol{x}, \boldsymbol{z}, \bar{\boldsymbol{\theta}} \mid S \quad \{s, s\}). \quad (18)$$

Unlike (13), this parameter ignores continuation values.

Conditioning in this fashion recognizes that the characteristics of people not making either final choice could be far away from the population making one of those two choices, and hence, might be far away from having any empirical or policy relevance.[49] One can also compute parameters of $ATE_{s,s'}$ for other conditioning sets, such as $S = s'$ (treatment on the treated). We report estimates of different versions of these treatment effects in the Web Appendix A.14.1.

3.5 Decomposing Observed Differences in Outcomes into Selection Bias, Sorting Gains, and Average Treatment Effects

Using our model, we interpret "ability bias" (really selection bias) and sorting on gains using the traditional Becker-Chiswick-Mincer model (1) and its extensions as a benchmark. To simplify the exposition, we focus on the upper branch of Figure 1 ($D_0 = 0$) and analyze continuous outcomes.[50]

There are two basic models used in the empirical literature estimating returns to schooling. One version studies outcomes and selection bias in terms of pairwise final schooling levels (s_0, s) attained by agents ($D_{s_0} + D_s = 1$), $s_0 \quad s$. It is defined for the population at one of these two terminal schooling states. It does not include terminal values beyond s. Another version studies gains and ability bias in terms of benefits associated with attaining (and possibly exceeding) given schooling levels ($Q_j = 1$). This includes continuation values. In the text, we develop both widely-used versions.

The effect of additional schooling starting at s_0 and stopping at s is captured by $Y_s^k - Y_{s_0}^k = \rho_{s_0,s}^k$.[51] This is the direct gain of going from s_0 to s. It does not include any gains from transitions beyond s:

---

[49]The estimated differences in treatment effects for the conditional and total populations are not large for outcomes associated with the decision to enroll in college, but are substantial for the choice to graduate from college. See Web Appendix A.14.2.

[50]The analysis for discrete outcomes is straightforward.

[51]Note that Y_s^k can be log outcomes as in (1). We can also formulate the outcomes in terms of latent variables.

$$\rho^k_{s_0, s} = Y^k_s - Y^k_{s_0} = \tau^k_s(X) - \tau^k_{s_0}(X) + \theta\,(\alpha^k_s - \alpha^k_{s_0}) + \omega^k_s - \omega^k_{s_0}.$$

In this notation we may write the outcome Y^k relative to base state Y^k_0 as

$$Y^k = Y^k_{s_0} + \sum_s \rho^k_{s_0, s} D_s. \quad (19)$$

This is a version of (1) where schooling is discretized at final schooling attainment levels: $S = s$ if $D_s = 1$. $E(\rho^k_{s_0, s})$ is one version of the returns to schooling compared to benchmark s_0 defined for the entire population.

Except for knife-edge cases, if $\lambda_s \ne 0$, dependence between D_s and $\rho^k_{s_0, s}$ is generated if either $\tau^k_{s_0}(X) \ne \tau^k_s(X)$, or $\alpha^k_{s_0} \ne \alpha^k_s$, or both.[52] Sorting on gains (correlation between $\rho^k_{s_0, s}$ and D_s) may not appear in empirical estimates if agents are sorting on gains beyond s and not on direct effects (i.e., sorting on components of $R_{s,i}$ as defined in (2)). Only in the case where there is no continuation value can we conclude from empirical estimates that absence of sorting effects defined in this fashion implies absence of sorting on potential future gains.

The traditional Griliches (1977) analysis of returns to schooling ignores sorting on gains and only considers ability bias. Assuming analysts condition on X (in levels and in interactions with D_s), sorting gains arise only if $\alpha^k_s - \alpha^k_{s_0} \ne 0$ and $\lambda_s \ne 0$. Even if $\alpha^k_s - \alpha^k_{s_0} = 0$, as long as $\alpha^k_{s_0} \ne 0$, ability bias will arise in estimating the mean of the gains $\rho^k_{s_0, s}$ in (19), provided $\lambda_s \ne 0$.[53]

Note that the choice of a base state matters for estimating sorting gains in the general case where the magnitude of $\alpha^k_s - \alpha^k_{s_0}$ changes depending on the base state selected. Some representations may generate sorting gains that are absent from other representations with different base states.[54]

Within this framework, there are several meaningful ways to decompose the observed difference in outcomes between those at j who go on to $S = j + 1$. The observed difference can be decomposed as follows:

---

[52]Notice that even if there is no such dependence, some agents may still choose to go beyond s because of later gains in outcomes.

[53]For a given level of s, selection bias is defined as $E(Y^k_{s_0} \mid D_s = 1) - E(Y^k_{s_0} \mid D_{s_0} = 1)$, the mean difference in baseline outcomes for persons who stop at $S = s$ compared to those who stop at $S = s_0$.

[54]The extension of this analysis to more general model (5) with the GED is straightforward.

$$\underbrace{E[Y_{j+1}^k \mid S = j + 1] - E[Y_j^k \mid S = j]}_{\text{Observed difference}}$$

$$= \underbrace{E[Y_{j+1}^k - Y_j^k \mid S = j + 1]}_{\text{Treatment on the treated } TT_{j,\,j+1}} + \underbrace{E[Y_j^k \mid S = j + 1] - E[Y_j^k \mid S = j]}_{\substack{\text{Selection bias } SB_{j,\,j+1} \\ \text{from base state } j}}$$

$$= \underbrace{E[Y_{j+1}^k - Y_j^k \mid S \in \{j, j+1\}]}_{\substack{\text{Pairwise average treatment effect } ATE_{j,\,j+1} \\ \text{for people in conditioning set } \{j, j+1\}}} +$$

$$\underbrace{E[Y_{j+1}^k - Y_j^k \mid S = j + 1] - E[Y_{j+1}^k - Y_j^k \mid S \in \{j, j+1\}]}_{\text{Sorting gains } SG_{j,\,j+1}} + \underbrace{E[Y_j^k \mid S = j + 1] - E[Y_j^k \mid S = j]}_{\text{Selection bias } SB_{j,\,j+1}}.$$

[55]   (20)

Note that the ATE parameter depends on the distributions of characteristics of X and θ for persons at node j, as do the sorting on gains and selection bias parameters. These components can be further decomposed into selection on observed variables and selection on unobserved ability components θ, and the ability components can be further decomposed into cognitive and non-cognitive components.

These decompositions focus on gains *up to* final schooling states. They compare observed differences across pairs of final schooling levels. The empirical literature on the returns to schooling also compares the observed differences in outcomes between persons at a given node ($Q_j = 1$) who make a particular schooling transition with those who do not make that transition.

Thus, we can decompose the observed gain from going to $j + 1$ from j for those at j ($Q_j = 1$) into a gain for those who take the transition ($D_j = 0$) and a selection bias term (the difference in the mean outcomes between those who would have gone on ($D_j = 0$), but are stopped at j (Fix $D_j = 1$), and those who chose not to go on). We can further decompose the treatment on the treated parameter into a node-specific ATE (the mean difference between those for whom $Q_j = 1$ where we fix $D_j = 0$ and we fix $D_j = 1$, respectively), and a "sorting gains" term which is the difference between the node-specific treatment on the treated term and the node-specific ATE.

In Web Appendix A.15.3, we decompose the values of being at j into components associated with stopping at j and continuing beyond j where, for the upper branch of Figure 1 ($D_0 = 0$),

---

[55]Appendix A.15.1 gives the exact decomposition for our specific functional forms.

$$Y^k = Y^k_0 + \sum_{j=1}^{\bar{s}} \rho^k_{j-1,j} Q_{j'} \quad (21)$$

where $\rho^k_{j-1,j} = Y^k_j - Y^k_{j-1}$. The expected future gain for a person at $j(\geq 1)$ is

$$E_j\left(\sum_{l>j}^{\bar{s}} \rho^k_{l-1,l} Q_l \mid Q_j = 1 \right) = \sum_{l>j}^{\bar{s}} [E_j(\rho^k_{l-1,l} \mid Q_l = 1) P(Q_l = 1 \mid Q_j = 1)], \quad j \geq 1,$$

where the conditioning $D_0 = 0$ is kept implicit.[56]

Analogous to decomposition (20), we can decompose the observed difference between those with $D_j = 0$ and those with $D_j = 1$, i.e., the observed difference between those that do and do not make a particular transition conditional on making that transition. $E(\rho^k_{j,j+1})$ is the expected incremental gain of proceeding to the next stage. For the upper branch ($D_0 = 0$), we may write for the kth outcome at node j:

$$\underbrace{E[Y^k \mid D_j = 0, Q_j = 1] - E[Y^k \mid D_j = 1, Q_j = 1]}_{\text{Observed difference}}$$

$$= \underbrace{E[Y^k \mid D_j = 0, Q_j = 1] - E[Y^k \mid D_j = 0, Q_j = 1, Fix\, D_j = 1]}_{\text{Dynamic treatment on the treated for those at } j}$$

$$+ \underbrace{E[Y^k \mid D_j = 0, Q_j = 1, Fix\, D_j = 1] - E[Y^k \mid D_j = 1, Q_j = 1]}_{\text{Selection bias for those at } j}$$

$$= \underbrace{E[Y^k \mid Q_j = 1, Fix\, D_j = 0] - E[Y^k \mid Q_j = 1, Fix\, D_j = 1]}_{\text{ATE for those at } j} \quad (22)$$

$$+ \underbrace{\left\{ \begin{array}{l} (E[Y^k \mid D_j = 0, Q_j = 1] - E[Y^k \mid D_j = 0, Q_j = 1, Fix\, D_j = 1]) \\ -(E[Y^k \mid Q_j = 1, Fix\, D_j = 0] - E[Y^k \mid Q_j = 1, Fix\, D_j = 1]) \end{array} \right\}}_{\text{TT} - \text{ATE: Sorting gain at } j \text{ for those who transit to } j+1}$$

$$+ \underbrace{E[Y^k \mid D_j = 0, Q_j = 1, Fix\, D_j = 1] - E[Y^k \mid D_j = 1, Q_j = 1]}_{\text{Selection bias}}.$$

The node-specific ATE is defined for the population at $Q_j = 1$ and considers either forcing population members to stay at j, or moving the entire group from j to $j + 1$ (i.e, $Fix\, D_j = 1$ and $Fix\, D_j = 0$, respectively). The sorting gain is the average net gain beyond ATE to those who actually take the transition ($D_j = 0$).

---

[56]The more general expression incorporating D_0 is presented as Equation (A.10) in the Web Appendix.

4 Identification and Model Likelihood

The treatment effects defined in Section 3 can be identified using alternative empirical approaches. The main approach used in this paper exploits the fact that, conditional on θ, X, Z, outcomes and choices are statistically independent where X and Z are observed and θ is not. If θ were observed, one could condition on θ, X, Z and identify the model of Equations (3)–(9) and the treatment effects that can be generated from it. We use factor model (10) to proxy θ using measurements M.

Under the conditions presented in Heckman et al. (2016), we can non-parametrically identify the model of Equations (3)–(7) including the distribution of θ, as well as the Φ functions and the distribution of e (which can be interpreted as measurement errors). Effectively, we match on proxies for θ and correct for the effects of measurement error (e) in creating the proxies. Such corrections are possible because with multiple measures on θ we can identify the distribution of e.[57] We can identify treatment effects even though we do not isolate individual factors. We only need that the factors θ are spanned by M, not that Equations (10) are separately identified.[58]

Another approach to identification uses instrumental variables which, if available, under the conditions presented in Heckman et al. (2016) can be used to identify the structural model (3)–(9) without invoking the factor structure (8) and (9) or the postulated conditional independence assumptions.

The precise parameterization and the likelihood function for the model we estimate is presented in Web Appendix A.4. While, in principle, it is possible to identify the model non-parametrically, in this paper we make parametric assumptions in order facilitate computation. We subject the estimated model to rigorous goodness-of-fit tests which the model passes.[59]

5 Our Data, A Benchmark OLS Analysis of the Outcomes We Study, and Our Exclusion Restrictions

We estimate our model on a sample of males extracted from the widely-used National Longitudinal Sample of Youth (NLSY 79).[60] Before discussing estimates from our model, it is informative to set the stage for what follows and present adjusted and unadjusted associations between the outcomes we study and schooling. Figure 2 presents estimated linear regression relationships between different levels of schooling relative to high school dropouts and the four outcomes analyzed in this paper: wages, log present value of wages (or PV of wages), health limitations, and smoking.[61] These are least squares regressions

---

[57]Under linear specifications for (10), we can directly estimate the θ and use factor regression methods. See, e.g., Heckman et al. (2013a), Heckman et al. (2016), and the references cited therein.

[58]As noted in Heckman et al. (2011b), we do not need to solve classical identification problems associated with estimating equation system (10) in order to extract measure-preserving transformations of θ on which we can condition in order to identify treatment effects. In the linear factor analysis literature these are the classical rotation and normalization problems.

[59]See Web Appendix A.5.

[60]Web Appendix A.2 presents a detailed discussion of the data we analyze and our exclusion restrictions.

[61]Adjustments are made through linear regression. This decomposition uses high school dropout as the base category (s_0).

using the regressors indicated at the base of the figure, including our proxies for ability. They do not separate out the roles of X and θ in contributing to the causal and selection bias components of the observed differences. Least squares estimates of this form are commonly reported in the literature that investigates the effects of schooling controlling for X, Z, θ.

The black bars in each panel show the unadjusted mean differences in outcomes for persons at the indicated levels of educational attainment compared to those for high school dropouts. Higher ability is associated with higher earnings and more schooling. However, as shown by the grey bars in Figure 2, adjusting for family background and adolescent measures of ability attenuates, but does not eliminate, the estimated least squares estimates of the effects of education.

Figure 2 shows that controlling for proxied ability substantially reduces the observed differences in earnings across educational groups. These regression estimates suggest, but do not identify, substantial causal effects which we report below.

Entering θ as a regressor is a traditional way to control for ability bias. It eliminates the ability bias emphasized by Griliches (1977). If there is sorting on gains that depend on X and θ, this approach over-controls for those variables that are components of the causal effect of treatment on the treated as defined in (20).[62] Figure 2 reports traditional measures of regression-adjusted causal effects of schooling. At the same time, such regressions do not discriminate among the components of (20) which have different causal interpretations. In Web Appendix A.17.2, we compare the OLS pairwise causal effects implicit in the estimates reported in Figure 2 with the estimates from our version of a structural model discussed below.[63]

It is sometimes claimed that a linear-in-years-of-schooling model fits the data well.[64] The white bar in Figure 2 displays the OLS-adjusted effect of schooling controlling for years of completed schooling as in Equation (1).[65] The white bars in all figures show that, even after controlling for years of schooling, the educational indicators still play an important role. OLS estimates of Mincer specification (1) do not precisely describe the data. There are effects of schooling beyond those captured by a linear years-of-schooling specification.[66]

5.1 Control Variables and Exclusion Restrictions

As previously noted, identification of our model and the associated treatment effects does not depend exclusively on conditional independence assumptions associated with our factor model.[67] Node-specific instruments can non-parametrically identify treatment effects without invoking the full set of conditional independence assumptions.[68] We have a variety

---

[62]See also the decompositions in (A.7).

[63]The OLS estimates do not identify treatment on the treated parameters. They are in rough agreement with ATEs except for the log present value of earnings.

[64]See, e.g., Card (1999, 2001). Heckman et al. (2006a) dispute this claim.

[65]Mis-measurement of schooling is less of a concern in our data, as the survey asks numerous educational questions every year which we use to determine an individual's final schooling state.

[66]Using our estimated model, we find, however, that population ATEs are well described by a linear-in-schooling specification. See Web Appendix A.8.

[67]See Carneiro et al. (2003).

[68]See Heckman et al. (2016).

of exclusion restrictions that affect choices but not outcomes. Table 1 documents the control variables (**X**) and the exclusion restrictions (components of **Z** not in **X**) used in this paper. Our instruments are traditional in the literature that estimates the causal effects of education. [69]

6 Estimated Causal Effects

In this section of the paper, we move beyond OLS analyses of causal effects of schooling and present the estimated causal effects of schooling from our model. Since the model is non-linear and multidimensional, in the main body of the paper we only report the treatment effects derived from it.[70] We randomly draw sets of regressors from our sample and a vector of factors from the estimated factor distribution to simulate the reported treatment effects.[71]

Section 6.1 presents estimated treatment effects across final schooling levels. These are based on Equation (18) and extend Becker (1964) by controlling for observed and proxied unobserved variables. Section 6.2 presents the main empirical analysis of this paper. We estimate dynamic treatment effects, inclusive of continuation values. We analyze the contribution of continuation values, sorting on gains, and selection bias to measured differences in education across levels. Section 6.3 analyzes the effects of cognitive and non-cognitive endowments on estimated treatment effects. Section 6.4 presents estimates of distributions of treatment effects. Section 6.5 examines the implications of our analysis for the validity of the Becker-Chiswick-Mincer model. Section 6.6 summarizes our analysis. In the text of our paper, we focus on the transitions in the upper branch of Figure 1, although our model is estimated over both branches.

6.1 The Estimated Average Causal Effect of Educational Choices by Pairwise Final Schooling Levels

We first present estimates of average treatment effects $ATE_{s-1,s}$ (18) for the four outcomes studied in this paper at final schooling level s compared to final schooling level $s-1$.[72] They ignore continuation values.

The shaded regions labeled "Observed" in Figure 3 are the raw differences found in our data. The estimated average causal effects (displayed in the light blocks) are large and statistically significant for all outcomes except for the log PV wages for graduating high school (compared to dropping out).[73] For example, the leftmost bar in panel 3a can be interpreted as follows: while high school graduates make on average 24 log points higher wages than high school dropouts, we find that the average causal effect of graduating high school is on average 12 log points for the same population.

---

[69]For example, presence of a nearby college or distance to college is used by Cameron and Taber (2004), Kling (2001), Carneiro et al. (2013), Cawley et al. (1997), Heckman et al. (2011a), and Eisenhauer et al. (2015b). Local tuition at two- or four-year colleges is used as an instrument by Kane and Rouse (1993), Heckman et al. (2011a), Eisenhauer et al. (2015b), and Cameron and Taber (2004). Local labor market shocks are used by Heckman et al. (2011a) and Eisenhauer et al. (2015b).

[70]Parameter estimates for individual equations are reported in Web Appendix A.6.

[71]We randomly draw an individual and use their full set of regressors.

[72]This is Expression (18) for the case $s' = s + 1$.

[73]For that group, the delay in receiving high school wage rates is not sufficiently compensated by higher wage rates.

Web Appendix A.14.1 reports traditional treatment effects (treatment on the treated, treatment on the untreated, as well as the ATEs displayed in Figure 3). Web Appendix A. 15.2 presents estimates of decomposition (20) for all four outcomes. The decompositions show substantial gains for high-ability persons who graduate college. A large component of the observed difference is properly attributed to selection bias for most outcomes.[74]

6.2 Dynamic Treatment Effects

A major contribution of this paper is the estimation of dynamic treatment effects that include continuation values. These are defined for populations that achieve a node ($Q_j = 1$) which includes people who might go beyond j and even $j + 1$. Specifically, we calculate the average gains to fixing $D_j = 0$ (and possibly going beyond $j + 1$) compared to those at j ($Q_j = 1$) who stop at j ($D_j = 1$). See Equation (11) for the precise expression. Figure 4 plots these treatment effects by the level of educational decision faced by the agent. These treatment effects are also broken down into those for low-ability and high-ability populations using the ability categories defined at the base of the figure. The figure also reports AMTE for individuals at the margin of indifference at each transition.[75]

There are large and statistically significant average causal effects of education for all wage outcomes.[76] Disaggregating by ability, the effects are strong for high-ability people who enroll in college. They are especially strong for those who graduate college. We find little to no evidence of any benefit of graduating college for low-ability individuals.[77] In fact, the point estimates are negative, albeit imprecisely estimated. Although there are wage rate benefits to low-ability people for enrolling in college (Figure 4A), the benefits in terms of the log present value of wages are minimal. For these people, the wage benefits of attending college barely offset the lost work experience and earnings from attending school.

At all levels of education, the estimated AMTE is substantial: there are marginal benefits to additional education at every transition node for individuals at or near the margin of indifference for that transition. The marginal benefits are close to (but generally somewhat below) the average benefits. This is consistent with diminishing benefits of educational expansion.[78] For people at all margins, there are benefits to taking the next transition that are especially pronounced for high school graduation. There are unrealized potential gains in the current system.

We probe more deeply in the Web Appendices. In A.14.2, we present a variety of treatment effects, including treatment on the treated (TT), treatment on the untreated (TUT), and the average treatment effect defined for the entire population, and not just for those at a particular node in our decision tree. This enables us to examine the extent of sorting on gains. In A.15.3, we go further and decompose observed differences in the data into average treatment effects, sorting gains, and selection bias (Equation (22)).

---

[74]For a comparison of the treatment effects implicit in Figure 2 with those implicit in Figure 3, see Web Appendix A.17.2, Tables A71–A74. For most outcomes, the agreement is rather close, except for the ln PV of wages.

[75]We define the margin of indifference to be $\|I_j/\sigma_j\| \quad 0.01$, where σ_j is the standard deviation of I_j.

[76]Across all outcomes, the GED has no benefit.

[77]These estimates are imprecisely determined, in part, because there are few low-ability persons in this category.

[78]A notable exception is for the AMTE for log PV of wages for high school graduation, for which the marginal benefits greatly exceed the average benefits.

Broadly speaking, for wage outcomes we find sorting on gains for college graduation. This arises primarily from the gains to high-ability people documented in Figure 4. We find the reverse pattern for high school graduation: *negative* sorting for graduating high school that is especially pronounced for log present value of our earnings.[79] Consistent with our analysis of AMTE, there are unrealized gains available in the system for a policy of promoting high school graduation for low-ability individuals. For all wage outcomes, there are substantial selection effects, ranging from 50–70% of the observed differences.

The story is different for the estimated educational causal effects on non-market outcomes. We first discuss smoking. There are strong average causal effects on reducing smoking. They are particularly strong for graduating high school for low-ability individuals. Nonetheless, there are substantial negative average causal effects from education on smoking for all nodes. There is little evidence of sorting on gains. Unlike the evidence for wages, there is substantially less evidence of selection bias at any transition.

The evidence for "Health Limits Work" indicates strongly beneficial causal effects for high school graduation, and weak—but generally precisely determined—causal effects for college graduation, which essentially vanish for low-ability persons. There is little evidence of causal effects of attending some college. There is no evidence of sorting on gains. Selection bias is a strong component of observed differences.

6.2.1 Continuation Values—We next decompose the node-specific average treatment effects, just discussed, into continuation value components. Figure 5 presents (in the white bars) the estimated continuation value components, while the bars behind the white boxes are the full treatment effects from Figure 4. We only display continuation values for nodes where these are possible.

For all outcomes except "Health Limits Work," there are large continuation values for high-ability individuals. While returns to high school are roughly the same across ability levels, the mechanisms producing these effects are different. The benefits for low-ability persons come through direct values. The benefits for high-ability persons come through continuation values. For most nodes and treatment effects, the continuation values are statistically significant as indicated by the "x" in Figure 5.[80]

6.3 The Effects on Cognitive and Non-Cognitive Endowments on Treatment Effects

Disaggregating the treatment effects for "high-" and "low-" endowment θ individuals in Figure 4 is a coarse approach. A byproduct of our analysis is that we can determine the contribution of cognitive and non-cognitive endowments (θ) to the explanation of estimated treatment effects. We can decompose the overall effects of θ into their contribution to the causal effects at each node and the contribution of endowments to attaining that node. We find substantial contributions of θ to each component at each node.

---

[79]TUT > ATE > TT.
[80]Note that if direct effects are negative, continuation values may be larger than treatment effects.

To illustrate, the panels in Figure 6 display the estimated average treatment effect of getting a four-year degree (compared to stopping with some college) for each decile pair of cognitive and non-cognitive endowments.[81],[82] Treatment effects, in general, depend on both measures of ability. Moreover, different outcomes depend on the two dimensions of ability in different ways. For example, the treatment effect of graduating college is increasing in both dimensions for the present values of wages, but the reductions in health limitations with education depend mostly on cognitive endowments.

6.4 Distributions of Treatment Effects

One benefit of our approach over the standard IV approach is that we can identify the distributions of expected treatment effects—a feature missing from the standard treatment effect literature. Figure 7 plots the distribution of gains for persons who graduate from college (compared to attending college but not attaining a four-year degree) along with the mean treatment effects.[83]

The graphs provide a nice summary of our main findings for all dynamic treatment effects for college graduation. There are strong causal effects for all outcomes. There is also substantial heterogeneity among persons. Sorting on gains is pronounced for wage outcomes but less so for health and smoking. This is consistent with the analysis in Appendix A.15.3, where we report estimates of sorting on gains.

A byproduct of our analysis is that we can test the rank-invariance of counterfactual outcomes across states. The assumption of rank invariance is the basis for the numerous analyses based on quantile treatment effects.[85] It implies that the Spearman correlations are 1 across any pair of counterfactual states. In our simulations, we find that the Spearman correlations are large but are also not 1. They are between 0.70 and 0.85 for log wages, 0.60 and 0.90 for present value of wages, and notably smaller for smoking and health limitations. [86] Rank invariance is an especially poor assumption for those outcomes.

6.5 Taking Stock of the Becker-Chiswick-Mincer Model

We have tested many features of the widely-used model of Equation (1) to determine its robustness. Some features of Mincer model (1) are broadly consistent with our estimated structural model. While OLS-regression-adjusted versions are not linear in years of schooling (see the evidence in Web Appendix A.8), our estimated ATEs are roughly consistent with linearity for most outcomes.

The correlation between ρ_i and S_i is a centerpiece of the modern IV literature.[87] It varies across transitions (see Web Appendix A.13). However, ρ_i turns out to be node-specific ($\rho_{s,s',i}$) and not the same across transitions.

---

[81]Web Appendix A.7 reports a full set of results across all nodes.
[82]These figures show average benefits by decile over the full population, rather than for the population that reaches each node.
[83]Web Appendix A.9 reports a full set of distributions of treatment effects for all outcomes. Expectations are computed over the idiosyncratic error terms (ω_s^k).

[85]See, e.g., Bitler et al. (2006).
[86]See Web Appendix Tables A22–A25.
[87]This is the correlated random coefficient model. See Heckman and Vytlacil (1998).

Sorting on gains, measured either for specification as in (19) ($COV(\rho_{s,s'}, D_s) \quad 0$), or for specification as in (21) ($COV(\rho_{j-1,j}, Q_j) \quad 0$), reveals that there is positive sorting on wage gains only at the higher levels of education. Our estimated correlation patterns are consistent with our evidence on sorting gains presented in Web Appendix A.15.

6.6 Summarizing Our Analysis of Causal Effects of Education

In this section and in our Web Appendix, we have analyzed a variety of economically interpretable treatment effects. We reach the following broad conclusions.

1. There are substantial causal benefits for all outcomes analyzed from education, except for GED certification.

2. Continuation values are an important component of causal effects for most outcomes except health limits work.

3. There are substantial benefits from graduating high school that are especially strong for the less able, many of whom currently do not graduate. This suggests strong gains from programs promoting high school graduation.

4. For the wage outcomes we study, there is evidence on sorting on gains from graduating college for high-ability persons.[88] There are no causal effects of college graduation for low-ability persons. College graduation is not for all.

5. There are strong benefits of education for those at the margin of indifference at all nodes. These are largely direct effects with little contribution from continuation values.

6. We estimate strong causal effects for the non-monetary outcomes studied. They are particularly strong for high school graduation. There is little evidence of sorting on gains in either non-monetary outcome examined. Continuation values are largely absent for our measure of health. For smoking, continuation values are most pronounced among higher-ability persons. Selection bias is less empirically important for smoking, but is substantial for health limits work.

7 Policy Simulations from Our Model

Using our model, it is possible to conduct a variety of counterfactual policy simulations, a feature not shared by standard treatment effect models. We achieve these results without imposing strong assumptions on the choice model. We consider two policy experiments: (i) a tuition subsidy; and (ii) an increase in the cognitive and non-cognitive endowments of those at the bottom of the endowment distribution. The first policy experiment is similar to what is estimated by LATE only in the special case where the instrument in LATE corresponds to the exact policy experiment. The second policy experiment is of interest because early childhood programs boost these endowments (Heckman et al., 2013a). Neither

---

[88]Part of the relationship between ability and returns to college could operate through college quality. For example, Dillon and Smith (2015) show that ability is an important determinant of college quality and that college quality improves wages even after controlling for ability.

set of counterfactuals generated can be estimated from instrumental variable estimands. We ignore general equilibrium effects in these simulations.

7.1 Policy-Relevant Treatment Effects

Unless the instruments correspond to policies, IV does not identify policy-relevant treatment effects. The PRTE allows us to identify who would be induced to change educational choices under specific policy changes, and how these individuals would benefit on average. As an example of the power of our methodology, we simulate the response to a policy intervention that provides a one standard deviation subsidy to early college tuition (approximately $850 dollars per year of college). Column 1 of Table 2 presents the average treatment effect (including continuation values) in our estimated model for those who are induced to change education levels by the tuition subsidy. Since we do not find evidence that college tuition affects high school graduation rates, the subsidy only induces high school graduates to change their college enrollment decisions and does not affect high school graduation decisions. Those induced to enroll may then go on to graduate with a four-year degree.[89] Columns 2 and 3 of Table 2 decompose the PRTE into the average gains for those induced to enroll and then go on to earn four-year degrees and the average gains for those who do not. For the most part, the PRTE is larger for those who go on to earn four-year degrees.

Figure 8 shows which individuals are induced to enroll in college within the deciles of the distribution of the unobservable in the choice equation for node 2,[90] conditional on $Q_2 = 1$ (the node determining college enrollment). These are the unobserved components of heterogeneity acted upon by the agent but unobserved by the economist.

The policy induces some individuals at every decile to switch, but places more weight on those in the middle deciles of the distribution. The figure decomposes the effect of those induced to switch into effects for those who go on to graduate with four-year degrees and effects for those who do not. Those induced to switch in the top deciles are more likely to go on to graduate with a four-year college degree.

The $850 subsidy induces 12.8% of high school graduates who previously did not attend college to enroll in college. Of those induced to enroll, more than a third go on to graduate with a four-year degree. For outcomes such as smoking, the benefits are larger for those who graduate with a four-year degree. The large gains for marginal individuals induced to enroll is consistent with the empirical literature, that finds large psychic costs are necessary to justify college schooling choices, and the failure of agents to respond to strong monetary incentives.

Using the estimated benefits, we can determine if the monetary gains in the present value of wages at age 18 is greater than the $850 subsidy.[91] Given a PRTE of 0.13 for log present

---

[89]Models were estimated that include tuition as a determinant of the high school graduation decision. However, the estimated effects of tuition on high school graduation are small and statistically insignificant. We do not impose the requirement that future values of costs affect current educational choices. This highlights the benefits of our more robust approach. We do not impose the requirement that agents know and act on publicly available information.

[90]The unobservable is the bundle $\eta_1 = -(\boldsymbol{\theta}'\boldsymbol{\lambda}_1 - \nu_1)$.

[91]However, a limitation of our model is that we can only estimate the monetary costs and do not estimate psychic costs.

value of wage income, the average gains for those induced to enroll is $36,401 in year 2000 dollars. If the subsidy is given for the first two years of college, then the policy clearly leads to monetary gains for those induced to enroll. If the subsidy is also offered to those already enrolled, the overall monetary costs of the subsidy is much larger because it is given to more than 8 students previously enrolled for each new student induced to enroll (dead weight).

7.2 Boosting Cognitive and Non-Cognitive Endowments

Using simulation methods, it is possible to construct counterfactual policy simulations unrelated to any particular set of instruments. For example, some early childhood programs have been shown to have lasting impacts on the cognitive and non-cognitive endowments of low-ability children (see Heckman et al., 2013a). We simulate two policy experiments: (i) increasing the cognitive endowment of those in the lowest decile; and (ii) increasing the non-cognitive endowment of those in the lowest decile.[92]

The panels of Figure 9 show the average gains for increasing the cognitive or non-cognitive endowments of those in the lowest decile of each ability. Increased cognition helps individuals across the board. Increasing socio-emotional endowments has a smaller effect on labor market outcomes but has substantial effects on health.[93]

8 Testing the Two-Factor Assumption

Throughout this paper, we have assumed that selection of outcomes occurs on the basis of a two-component vector θ, where the components can be proxied by our measures of cognitive and non-cognitive endowments. An obvious objection to this approach is that there may be unproxied endowments that affect both choices and outcomes that we do not measure. For example, one could imagine that a component of the idiosyncratic error terms in the educational choices (v_j) represent taste for schooling. This could generate correlations between the unobservables in the different educational choices and bias our results.

In order to test for the presence of a third factor that influences both choices and outcomes, we test whether the simulated model fits the sample covariances between Y^k and D_j, $j = 1$, …, $\bar{s}$, $k = 1$, …, 4. If an important third factor common to both outcome and choice equations has been omitted, the sample fit should be poor. In fact, we find a good fit.

Cunha and Heckman (2016) estimate a related model using the same data source. They find that a three-factor model explains wages and present value of wages. Two of their factors correspond to the factors used in this paper. Their third factor improves the fit of the wage outcome data but does *not* enter agent decision equations or affect selection or sorting bias. Our evidence is consistent with their findings.

---

[92]The details of how these simulations were conducted are presented in Web Appendix A.11. Our model does not address general equilibrium effects of such a change in the endowment distribution.
[93]We present additional policy simulations in Web Appendix A.11.

9 Comparisons with Simple Treatment Effect Estimators

In this paper, we have exploited the assumption of conditional independence of outcomes and choices given X, Z, θ. This raises the question of how similar our results would be if we had used simple matching and regression methods.

Table 3 presents two sets of estimates for the models discussed in detail at the base of the table. The first four columns of numbers are node-specific linear-regression estimators of Equation (11), using as regressors the background variables reported in Table 1 but not the "exclusion restriction" variables. The first column of estimates come from a model without any control for θ. The estimates for the other three models control for θ in various ways as noted at the base of Table 3.

We use two versions of nearest-neighbor matching estimators based on the full set of control variables listed in Table 1. Details of the matching procedures are given at the base of Table 3 and in Web Appendix A.18.

The OLS estimates differ greatly from model estimates when there is no adjustment for ability. Controlling for ability has substantial effects on the estimated average treatment effects. Across schooling nodes, all of the estimates that control for θ are "within the ball park" of the estimates produced from our model, although some discrepancies are substantial. This is good news for applied economists mainly interested in using simple methods to estimate node-specific average treatment effects. However, these simple methods do not estimate decision rules, do not enable analysts to estimate AMTE and PRTE, or address many of the other questions addressed in this paper.[94]

10 Summary and Conclusion

Gary Becker's pioneering research on human capital launched a large and active industry estimating causal effects and returns to schooling. Multiple methodological approaches have been used to secure these estimates ranging from reduced-form treatment effect methods to fully structural methods. Each methodology has its benefits and limitations.

The early literature on human capital ignored the dynamics of schooling choices. This paper develops and estimates a robust dynamic model of schooling and its causal consequences for earnings, health, and smoking. Our model recognizes the sequential dynamic nature of educational decisions. We borrow features from both the reduced-form treatment effect literature and the structural literature. Our estimated model passes a variety of goodness-of-fit and model specification tests.

We allow agents to be irrational and myopic in making schooling decisions. Hence, we can use our model to test some of the rationality and information processing assumptions maintained in the dynamic discrete choice literature on education.

---

[94]Table A70 of the Web Appendix compares OLS estimates of dynamic treatment effects and continuation values with our model estimates. The OLS estimates are "within ballpark" for smoking and health limits work, but they are wide off the mark for wages and PV wages.

We use our dynamic choice model to estimate causal effects arising from multiple levels of schooling rather than just the binary comparisons typically featured in the literature on treatment effects and in many structural papers.[95] By estimating a sequential model of schooling in a unified framework, we are able to analyze the *ex post* returns to education for people at different margins of choice and analyze a variety of economically interesting policy counterfactuals. We are able to characterize who benefits from education for a variety of market and non-market outcomes.

We decompose the benefits of schooling at different levels into direct components and indirect components arising from continuation values. We estimate substantial continuation value components of graduating high school and completing college for high-ability individuals. For them, schooling opens up valuable options for future schooling. Standard estimates of the benefits of education based only on direct components of dynamic treatment effects underestimate the full benefits of education. For low-ability individuals, there are substantial direct effects of graduating high school, but little continuation value.

Without imposing rationality, we nonetheless find evidence consistent with it. We find positive sorting into schooling based on gains, especially for higher schooling levels. Schooling has strong causal effects on earnings, health, and healthy behaviors. Both cognitive and non-cognitive endowments affect schooling choices and outcomes at each level of schooling.

We link the structural and matching literatures using conditional independence assumptions. We investigate how simple methods used in the treatment effect literature perform in estimating average treatment effects.[96] They roughly approximate our model estimates of average treatment effects, provided we condition on endowments of cognitive and non-cognitive skills. However, these simple methods *do not* identify the treatment effects for persons at the margins of different choices (the average marginal treatment effects).[97] We test the empirical foundations of the Mincer model and find it wanting. A richer specification of the schooling earnings decision is warranted to generate empirically supported estimates of causal effects.

We use our estimated model to conduct two policy experiments. We determine the groups that benefit from a tuition reduction policy and what those benefits are. We also examine how the impact of boosts in cognitive and non-cognitive skills affects educational choices and outcomes.

Our analysis enriches the pioneering analysis of Becker (1964). The early research on human capital was casual about agent heterogeneity. It ignored selection bias and sorting gains from schooling. Later work by Griliches (1977) focused on selection bias ("ability bias"), but ignored sorting gains. In this paper, we quantify both components of outcome equations. We find evidence of selection bias at all levels of schooling for all outcomes and sorting gains at higher levels of schooling for wage outcomes.

---

[95]See, e.g., Willis and Rosen (1979).
[96]IV estimates are very different from our model estimates. See Heckman et al. (2016).
[97]We can roughly approximate continuation values using simple methods. See Table A70 in the Web Appendix.

Our findings thus support the basic insights of Becker (1964). Schooling has strong causal effects on market and non-market outcomes. Both cognitive and non-cognitive endowments affect schooling choices and outcomes. People sort into schooling based on realized incremental gains.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

References

Abbring Jaap H, Heckman James J. Econometric Evaluation of Social Programs, Part III: Distributional Treatment Effects, Dynamic Treatment Effects, Dynamic Discrete Choice, and General Equilibrium Policy Evaluation. In: Heckman James J, Leamer Edward E, editorsHandbook of Econometrics. Vol. 6B. Amsterdam: Elsevier Science B. V; 2007. 5145–5303. chap. 72

Adda JéerômeCooper Russell W. Dynamic Economics: Quantitative Methods and Applications. Cambridge, MA: The MIT Press; 2003.

Almlund MathildeDuckworth AngelaHeckman James J, Kautz Tim. Personality Psychology and Economics. In: Hanushek Eric A, Machin StephenWößmann Ludger, editorsHandbook of the Economics of Education. Vol. 4. Amsterdam: Elsevier; 2011. 1–181. chap. 1

Altonji Joseph G. The Demand for and Return to Education When Education Outcomes Are Uncertain. Journal of Labor Economics. 1993; 11(1):48–83.

Angrist Joshua D, Imbens Guido W. Two-Stage Least Squares Estimation of Average Causal Effects in Models with Variable Treatment Intensity. Journal of the American Statistical Association. 1995; 90(430):431–442.

Angrist Joshua D, Pischke Jörn-Steffan. Mostly Harmless Econometrics: An Empiricist's Companion. Princeton, NJ: Princeton University Press; 2009.

Arcidiacono PeterMiller Robert A. Conditional Choice Probability Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity. Econometrica. 2011; 7(6):1823–1868.

Bamberger Gustavo. PhD thesis. University of Chicago, Graduate School of Business; 1987. Occupational Choice: the Role of Undergraduate Education.

Bartlett Maurice S. The Statistical Conception of Mental Factors. British Journal of Psychology. 1937; 28(1):97–104.

Bartlett Maurice S. Methods of Estimating Mental Factors. Nature. 1938; 141:609–610.

Becker Gary S. Investment in Human Capital: A Theoretical Analysis. Journal of Political Economy. 1962; 70(5 Part 2):9–49.

Becker Gary S. Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education. 1. Chicago: University of Chicago Press for the National Bureau of Economic Research; 1964.

Becker Gary S, Chiswick Barry R. Education and the Distribution of Earnings. American Economic Review. 1966; 56(1/2):358–369.

Bitler Marianne P, Gelbach Jonah B, Hoynes Hilary W. What Mean Impacts Miss: Distributional Effects of Welfare Reform Experiments. American Economic Review. 2006; 96(4):988–1012.

Blevins Jason R. Nonparametric Identification of Dynamic Decision Processes with Discrete and Continuous Choices. Quantitative Economics. 2014; 5(3):531–554.

Borghans LexDuckworth Angela L, Heckman James J, ter Weel Bas. The Economics and Psychology of Personality Traits. Journal of Human Resources. 2008; 43(4):972– 1059.

Cameron Stephen V, Heckman James J. The Nonequivalence of High School Equivalents. Journal of Labor Economics. 1993; 11(1 Part 1):1–47.

Cameron Stephen V. The Dynamics of Educational Attainment for Black, Hispanic, and White Males. Journal of Political Economy. 2001; 109(3):455–499.

Cameron Stephen V, Taber Christopher. Estimation of Educational Borrowing Constraints Using Returns to Schooling. Journal of Political Economy. 2004; 112(1):132–182.

Card David. The Causal Effect of Education on Earnings. In: Ashenfelter Orley C, Card David, editorsHandbook of Labor Economics. Vol. 3A. Amsterdam: Elsevier Science B.V; 1999. 1801– 1863. chap. 30

Card David. Estimating the Return to Schooling: Progress on Some Persistent Econometric Problems. Econometrica. 2001; 69(5):1127–1160.

Carneiro PedroHansen KarstenHeckman James J. Estimating Distributions of Treatment Effects with an Application to the Returns to Schooling and Measurement of the Effects of Uncertainty on College Choice. International Economic Review. 2003; 44(2):361–422.

Carneiro PedroHeckman James J, Vytlacil Edward J. Evaluating Marginal Policy Changes and the Average Effect of Treatment for Individuals at the Margin. Econometrica. 2010; 78(1):377–394. [PubMed: 20209119]

Carneiro PedroJames J. Estimating Marginal Returns to Education. American Economic Review. 2011; 101(6):2754–2781. [PubMed: 25110355]

Carneiro PedroMeghir CostasParey Matthias. Maternal Education, Home Environments, and the Development of Children and Adolescents. Journal of the European Economic Association. 2013; 11(S1):123–160.

Cawley JohnConneely KarenHeckman James J, Vytlacil Edward J. Cognitive Ability, Wages, and Meritocracy. In: Devlin BernieFienberg Stephen E, Resnick Daniel P, Roeder Kathryn, editorsIntelligence, Genes, and Success: Scientists Respond to The Bell Curve. Vol. chap. 8. New York: Springer Verlag; 1997. 179–192.

Comay YochananMelnik AriePollatschek Moshe A. The Option Value of Education and the Optimal Path for Investment in Human Capital. International Economic Review. 1973; 14(2):421–435.

Cunha FlávioHeckman James J. Decomposing Trends in Inequality in Earnings into Forcastable and Uncertain Components. Journal of Labor Economics. 2016 Forthcoming.

Cunha FlávioHeckman James J, Navarro Salvador. The Identification and Economic Content of Ordered Choice Models with Stochastic Cutoffs. International Economic Review. 2007; 48(4): 1273–1309.

Cutler David M, Lleras-Muney Adriana. Understanding Differences in Health Behaviors by Education. Journal of Health Economics. 2010; 29(1):1–28. [PubMed: 19963292]

Dillon Eleanor WiskeSmith Jeffrey Andrew. IZA Discussion Paper 9080. Institute for the Study of Labor; Bonn: 2015. The Consequences of Academic Match between Students and Colleges.

Dothan UriWilliams Joseph. Education as an Option. The Journal of Business. 1981; 54(1):117–139.

Eckstein ZviWolpin Kenneth I. The Specification and Estimation of Dynamic Stochastic Discrete Choice Models: A Survey. Journal of Human Resources. 1989; 24(4):562–598.

Eisenhauer PhilippHeckman James J, Mosso Stefano. Estimation of Dynamic Discrete Choice Models by Maximum Likelihood and the Simulated Method of Moments. International Economic Review. 2015a; 56(2):331–357. [PubMed: 26494926]

Eisenhauer PhilippHeckman James J, Vytlacil Edward J. Generalized Roy Model and Cost-Benefit Analysis of Social Programs. Journal of Political Economy. 2015b; 123(2):413–433. [PubMed: 26709315]

Geweke JohnKeane Michael. Computationally Intensive Methods for Integration in Econometrics. In: Heckman James J, Leamer Edward E, editorsHandbook of Econometrics. Vol. 5. Amsterdam: Elsevier Science B. V; 2001. 3463–3568. chap. 56

Griliches Zvi. Estimating the Returns to Schooling: Some Econometric Problems. Econometrica. 1977; 45(1):1–22. [PubMed: 12278790]

Grossman Michael. The Human Capital Model. In: Culyer Anthony J, Newhouse Joseph P, editorsHandbook of Health Economics. Vol. 1. Amsterdam: Elsevier Science B. V; 2000. 347–408. chap. 7

Haavelmo Trygve. The Statistical Implications of a System of Simultaneous Equations. Econometrica. 1943; 11(1):1–12.

Heckman James J. Unpublished manuscript. University of Chicago, Department of Economics; 1981. The Empirical Content of Alternative Models of Labor Earnings.

Heckman James J. Micro Data, Heterogeneity, and the Evaluation of Public Policy: Nobel Lecture. Journal of Political Economy. 2001; 109(4):673–748.

Heckman James J. Econometric Causality. International Statistical Review. 2008; 76(1):1–27.

Heckman James J. Building Bridges Between Structural and Program Evaluation Approaches to Evaluating Policy. Journal of Economic Literature. 2010; 48(2):356–398. [PubMed: 21743749]

Heckman James J, Carneiro PedroVytlacil Edward. Estimating Marginal Returns to Education. American Economic Review. 2011a; 101(6):2754–2871. [PubMed: 25110355]

Heckman James J, Humphries John EricKautz Tim, editorsThe Myth of Achievement Tests: The GED and the Role of Character in American Life. Chicago: University of Chicago Press; 2014.

Heckman James J, Humphries John EricVeramendi Gregory. Dynamic Treatment Effects. Journal of Econometrics. 2016; 191(2):276–292. [PubMed: 27041793]

Heckman James J, Ichimura HidehikoTodd Petra E. Matching as an Econometric Evaluation Estimator. Review of Economic Studies. 1998; 65(2):261–294.

Heckman James J, Lochner Lance J, Todd Petra E. Earnings Functions, Rates of Return and Treatment Effects: The Mincer Equation and Beyond. In: Hanushek Eric A, Welch Frank, editorsHandbook of the Economics of Education. Vol. 1. Amsterdam: Elsevier; 2006a. 307–458. chap. 7

Heckman James J. Earnings Functions and Rates of Return. Journal of Human Capital. 2008; 2(1):1–31.

Heckman James J, Navarro Salvador. Dynamic Discrete Choice and Dynamic Treatment Effects. Journal of Econometrics. 2007; 136(2):341–396.

Heckman James J, Pinto Rodrigo. Causal Analysis after Haavelmo. Econometric Theory. 2015; 31(1):115–151. [PubMed: 25729123]

Heckman James J, Pinto RodrigoSavelyev Peter A. Understanding the Mechanisms Through Which an Influential Early Childhood Program Boosted Adult Outcomes. American Economic Review. 2013a; 103(6):2052–2086. [PubMed: 24634518]

Heckman James J, Schennach Susanne M, Williams Benjamin. Unpublished manuscript. University of Chicago, Department of Economics; 2011b. Matching with Error-Laden Covariates.

Heckman James J. Unpublished Manuscript. University of Chicago, Department of Economics; 2013b. Matching on Proxy Variables.

Heckman James J, Smith Jeffrey A, Clements Nancy. Making the Most Out Of Programme Evaluations and Social Experiments: Accounting for Heterogeneity in Programme Impacts. Review of Economic Studies. 1997; 64(4):487–535.

Heckman James J, Stixrud JoraUrzúa Sergio. The Effects of Cognitive and Noncognitive Abilities on Labor Market Outcomes and Social Behavior. Journal of Labor Economics. 2006b; 24(3):411–482.

Heckman James J, Urzúa Sergio. Comparing IV With Structural Models: What Simple IV Can and Cannot Identify. Journal of Econometrics. 2010; 156(1):27–37. [PubMed: 20440375]

Heckman James J, Urzúa SergioVytlacil Edward J. Understanding Instrumental Variables in Models with Essential Heterogeneity. Review of Economics and Statistics. 2006c; 88(3):389–432.

Heckman James J, Vytlacil Edward J. Instrumental Variables Methods for the Correlated Random Coefficient Model: Estimating the Average Rate of Return to Schooling When the Return Is Correlated with Schooling. Journal of Human Resources. 1998; 33(4):974–987.

Heckman James J. Local Instrumental Variables and Latent Variable Models for Identifying and Bounding Treatment Effects. Proceedings of the National Academy of Sciences. 1999; 96(8):4730–4734.

Heckman James J. Structural Equations, Treatment Effects and Econometric Policy Evaluation. Econometrica. 2005; 73(3):669–738.

Heckman James J. Econometric Evaluation of Social Programs, Part I: Causal Models, Structural Models and Econometric Policy Evaluation. In: Heckman James J, Leamer Edward E, editorsHandbook of Econometrics. Vol. 6B. Vol. chap. 70. Amsterdam: Elsevier Science B. V; 2007a. 4779–4874.

Heckman James J. Econometric Evaluation of Social Programs, Part II: Using the Marginal Treatment Effect to Organize Alternative Economic Estimators to Evaluate Social Programs and to Forecast Their Effects in New Environments. In: Heckman James J, Leamer Edward E, editorsHandbook of Econometrics. Vol. 6B. Amsterdam: Elsevier Science B. V; 2007b. 4875–5143. chap. 71

Imbens Guido W, Angrist Joshua D. Identification and Estimation of Local Average Treatment Effects. Econometrica. 1994; 62(2):467–475.

Kane Thomas J, Rouse Cecilia E. Working Paper 4268. National Bureau of Economic Research; 1993. Labor Market Returns to Two- and Four-Year Colleges: Is a Credit a Credit and Do Degrees Matter?.

Keane Michael P, Wolpin Kenneth I. The Career Decisions of Young Men. Journal of Political Economy. 1997; 105(3):473–522.

Kling Jeffrey R. Interpreting Instrumental Variables Estimates of the Returns to Schooling. Journal of Business and Economic Statistics. 2001; 19(3):358–364.

Lochner Lance. Working Paper 16722. National Bureau of Economic Research; 2011. Non-Production Benefits of Education: Crime, Health, and Good Citizenship.

Manski Charles F. Measuring Expectations. Econometrica. 2004; 72(5):1329–1376.

McMahon Walter W. The Appraisal of Investments in Educational Facilities. Paris: European Investment Bank/OECD; 2000. Externalities, Non-Market Effects, and Trends in Returns to Educational Investments; 51–83.

McMahon Walter W. Higher Learning, Greater Good. Baltimore, MD: Johns Hopkins University Press; 2009.

Mincer Jacob. Schooling, Experience, and Earnings. New York: Columbia University Press for National Bureau of Economic Research; 1974.

Oreopoulos PhilipPetronijevic Uros. Working Paper 19053. National Bureau of Economic Research; 2013. Making College Worth It: A Review of Research on the Returns to Higher Education.

Oreopoulos PhilipSalvanes Kjell G. Priceless: The Nonpecuniary Benefits of Schooling. Journal of Economic Perspectives. 2011; 25(1):159–184.

Quandt Richard E. The Estimation of the Parameters of a Linear Regression System Obeying Two Separate Regimes. Journal of the American Statistical Association. 1958; 53(284):873–880.

Quandt Richard E. A New Approach to Estimating Switching Regressions. Journal of the American Statistical Association. 1972; 67(338):306–310.

Rust John. Structural Estimation of Markov Decision Processes. In: Engle Robert F, McFadden Daniel L, editorsHandbook of Econometrics. Vol. 4. New York, NY: North-Holland; 1994. 3081–3143. chap. 51

Schennach Susanne M, White HalbertChalak Karim. Local Indirect Least Squares and Average Marginal Effects in Nonseparable Structural Systems. Journal of Econometrics. 2012; 166(2):282–302.

Stange Kevin M. An Empirical Investigation of the Option Value of College Enrollment. American Economic Journal: Applied Economics. 2012; 4(1):49–84.

Vytlacil Edward J. Independence, Monotonicity, and Latent Index Models: An Equivalence Result. Econometrica. 2002; 70(1):331–341.

Weisbrod Burton A. Education and Investment in Human Capital. Journal of Political Economy. 1962; 70(5 Part 2):106–123. Investment in Human Beings.

White HalbertChalak Karim. Settable Systems: An Extension of Pearl's Causal Model with Optimization, Equilibrium, and Learning. Journal of Machine Learning Research. 2009; 10:1759–1799.

Willis Robert J, Rosen Sherwin. Education and Self-Selection. Journal of Political Economy. 1979; 87(5 Part 2):S7–S36.
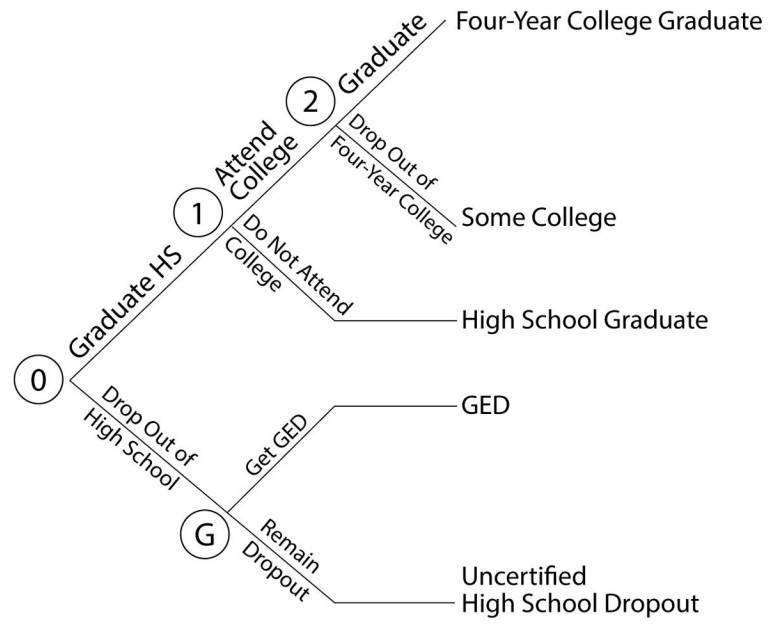
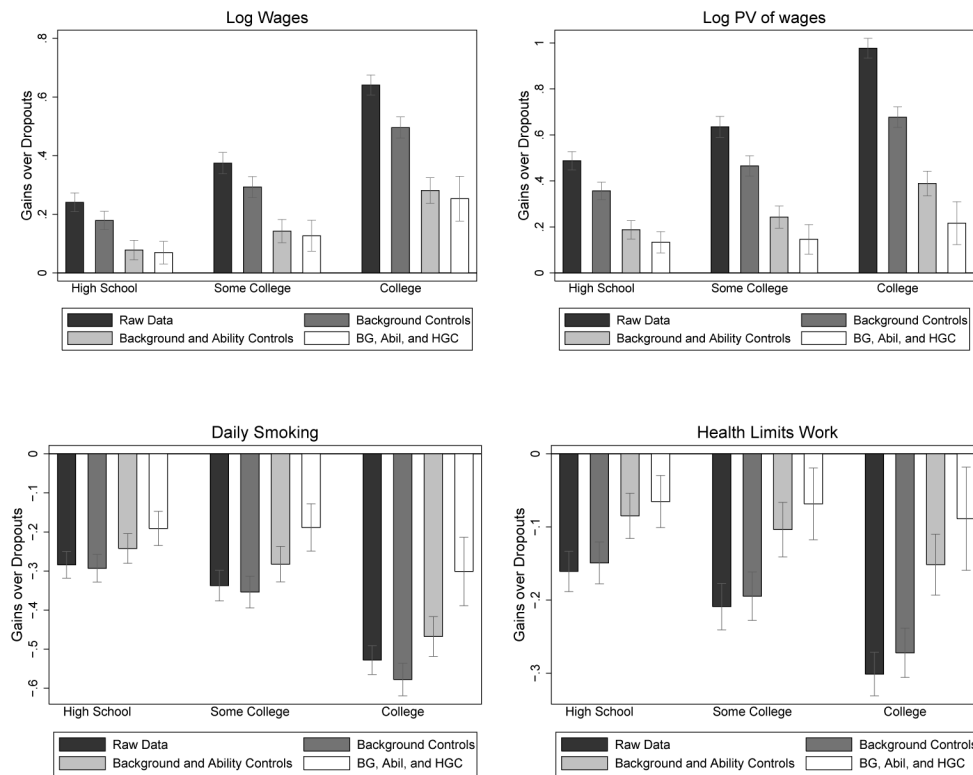**Figure 1.**
A Multistage Dynamic Decision Model

Figure 2.

Observed and Adjusted Benefits from Education

Notes: The bars represent the coefficients from a regression of the designated outcome on dummy variables for educational attainment, where the omitted category is high school dropout. Regressions are run adding successive controls for background and proxies for ability. Background controls include race, age in 1979, region of residence in 1979, urban status in 1979, broken home status, number of siblings, mother's education, father's education, and family income in 1979. Proxies for ability are average score on the Armed Services Vocational Aptitude Battery (ASVAB) tests and ninth grade GPA in core subjects (language, math, science, and social science). See the discussion surrounding Table 1 (below) and Web Appendix A.2 for additional details. "Some College" includes anyone who enrolled in college, but did not receive a four-year college degree. The white bars additionally control for highest grade completed (HGC). *Source:* NLSY79 data.
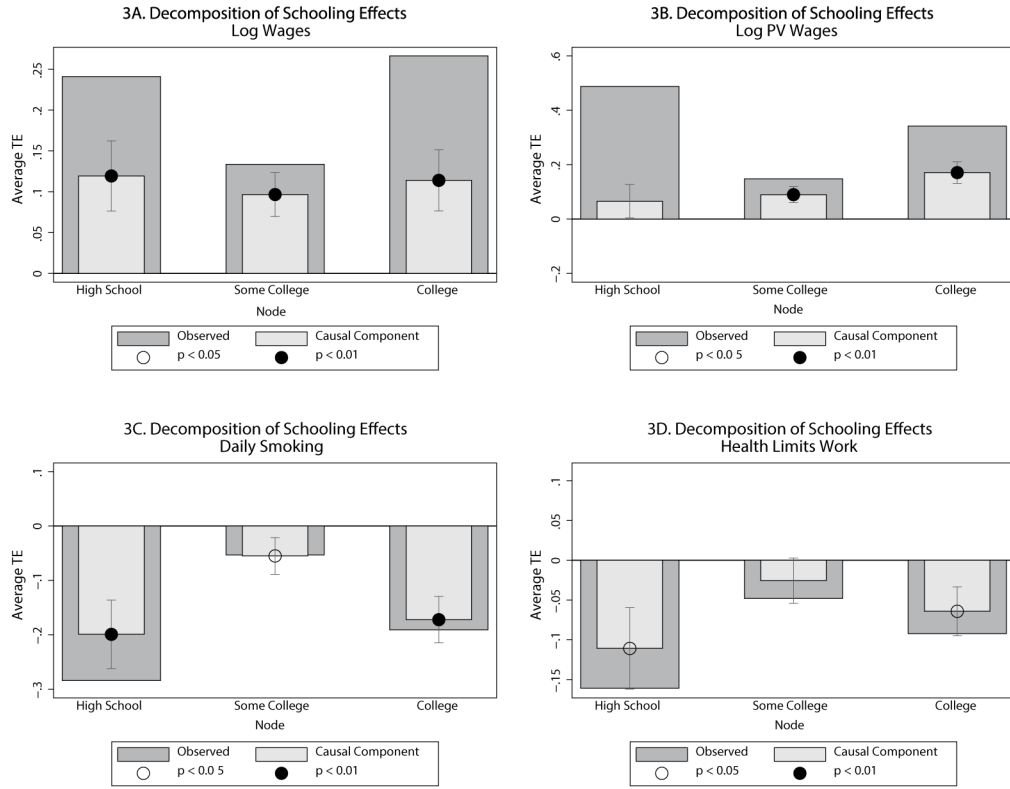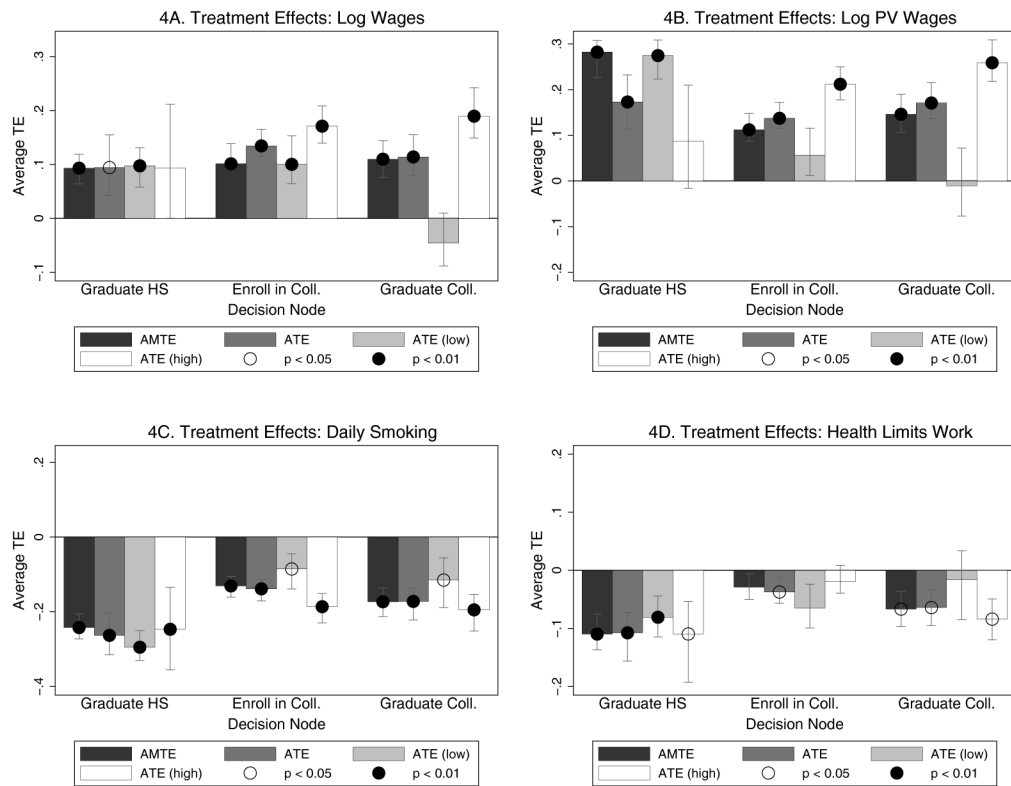
Figure 3.
Causal Versus Observed Differences by Final Schooling Level (compared to next lowest level)

Notes: These figures report pairwise treatment effect (18) for the indicated schooling nodes. Each bar compares the mean outcomes from a particular schooling level *j* and the next lowest level *j*–1 defined for the set of persons who complete schooling at *j*–1 or *j*. The "Observed" bar displays the observed differences in the data. The "Causal Component" bar displays the estimated average treatment effect to those who get treated (ATE) for the indicated group. The difference between the observed and causal treatment effect is attributed to the effect of selection and ability. Selection includes sorting on gains. The error bars and significance levels for the estimated ATE are calculated using 200 bootstrap samples. Error bars show one standard deviation and correspond to the 15.87th and 84.13th percentiles of the bootstrapped estimates, allowing for asymmetry. Significance at the 5% and 1% levels is shown by open and filled circles on the plots, respectively.

4A. Treatment Effects: Log Wages

4B. Treatment Effects: Log PV Wages

4C. Treatment Effects: Daily Smoking

4D. Treatment Effects: Health Limits Work

Sorting on Ability

	Low Ability	High Ability
D_1: Dropping from HS vs. Graduating from HS	0.31	0.31
D_2: **HS Graduate** vs. College Enrollment	0.22	0.38
D_3: **Some College** vs. **Four-Year College Degree**	0.13	0.51

Figure 4.

Treatment Effects of Outcomes by Decision Node

$E(Y^k | \text{Fix } D_j = 0, Q_j = 1) - E(Y^k | \text{Fix } D_j = 1, Q_j = 1)$

Notes: The nodes in the table correspond to the next stage of the transition analyzed. Thus, "Graduate HS" refers to the decision node of whether or not the agent will graduate high school, and refers to the base state of not graduating high school. The error bars and significance levels for the estimated ATE Equation (13) are calculated using 200 bootstrap samples. Error bars show one standard deviation and correspond to the 15.87th and 84.13th percentiles of the bootstrapped estimates, allowing for asymmetry. Significance at the 5% and 1% level are shown by hollow and black circles on the plots, respectively. The figure reports various treatment effects for those who reach the decision node, including the estimated ATE conditional on endowment levels. The high- (low-) ability group is defined as those individuals with cognitive and socio-emotional endowments above (below) the median in the overall population. These categories are not mutually exclusive, as some people may be high-ability in one dimension but low-ability in another. The table below the figure shows the proportion of individuals at each decision ($Q_j = 1$) that are high- and low-ability. The

larger proportion of the individuals are high-ability and a smaller proportion are low-ability in later educational decisions. In this table, final schooling levels are highlighted using bold letters.
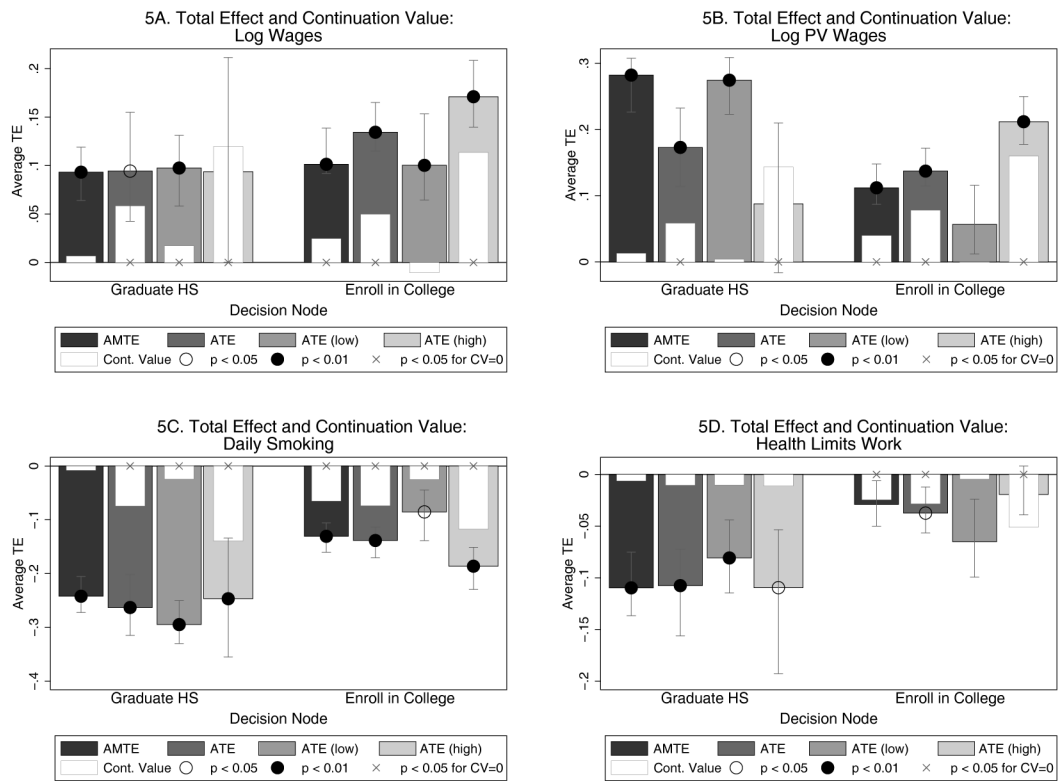
Figure 5.
Dynamic Treatment Effects:

Continuation Values and Total Treatment Effects by Node

Notes: High-ability individuals are those in the top 50% of the distributions of both cognitive and socio-emotional endowments. Low-ability individuals are those in the bottom 50% of the distributions of both cognitive and socio-emotional endowments. The error bars and significance levels for the estimated ATE are calculated using 200 bootstrap samples. Error bars show one standard deviation and correspond to the 15.87th and 84.13th percentiles of the bootstrapped estimates, allowing for asymmetry. Significance at the 5% and 1% level are shown by hollow and black circles on the plots, respectively. Statistical significance for continuation values at the 5% level are shown by "x." Section 3 provides details on how the continuation values and treatment effects are defined.
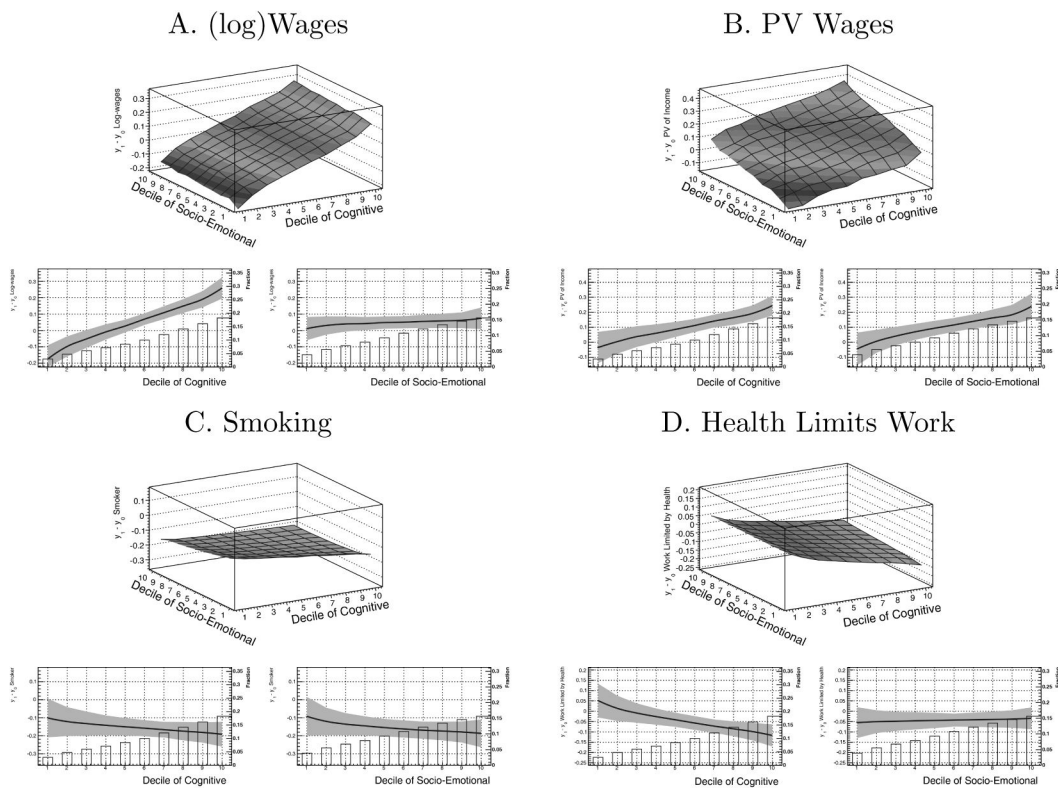
Figure 6.

Average Treatment Effect of Graduating from a Four-Year College by Outcome

Notes: Each panel in this figure studies the average effects of graduating with a four-year college degree on the outcome of interest. The effect is defined as the differences in the outcome between those with a four-year college degree and those with some college. For each panel, let $Y_{\text{some college}}$ and $Y_{\text{four-year degree}}$ denote the outcomes associated with attaining some college and graduating with a four-year degree, respectively. For each outcome, the first figure (top) presents $E(Y_{\text{four-year degree}} - Y_{\text{some college}}|d^C, d^{SE})$ where d^C and d^{SE} denote the cognitive and socio-emotional deciles computed from the marginal distributions of cognitive and socio-emotional endowments. The second figure (bottom left) presents $E(Y_{\text{four-year degree}} - Y_{\text{some college}}|d^C)$ so that the socio-emotional factor is integrated out. The bars in this figure display, for a given decile of cognitive endowment, the fraction of individuals visiting the node leading to the educational decision involving graduating from a four-year college. The last figure (bottom right) presents $E(Y_{\text{four-year degree}} - Y_{\text{some college}}|d^{SE})$ and the fraction of individuals visiting the node leading to the educational decision involving graduating from a four-year college for a given decile of socio-emotional endowment.
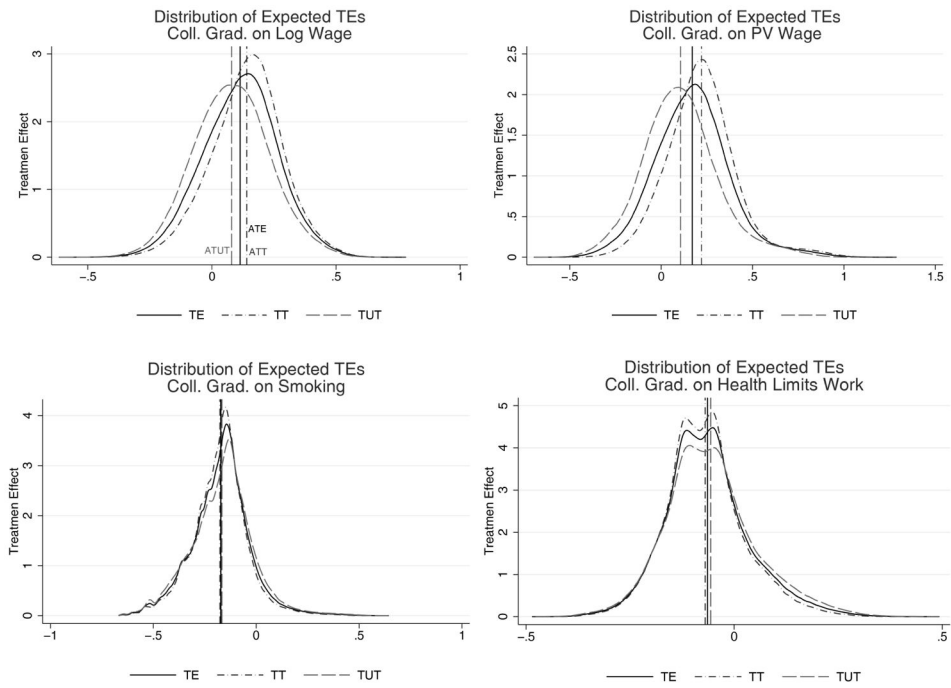
Figure 7.
Distributions of Expected Treatment Effects: College Graduation

Notes: Distributions of treatment effects including continuation values for those who reach the educational choice. The vertical lines represent the average treatment effects (ATE, ATT, and ATUT) for each of the distributions.
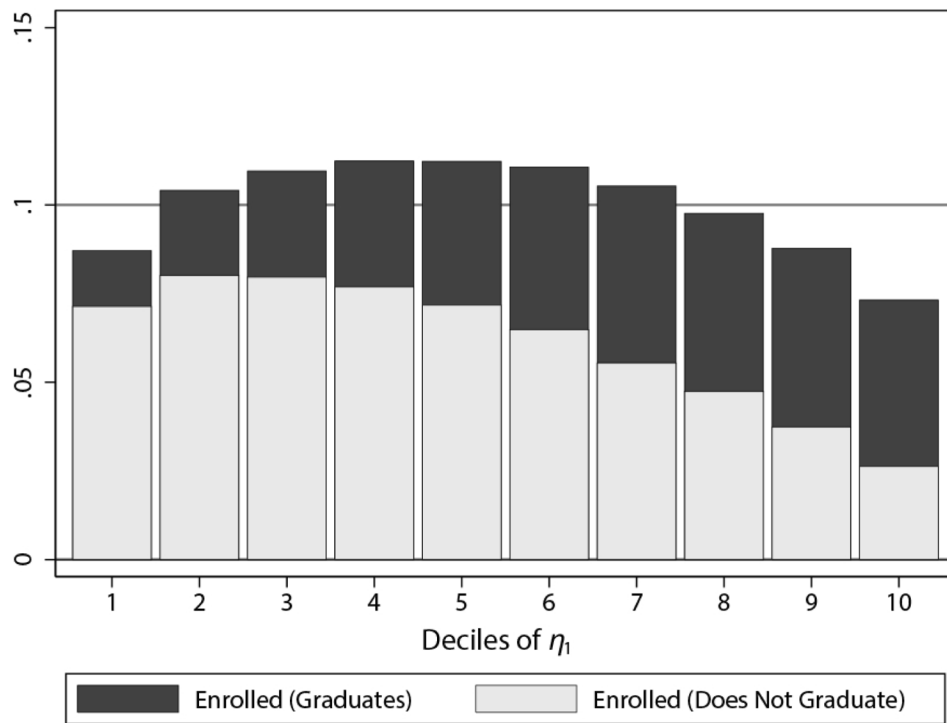
Figure 8.
PRTE: Who Is Induced to Switch?

Notes: The figure plots the proportion of individuals induced to switch from the policy that lay in each decile of η_2, where $\eta_1 = -(\boldsymbol{\theta}'\boldsymbol{\lambda}_1 - \nu_1)$. η_1 is the unobserved component of the educational choice model. The deciles are conditional on $Q_1 = 1$, so η_2 for individuals who reach the college enrollment decision. The bars are further decomposed into those that are induced to switch that then go on to earn four-year degrees and those that are induced to switch but do not go on to graduate.

Figure 9.
Policy Experiments
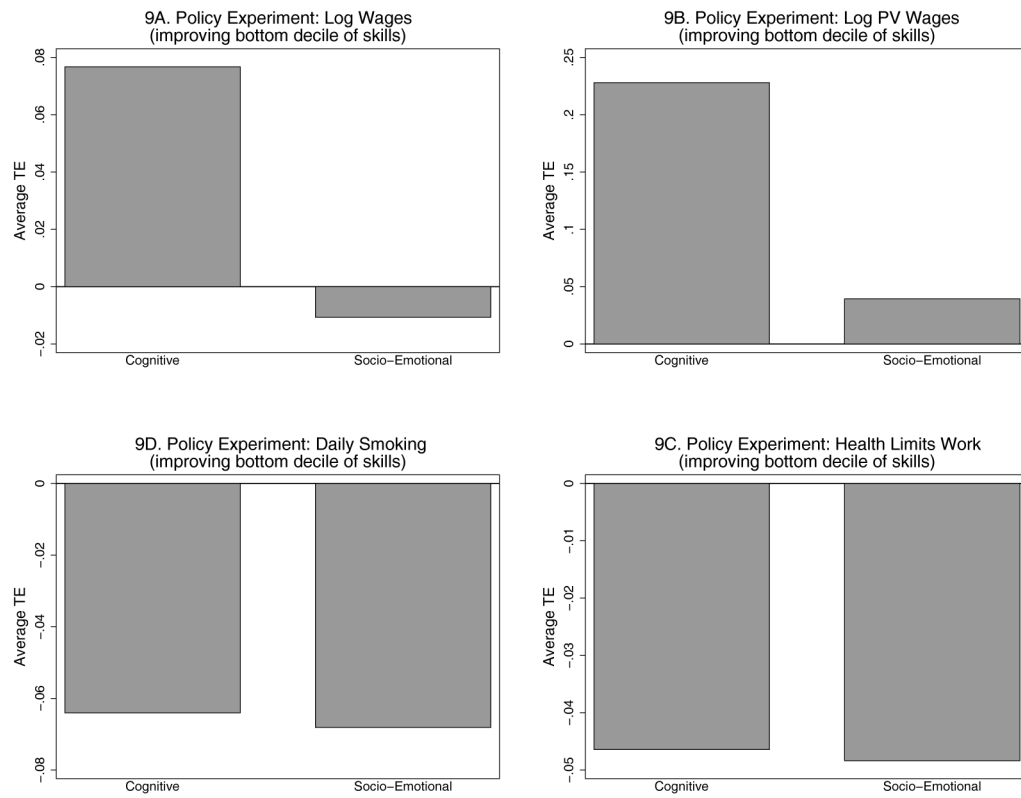*Notes*: This plot shows the average gains for those in the bottom deciles of cognitive ability (left) and socio-emotional ability (right), from an increase in the endowment.

Table 1

Control Variables and Instruments Used in the Analysis

Control Variables	Measurement Equations	Choice	Outcomes
Race	x	x	x
Broken Home	x	x	x
Number of Siblings	x	x	x
Parents' Education	x	x	x
Family Income (1979)	x	x	x
Region of Residence[a]	x	x	x
Urban Status[a]	x	x	x
Age[b]	x	x	x
Local Unemployment[c]			x
Local Long-Run Unemployment		x	
Instruments (Exclusion Restrictions)			
Local Unemployment at Age 17[d]		x	
Local Unemployment at Age 22[e]		x	
College Present in County 1977[f]		x	
Local College Tuition at Age 17[g]		x	
Local College Tuition at Age 22[h]		x	

Notes:

[a] Region and urban dummies are specific to the age that the measurement, educational choice, or outcome occurred.

[b] Age in 1979 is included as a cohort control. We also included individual cohort dummies which did not change the results.

[c] For economic outcomes, local unemployment at the time the outcome is measured.

[d] This is an instrument for choices at nodes 0 and 1. It represents opportunity costs at the time schooling decisions are made.

[e] This is an instrument for the choice at node 2.

[f] Presence of a four-year college in the county in 1977 is constructed from Kling (2001) and enters the choice to enroll and the choice to graduate from college.

[g] Local college tuition at age 17 only enters the college enrollment graduation decisions.

[h] Local college tuition at age 22 only enters the college completion equation. The measurement system includes the arithmetic reasoning, coding speed, paragraph comprehension, word knowledge, mathematical knowledge, and numerical operations sub-tests of the ASVAB, 9th grade GPA in math, English, science, and social studies, and early risky and reckless behavior. We assume ASVAB only loads on the cognitive factor. See Web Appendix Section A.2 for details.

Table 2

PRTE: Standard Deviation Decrease in Tuition

	PRTE	**Four-Year Degree**	**No Four-Year Degree**
Log Wages	0.125 (0.023)	0.143 (0.027)	0.114 (0.027)
PV Log Wages	0.129 (0.03)	0.138 (0.033)	0.123 (0.028)
Health Limits Work	−0.036 (0.022)	−0.025 (0.021)	−0.043 (0.023)
Smoking	−0.131 (0.029)	−0.166 (0.030)	−0.108 (0.030)

Notes: The table shows the policy-relevant treatment effect (PRTE) of reducing tuition for the first two years of college by a standard deviation (approx. $850 per annum). The PRTE is the average treatment effect of those induced to change educational choices as a result of the policy:

$PRTE_{p,\,p}^{k} := \quad E(Y^k(p\;) - Y^k(p) \quad X = x, Z = z, \theta = \bar{\theta}) dF_{X,\,Z,\,\theta}(x, z, \bar{\theta} \quad S(p) \quad S(p\;))$. Column 1 shows the overall PRTE.

Column 2 shows the PRTE for those induced to enroll by the policy who then go on to complete four-year college degrees. Column 3 shows the PRTE for individuals induced to enroll but who do not complete four-year degrees.

Table 3

Average Treatment Effects - Comparison of Estimates from Our Model to Those from Simpler Methods

HS Grad.	Linear Regression				Matching		Model
	OLS	OLS-P	OLS-F	OLS-FI	NNM(3)-F	PSM-F	ATE_j^{k*}
Wages	0.205	0.073	0.155	0.159	0.098	0.132	0.094
SE	(0.025)	(0.026)	(0.025)	(0.035)	(0.037)	(0.051)	(0.056)
PV-Wage	0.380	0.213	0.318	0.277	0.196	0.226	0.173
SE	(0.030)	(0.031)	(0.030)	(0.041)	(0.053)	(0.058)	(0.059)
Smoking	−0.327	−0.246	−0.281	−0.301	−0.260	−0.271	−0.263
SE	(0.028)	(0.029)	(0.027)	(0.028)	(0.058)	(0.060)	(0.056)
Health-Limits-Work	−0.178	−0.115	−0.151	−0.150	−0.048	−0.095	−0.108
SE	(0.022)	(0.024)	(0.023)	(0.033)	(0.029)	(0.036)	(0.042)
Coll. Enroll	OLS	OLS-P	OLS-F	OLS-FI	NNM(3)-F	PSM-F	ATE_j^{k}
Wages	0.223	0.121	0.186	0.190	0.177	0.207	0.134
SE	(0.023)	(0.024)	(0.024)	(0.023)	(0.029)	(0.031)	(0.025)
PV-Wage	0.221	0.109	0.176	0.171	0.188	0.226	0.137
SE	(0.027)	(0.029)	(0.028)	(0.027)	(0.030)	(0.032)	(0.029)
Smoking	−0.177	−0.138	−0.165	−0.170	−0.129	−0.144	−0.139
SE	(0.026)	(0.028)	(0.027)	(0.028)	(0.029)	(0.058)	(0.028)
Health-Limits-Work	−0.085	−0.037	−0.066	−0.057	−0.029	−0.042	−0.037
SE	(0.020)	(0.022)	(0.021)	(0.021)	(0.022)	(0.029)	(0.022)
Coll. Grad	OLS	OLS-P	OLS-F	OLS-FI	NNM(3)-F	PSM-F	ATE_j^{k}
Wages	0.210	0.146	0.184	0.185	0.173	0.143	0.114
SE	(0.032)	(0.034)	(0.033)	(0.035)	(0.041)	(0.051)	(0.037)
PV-Wage	0.243	0.163	0.208	0.228	0.191	0.269	0.171
SE	(0.037)	(0.040)	(0.038)	(0.037)	(0.039)	(0.042)	(0.040)

HS Grad.	Linear Regression				Matching		Model
	OLS	OLS-P	OLS-F	OLS-FI	NNM(3)-F	PSM-F	ATE_j^{k*}
Smoking	−0.209	−0.171	−0.195	−0.192	−0.132	−0.161	−0.172
SE	(0.032)	(0.035)	(0.033)	(0.035)	(0.039)	(0.039)	(0.043)
Health-Limits-Work	−0.085	−0.069	−0.078	−0.077	−0.048	−0.051	−0.064
SE	(0.024)	(0.026)	(0.025)	(0.026)	(0.026)	(0.027)	(0.031)

Notes: We estimate the ATE inclusive of continuation values for each outcome and and educational choice using a variety of methods. All models are estimated for populations that reach the node being analyzed ($Q_j = 1$), inclusive of those who go on to further schooling in order to make them comparable to the ATE from our model that includes continuation values (Equation (21)). All OLS models use the full set of controls listed in Table 1. "OLS" estimates a linear model using a schooling dummy (Q_{j+1}), and controls ($Y = Q_{j+1}bj + X'\beta + e$). "OLS-P" estimates a linear model using a schooling dummy, a vector of controls, and three measures of abilities arrayed in a vector A: summed ASVAB scores, GPA, and an indicator of risky behavior ($Y = Q_{j+1}bj + X'\beta + A'\alpha + e$). All models ending in "-F" are estimated using Bartlett factor scores (Bartlett, 1937, 1938) estimated using our measurement system, but using the built-in routine for estimating factor models in STATA via maximum likelihood, not accounting for schooling at the time of the test. "OLS-F" estimates the model $Y = Q_{j+1}bj + X'\beta + \hat{\theta}'\alpha + e$ where $\hat{\theta}$ are the Bartlett factor scores described above. "OLS-FI" is similar to "OLS-F" except that Q_{j+1} is interacted with the X and $\hat{\theta}$ allowing the coefficients on the controls and abilities to vary by education level. "NNM(3)-F" is the estimated treatment effect from nearest-neighbor matching with 3 neighbors. Neighbors are matched on their Bartlett cognitive factor, Bartlett non-cognitive factor, and an index constructed from their observed characteristics (Z) generating choices as described in Web Appendix A.18. "PSM-F" presents the estimated average treatment effect from propensity score matching where propensity scores are estimated using Bartlett cognitive factors, Bartlett non-cognitive factors, the full set of control variables, and the full set of node-specific instruments. "ATE_j^k" presents the estimated average treatment effect from the model presented in this paper (inclusive of continuation value), corresponding to Equation (13).